A Comprehensive Investigation of Flight and Duty Time Limitations and their Ability to Control Crew Fatigue

In commercial aviation, crew schedules are regulated by laws that define maximum duty time limits, flight time limits or minimum rest periods and other constraints. These rules and limits, collectively referred to as Flight Time Limitations or FTLs, were intended as a simple method of limiting and accounting for fatigue. The differences between FTL regulations around the world are significant and can affect, to varying extent, crew productivity as well as the levels of crew alertness.

This report represents the most comprehensive quantitative analysis on FTLs performed to-date, measuring seven different FTL regulations and the effects they have on crew alertness and crew scheduling efficiency. Each FTL has been applied to more than 300 large airline fleets world-wide, using a crew schedule optimization system also used by many airlines globally.

The results of this study show that FTLs have limited effectiveness in reducing crew fatigue. Even the latest proposed changes to the FTLs by FAA and EASA provide only marginal improvement. The results show that crew fatigue in two-pilot operations is mainly dependant on the nature of the underlying flight schedule and that the newly proposed FTLs do not effectively reduce crew fatigue.

This paper proposes a shift in focus for the aviation industry to place a greater emphasis on fatigue models as part of crew scheduling to further reduce crew fatigue risk.

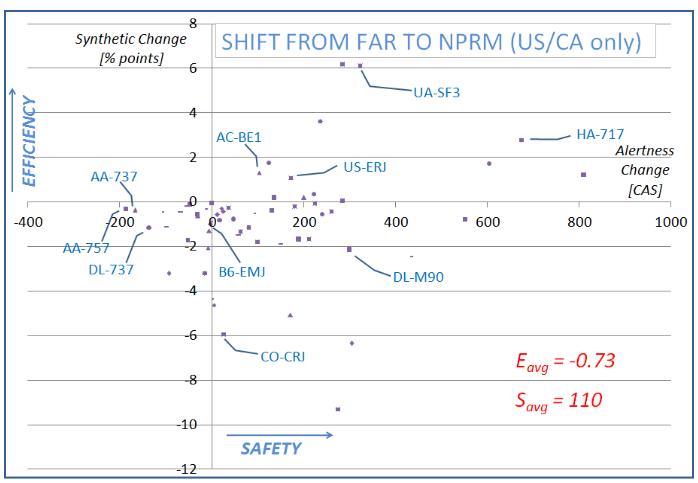

The quantified effect on efficiency and safety for two-pilot operation should US and Canada airline fleets be planned with NPRM instead of current FAR rules.

ÅKE OLBERT ake.olbert@jeppesen.com
DAVID HELLERSTRÖM david.hellerstrom@jeppesen.com
TOMAS KLEMETS tomas.klemets@jeppesen.com

INTRODUCTION

Fatigue in crew schedules is difficult to adequately measure and capture using rules that only limit working hours and safeguard rest. In June 2010 Jeppesen published an article called "The Best Rest" [1] in Aerosafety World[1]. The article presented a large-scale comparison of the ability of different regulatory approaches in enabling efficiency and capturing crew fatigue. The comparison was based on three large airline networks.

The results presented in that article showed that the US FAR rules stood out from the other rule sets as being more efficient, while also being the least protective against fatiguing patterns. These results were novel in that they quantified efficiency and fatigue risk in actual production systems used daily by many large airlines, and also proposed a methodology by which rule sets can be improved.

However, the results also pointed to a rather large variation in fatigue that originated with the flight schedule itself - i.e., the actual fatigue levels came out quite differently for different networks using the exact same rules. The questions that followed thus became:

a) How large is this variance?
b) Can one set of regulatory rules, in their current form and applied by all operators, contribute to a reduction in crew fatigue in any major way?

Could it be that crew fatigue is largely pre-determined by the flight schedules, and that regulatory rules principally limit efficiency? This question is also motivated from a scientific point of view. *Time of day,* a major contributor to crew fatigue, is primarily set when designing the flight schedule itself.

[1] The Flight Safety Foundation Journal

THIS STUDY

This study involves using actual production systems to construct crew pairings (efficient sequences of flights) for over 300 different two-pilot operation fleets world-wide. The vast amount of data generated makes it possible to look at the average difference between any two or more FTLs across all 300 fleets. This delivers a comprehensive picture of the effectiveness of these regulations.

Flight Time Limitations Studied

The study was carried out using seven regulatory FTLs for each airline fleet;

- **FAR** (USA) Part 121 rules
- "**NPRM**" The proposed updates to FAR as defined in the FAA NPRM from 10 Sep 2010
- **EU-Ops** The current rules containing subpart Q
- **"NPA"** The proposed updates to EU-Ops as defined in the NPA 2010-14A as of 20 March 2011
- **CAP 371** (UK/Commonwealth)
- **CCAR** (China) CCAR 121 Rev 3
- **DGCA** (India)

Preconditions and Constraints

The flight schedule data used was taken from the public OAG file[2] for the period 02-08 MAY 2011. All airline fleets with more than 200 weekly flights were included in the analysis with only a few exceptions[3].

Pairing construction is the defining step for crew efficiency as well as fatigue. The resulting product of this step of the crew

[2] See gov.oag.com

[3] To save computer time the large 737 fleets of Ryanair and Southwest Airlines were excluded.

management process can only deteriorate further in the following rostering step and the later (often manual) roster maintenance step. Therefore, the actual efficiency (if using only FTLs) will be worse than reported upon here.

Most operators are bound by labor agreements in addition to FTLs, which reduce efficiency further but often increase alertness levels. The results should therefore be read as a pure comparison of the effectiveness of different regulatory approaches in limiting crew fatigue and allowing for efficiency and operational flexibility. **The results here presented do not reflect these operators'/fleets' <u>actual</u> fatigue levels.**

Another item to be aware of is flight schedule adaptation. The comparison is made using the same flight schedule for planning under different FTLs. In a long run, a change in FTL will influence the flight schedule. If, for example, the US market goes with the NPRM rules increasing rest from 8h to 9h on rest location, outstation turns will gradually be adjusted to avoid crew inefficiency.

For the pairing construction step, the following was applied:

- Crew bases were selected automatically by investigating the network and selecting the "natural hubs" for crew
- The planning was not bound by any base constraints on crew but instead assumed an optimal distribution of crew over the bases
- Positioning of crew (so-called deadheading) was allowed only on the same carrier
- Aircraft rotations were created using an algorithm based on the First-In- First-Out principle

In addition to the regulatory rules, a number of typical scheduling rules were used to ensure a degree of operational feasibility and robustness of the solutions. These were:

a) Briefing time: 0h45m
b) Debriefing time: 0h15m
c) Max 2 deadheads per duty
d) No "middle-deadheads"[4] allowed
e) Max 2 aircraft changes in duty
f) Min connection time for aircraft change: 0h45m
g) Max 5 duty days in a trip

The objective function for the planning was configured for a typical operator striving for efficiency and robust solutions that also minimize costs for layovers, per diem allowances and deadheading. Due to the different payment structures for the majority of US and Canadian carriers, compared to the rest of the world, runs for those fleets were configured to minimize additional *credit time,* i.e., pay time that is higher than flight time. The amount of credit time reflects the efficiency for a North American carrier, while for operators in other regions, active block hours per calendar day were maximized.

All plans were produced using the Jeppesen Crew Pairing optimizer [11], a tool used by many of the largest operators for their actual production planning.

Sleep opportunities for the crew commenced 2h after arrival and lasted until 2h before the next-coming departure. At outstations, this was reduced to 1h 45m in both cases. (Within these sleep opportunities, the fatigue model will predict sleep/wake, otherwise wakefulness is assumed.) No in-flight sleep was used as

[4] Deadheads are not allowed other than first or last in a duty day.

only two-pilot operations were investigated.

Comparison metrics

Several different metrics were used for comparison of the solutions.

Efficiency. For representing efficiency the choice was made to divide the airline fleets in two groups for simplicity;

a) **Operations based in the US or Canada** For this group the difference between credit and flight time was measured as percentage of flight time – here referred to as *synthetic*. Most fleets in this region have between 2-8% synthetic in their operation.

b) **Non US/Canada operations** For this group, which often operates on monthly salary with overtime pay, measurement of productivity was expressed in active block hours (flying hours) per calendar day.

The efficiency of a fleet is restricted not only by FTLs, but also by the flight schedule not providing enough efficient connections between flights. There are also imbalances of crew across bases compared to optimality, leading to excessive positioning of crew (dead-heading).

Safety. A low level of alertness on a flight is associated with higher risk. Although there are many different factors affecting human performance and also on the more complex interaction with mission difficulty[5], in this study, we used a simplified view of equating safety with the predicted level of alertness. The alertness properties of a solution are hard to map to a single descriptive value or statistical measure, but it is clearly the so-called tail of the alertness distribution of all of the flights (see Figure 1) where the risk is the greatest. The risk will grow exponentially as the alertness approaches zero – corresponding to a state where an average individual is "fighting sleep".

We have chosen to re-use the metrics from the Best Rest work [1] focusing on the KPI called PA5 (see Figure 1). PA5 is the average predicted level of alertness of the worst 5% flights in the overall solution for a fleet measured in CAS points (the Common Alertness Scale) ranging from 0 to 10,000.

Crew was assumed to be of intermediate diurnal[6] type with a habitual sleep length of 8h.

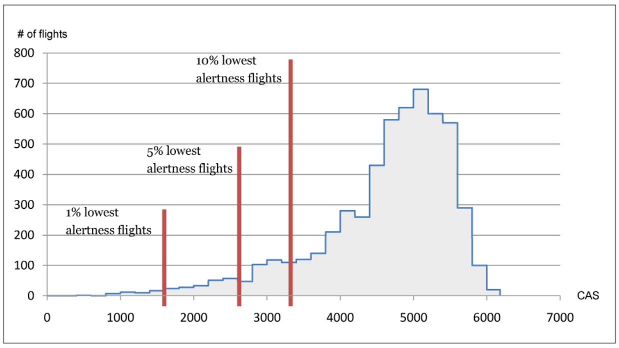

Figure 1. *An example distribution of the worst predicted alertness in a crew schedule comprised of 5,518 flights. The Y-axis depicts the number of flights in each alertness interval (bins being 200 CAS-points wide). The X-axis represents the lowest predicted alertness experienced on the flight, CAS-scale. Dividers for the 1, 5 and 10% lowest alertness flights are marked with vertical lines. PA5 thus constitutes the average of all flights below (to the left of) the mid divider.*

The prediction of alertness was made using the Boeing Alertness Model (BAM) in version 1.1.6. BAM is based on validated, openly published science[7], and

[5] The external threat factors of the flight like weather, runway length, surrounding terrain etc. See [2].

[6] The individual inclination towards morningness or eveningness

[7] Please see the work of Akerstedt et.al.:[3-8].

is used by many airlines and crew worldwide for evaluating crew schedules. BAM has also been used operationally by airlines since June 2011 for creating crew schedules. It should be noted that no commercially available fatigue model implementation has published peer-reviewed validation towards airline data that has not been part of improving the model itself. For this study, the BAM output on the CAS scale 0-10,000 was trusted to accurately represent the alertness perspective for an average individual.

The metric used for alertness has been the lowest level predicted during active block time for each flight. In total, 385,000 flights (766,000 block hours) were evaluated in this study.

ANALYSIS

After conducting more than 2,100 production runs with the Jeppesen optimizer, there are several possible ways of analyzing the results. Here we will focus on two perspectives:

a) Investigating the *average effect* on efficiency and safety for all fleets when/if they were to change over from one FTL to another. For example, what would the switch from EU-Ops to NPA result in for a fleet?

b) Looking into the *absolute numbers* on efficiency and safety for each operator using different regulatory rules. A few fleets of varying properties are illustrated due to space limitations.

Average effects

Starting with the US and Canada, what would the shift from FAR to the proposed NPRM mean to the industry?

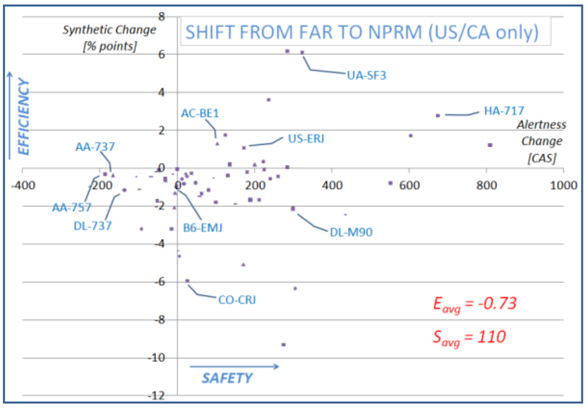

Figure 2. *The effect on efficiency (synthetic) and safety (PA5) when/if changing over from using FAR to NPRM. Note that the values represents relative difference between the rule – not the absolute levels operated on.*

The above scatter plot (Figure 2) illustrates the effect on each fleet when switching from planning purely on FAR rules to using NPRM. The vertical axis is the effect on efficiency expressed in percentage points on synthetic. The horizontal axis is the effect on safety as represented by points on the CAS scale for the 5% worst flights (PA5). The FAR reference is thereby placed in origin of the chart (0, 0) for each fleet.

This plot supports the results found in the Best Rest work with the NPRM rules constituting an improvement in terms of safety. The average effect for these fleets, weighted for the number of flights, increases PA5 with 110 CAS points. On a scale going to 10,000, this may sound insignificant, but it should be pointed out again that this is the average of the 5% worst flights and thus a significant improvement.

However, the weighted average effect on efficiency is less fortunate. On average, the industry will pay an additional 0.7% credit time for the above increase in safety. The network tested in the Best Rest work (Northwest A320 fleet) indicated a small gain in efficiency, but in this analysis, we can see that on average it is actually a loss.

Note that some fleets go against the trend and could become even less efficient and less safe - like the Delta 737 fleet. Other fleets, like the United SF3 fleet will gain up to 6% in credit and 350 CAS points.

What would then be the effect if this region, instead of adopting NPRM, applied the Chinese CCAR rules?

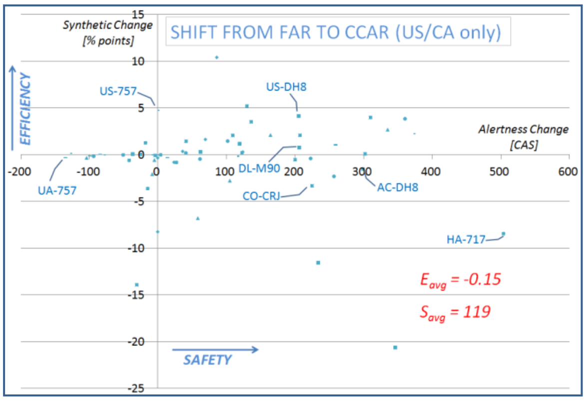

Figure 3. *The effect on efficiency (synthetic) and safety (PA5) when/if changing over from using FAR to CCAR.*

As seen in Figure 3 above, the CCARs would actually be more beneficial for this region than the NPRM, when looking only at the application of FTLs. The cost is a mere additional 0.15% credit time for achieving an increase of 119 points CAS in weighted average over all the networks measured.

Next, let's look at the European rules. Will the EASA NPA improve safety significantly in a cost-efficient way?

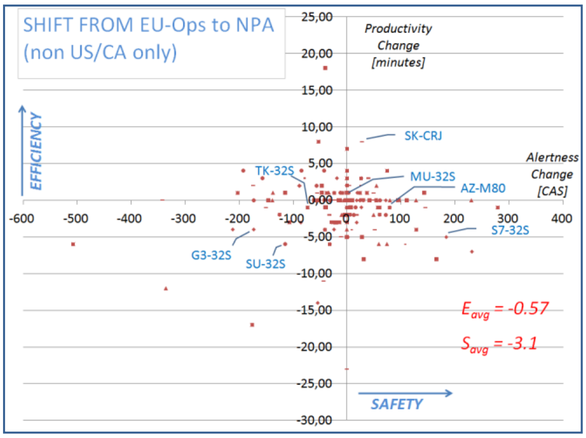

Figure 4. *The effect on efficiency and safety (PA5) when/if changing over from using EU-Ops to NPA.*

In Figure 4, the vertical axis uses the non US/Canada perspective on efficiency, which is productivity measured in block hours per calendar day. Note again that the axis illustrates the relative change for a fleet in minutes, not the absolute numbers operated on.

The average effect of the change is negligible – both in efficiency and in safety as represented by PA5. On average, these operators will lose half a minute of flying each day for virtually no gain in PA5.

Of course, comparing EU-Ops and FAR is also possible. A shift from EU-Ops to FAR for all the non US/Canada operators is illustrated below in Figure 5.

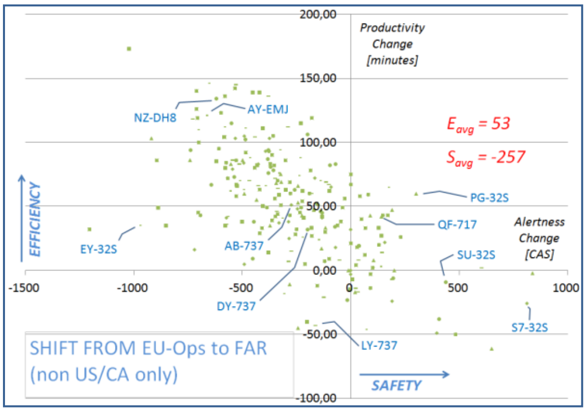

Figure 5. *The effect on efficiency and safety (PA5) when/if changing over from using EU-Ops to FAR.*

(Note that many of these fleets are not bound by EU-Ops today as they are based outside of the EU.)

This plot supports the results found in the Best Rest work with FAR allowing for higher efficiency than EU-Ops. The weighted average effect (by the number of flights) in each problem is 53 more minutes of flying per calendar day. This is a dramatic increase in efficiency.

On the horizontal axis, we also see that the shift over to FAR rules will generally affect safety negatively. The weighted average effect is –257 points on the CAS scale.

The data used for the Best Rest work implied that the NPRM would improve the FAR rules to be "on par" with the EU-Ops in ability to capture fatigue. From the new and more extensive data, we can now see that on average the gap will only be reduced by about half – a 110 point improvement of a 257 point gap.

Again, some fleets go against the trend in Figure 5 and become less efficient and more safe, such as the Aeroflot (SU) 32S fleet and Siberian (S7) 32S fleet.

The work conducted in evaluating seven FTLs lends itself to comparing an additional imaginary 17 switch-overs between rule sets and can be filtered out and applied to any region, country or set of operators. Subsequent publications to this white paper will be based on the data generated for this study.

Absolute numbers

The information produced allows for looking into each fleet individually and for studying the absolute numbers on efficiency and safety under each FTL. A few examples are visualized below to show the variance in results for individual fleets. Note that each fleet has the data points for the seven FTLs connected by a line.

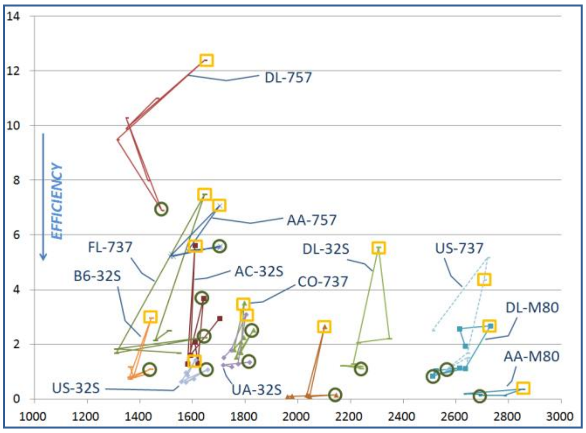

Figure 6. *US/Canada fleets synthetic (Y axis) versus CAS points (X axis) for seven rule sets each. Data points for FAR are circled, DGCA within a square.*

Looking at Figure 6 above, it is clear that data points (X, Y) are fairly close together for a fleet. The difference between two fleets is generally bigger than the variance in results between FTLs within one fleet.

FAR stands out in this data as the FTL allowing for the most efficiency. The DGCA rules drives most on the safety KPI but also create significant cost increase.

Figure 7 shows a similar visualization of a number of Asian fleets with productivity on the vertical axis.

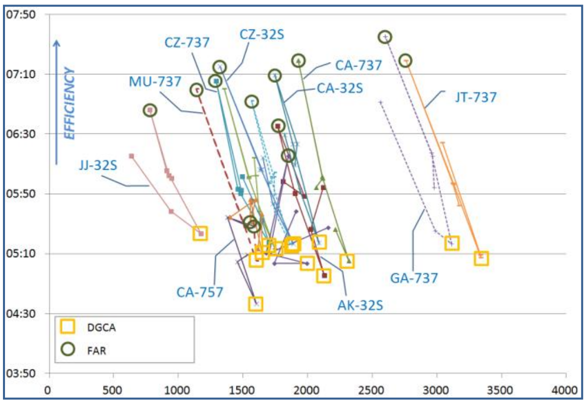

Figure 7. *A number of Asian fleets and the efficiency/safety perspective illustrated for the seven FTLs on their flight schedule. (FAR circled, and DGCA marked with a square).*

Looking at the CAS scale, the data points for one fleet, seems to fit together better than the data points for one FTL. Regarding efficiency however, FAR stands out as the more efficient rule set, and DGCA as the least efficient one.

Finally, a few fleets from areas other than Asia and US/Canada are shown in Figure 8 below.

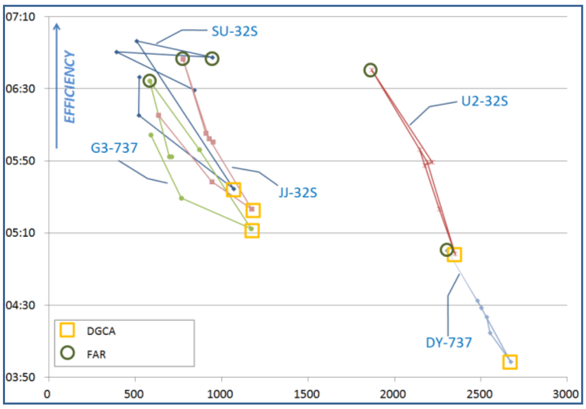

Figure 8. *Five fleets outside Asia and US/Canada and their, by FTLs allowed, efficiency and safety KPIs. (FAR circled, DGCA marked with a square.)*

Again, this graph emphasizes the con- clusion that the flight schedule is of higher importance to alertness levels than the FTL used. Norwegian (DY) and Easyjet (U2) are predominantly daytime operators and will come out significantly higher in CAS than Aeroflot (SU), GOL (G3) and TAM (JJ). Due to a small operation and fewer productive connections between flights, Norwegian 737 ends up lower in efficiency than the others.

The overall conclusion is that the current (and proposed) FTLs predominantly limit efficiency while only marginally reducing the fatigue levels determined by their business model.

The FTL Misalignment

After reviewing these high-level results for full planning solutions comprising crew pairings (flight sequences) for large numbers of crew, it could be of interest to assess the reasons behind this mis-alignment and lack of precision of the FTLs as shown by the large variance in the data above.

If the FTLs are built on scientific findings, at least after being revised (the NPA and NPRM), why they are so misaligned with fatigue models containing the very same science? The answer is two-fold: over-simplification and isolation. The FTLs do not contain a consistent representation of science, but several simplified rules applied in isolation.

Fatigue models, however, take into account the complex interaction and cumulative effects of several different causes of fatigue. For example, these include time of day, circadian disruption, time since awake, time on task and prior sleep deprivation. A fatigue model provides a continuous metric, which is a very important advantage to the binary (good/bad) output of FTLs which are applied isolated on portions of the crew schedule.

As a demonstration, let's look at a very simple and real example of FTL over-simplification.

A two-pilot operation throughout the night, starting at 8pm and lasting for almost 11 hours, will be quite challenging for the crew – especially around 4am in the "low" of the circadian phase. Since regulators would find it difficult to prohibit this type of operation altogether, they have instead implemented protection around such flights, such as instituting a preceding rest requirement of 12h, for example. What could then happen in real operation when this simple rule is applied?

Figure 9 is an example of the actual consequence for one operator.

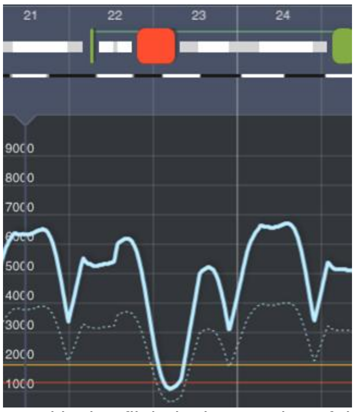

Figure 9. *A positioning flight in the morning of the 22nd inflicting on the night before a difficult flight, causing sleep deprivation. Modeling done using BAM in CrewAlert for iPhone.*

Departure in home-base time for the flight Gothenburg to Phuket is 7:40pm with an arrival at 6:30am, after flying for 10h 50m. The operator, with crew based in Stockholm, will need to position (deadhead) the crew into Gothenburg for staffing this flight since Gothenburg is not a crew base. Due to the regulatory 12h rest rule applied here, this deadhead flight takes place between 6:10 and 7:10 on the preceding morning to ensure a full 12 hours of "rest" in Gothenburg. Because of the regulatory rule, this operator will force the crew to rise at 04:00-04:30 the morning before a difficult flight through the night, which is exactly the opposite of what is wise for safeguarding alertness on the flight. Deadheading a few hours later would have been better for safety, for the pilots quality of life, and for efficiency as well.

The consequences for fatigue from new rules defined only in hours and minutes are often difficult to forecast due to the overall complexity in the crew management process. In addition to the example above, plenty of other situations occur where rules back-fire in a similar way.

Also with respect to efficiency, rules often fail. There are many flights flying through daytime conditions that are unnecessarily staffed with surplus pilots (in comparison). This occurs while only somewhat shorter night-duty flights that are in greater need of extra crew operate with only two pilots because they are shorter than the block-hour limit for augmentation.

Other examples include the over-simplification made on time zone travel that include rules such as “If crossing more than x time zones, a subsequent rest of y days are required.” The rules are well intended but, in reality, being acclimatized west from base makes a person more fit to fly late in the evening back to base time. So, the length of stay in the new “theater” also matters, as do the departure times and sector lengths involved.

FTLs are actually fatigue models, but they are binary in their approach - categorizing schedules as either perfectly safe or perfectly unsafe. They constitute a collection of many over-simplified rules that, when applied individually in isolation or, especially when combined and followed in reality, deviate from science and become superficial fixes. The precision is just not there.

This misalignment presents a large improvement opportunity for the airline industry. By reallocating resources and expenditure to where it matters most, we can build better crew schedules with lower risk without increased costs. In order to achieve this, there is need for a slightly different approach recognizing that rules will not resolve the situation. FTLs that regulate work hours are good at limiting working hours – but they will never be really good at reducing fatigue risk.

CONCLUSIONS

The conclusions based on the results above include:

- First and foremost, fatigue is significantly linked to the business model of the operator. A large portion of night time operation directly results in more fatiguing patterns.
- FTLs in their current form do not limit fatigue effectively.
- Current FTLs have a more significant effect on efficiency than on fatigue risk.
- FAR allows for the highest efficiency, but is also the least protective from fatigue.
- DGCA is the most protective rule set for fatigue risk but is generally most restrictive on efficiency.

So, could we not fix the problem by updating the FTLs to become more aligned with science?

A PERFECT FTL?

Let’s assume for a moment that a regulator does create a set of FTLs that is perfectly aligned with science and then prohibits flights below a certain alertness level. If this were to happen, the alertness level could not be the same for all operators without drastic consequences to their business models - at least not if this limit should make any <u>real</u> difference on the overall crew fatigue.

To illustrate this, let's look at alertness distributions for three operators with quite different business models in Figure 10.

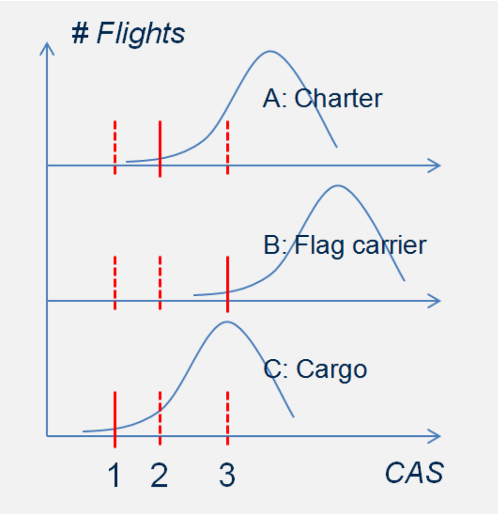

Figure 10. *Three business models and their alertness distribution over their operation.*

Due to individual variation, fatigue risk will exist in a significant part of the distributions' left-hand tails, and efforts should also be spent in managing this risk for operator B above.

If the regulator chooses level 1 as the lowest acceptable limit, operator A and B can comply without doing anything to reduce risk further. If level 3 is chosen, operator A and C will seriously need to change their business models.

From the perspective of protecting the public and further reducing <u>overall</u> risk, the regulator would need to implement one FTL per business model to secure a lower boundary for each.

To some extent, this is what now takes place, even if it is not the regulator doing the actual work. Under an FRMS, each operator can soon stray from FTLs and customize the rules carefully to their operation while still showing an equivalent level of safety. Thereby, each operator will need to untangle and align their FTL with science if the safety potential is to be released without sacrificing efficiency. This will not only be very expensive, but can also create further fragmentation and result in one FTL per operator. It will also be difficult to carry over experiences between operators.

FTLs, even if perfectly aligned with science, will still only categorize flights in legal/illegal, which leads to the last important conclusion:

- When acting within a set of FTLs there is a need for <u>direction</u> (incentives for planning) to drive the crew schedules away from fatigue.

A SHIFT IN FOCUS

The FTL updates from the FAA and EASA are somewhat misdirected with a one-size-fits-all approach. The result will either be rules that everyone can apply without any major consequences to their business models but that make little or no difference for fatigue, or rules that have tremendous impact on the business models for some operators and only a slightly better effect on reducing fatigue.

What other ways forward exist for ICAO and other regulators to address fatigue if (one set of) rules are not adequate? The authors' opinion is that much more effort should be spent on fatigue models. This should include more than what was expressed in the recently published FRMS manual and implementation guide [9] [10].

Fatigue models have the advantage of providing a metric and can be continuously improved through operational data. Models can also be used by operators for affecting flights within rule limits, shifting the tail of the distribution upwards in predicted alertness.

The introduction of mechanisms needed for efficient model development should be driven from the very top (ICAO/IATA) in order to control costs and maximize the effect.

It would be highly desirable if ICAO and the major regulators could consider the following:

a. Require operators to continuously collect and share data on crew fatigue in a standardized fashion meeting the needs of sleep and performance scientists.
b. Drive research on collected data with scientific entities to improve, and eventually endorse or recommend, the mathematical fatigue models to be used.
c. Leave the FTLs as-is for the time being. The work to align these with science has not been performed in a quantifiable way and the loopholes[8] and "adverse consequences" that could result from the changes have not been simulated.
d. Use fatigue models to improve the regulatory limits on flight and duty time best possible.
e. Require operators to identify their 5% worst flights per fleet and, over time, rank and track the development of a KPI reflecting the risk. Require the operator to report the basis for this risk assessment. Certainly, fatigue is one part, but mission difficulty may be present as well. Also require operators to report the measures taken to control and reduce the overall operational risk for these flights.

Using existing fatigue models and tools, operators can already consider:

a. Using a pre-flight risk assessment containing predicted alertness as a tool for reducing fatigue in planning as well as re-scheduling flights close to the day of operation. Taking crew off a flight due to a missing minute of minimum rest is no longer to be seen as a tool for addressing safety; placing a more suitable crew on the flight that reduces overall operational risk is.
b. Use the pre-flight risk assessment during pairing and roster construction to counter-balance the drive for efficiency, fulfilling crew preferences and reducing cost. This planning step in the process has the biggest impact on operational risk – apart from the choice of business model.
c. Monitor the 5% worst flights and take action to reduce risk by changing aircraft type/crewing levels, providing better hotel/hotel locations, increasing/reducing flight frequency, increasing/reducing layover time, providing awareness and training to the crew on possible risk reduction mitigations.
d. Track the risk KPI from month to month for each fleet and install control limits.
e. Evaluate changes to the scheduling rules/incentives using quantifiable metrics. Perform what-if analyses on changes made.

[8] Safety weaknesses in the FTL formulations that could be exploited in the crew management process when driving for efficiency.

This, in combination with training for crew and involved parties, would make an actual difference for crew fatigue.

FURTHER ANALYSIS AND WORK

The material built up in this study lends itself to extensive further analysis. It would be possible, for example, to look in detail at the worst patterns produced or "cross load" rule sets to investigate which patterns the NPA would prohibit but the NPRM allows for. By doing this, some of the differences could be understood on a more detailed level and the "rough edges" in the rules could be polished out.

It would be desirable to also measure all of the produced plans with fatigue models other than BAM. Several of the commercially available fatigue models are in the process of developing functionality and needed performance that is similar to BAM in enabling live connections to scheduling systems which should make this possible in the coming years.

In this investigation, we have only been looking at the defining step – the crew pairing construction. The FTLs also contain rules only affecting the rostering step of the process. Additional runs could be made in an automated large-scale rostering study that also measures the effect on rosters.

Furthermore, it will be of interest to compare the FTLs with the performance of a "pure" FRM approach, as done in the Best Rest work [1]. By setting a minimum allowed level of alertness, it would then be possible to measure the maximum efficiency advantage to FTLs maintaining safety; or the other way around would include measuring the safety advantage while maintaining efficiency.

As previously stated, this investigation was limited to re-plan, two-pilot operations for the largest networks. Work has also been done by Jeppesen to evaluate augmented operations that could be scaled up to the same type of comparison for those planning problems as well.

With considerable effort, developed capabilities could support the use of the methodology described in the Best Rest work to iterate forward and stress-test a best-fit rule set that performs well over all of these fleets. Such a rule set could be optimized for allowing for as much efficiency as possible while providing the best possible protection from fatigue for a certain business model.

IN CLOSING

Even if the situation is not acute (there are not that many fatigue-related incidents or accidents occurring each year), we should drive this issue forward with urgency to further improve flight safety. Crew fatigue is one of the bigger risk components remaining in a very safe industry. It is predominantly the responsibly acting pilots, and the fact that they are also at risk, that we have to thank for the current safety levels – not the FTLs.

REFERENCES

1. Hellerström, D, Romig, E, Eriksson, H and Klemets, T. *The Best Rest, Aerosafety World June 2010, 40-45. (The original white paper is found via www.jeppesen.com/frm)*

2. Romig, Emma and Klemets, Tomas. Fatigue Risk Management in Flight Crew Scheduling. *Aviation,Space, and Environmental Medicine, Volume 80, Number 12, December 2009 , pp. 1073-1074(2).*

3. Åkerstedt, T and Folkard, S. The Three-Process Model of Alertness and its Extension to Performance, Sleep Latency, and Sleep Length. *Chronobiology International 14(2), 115–123, 1997.*

4. Åkerstedt, T, Folkard, S and Portin, C. Predictions from the three-process model of alertness. *Aviation, Space, and Environmental Medicine 2004; 75(3, Suppl.):A75-83.*

5. Åkerstedt, T and Folkard, S. Predicting duration of sleep from the three process model of regulation of alertness. *Occupational and Environmental Medicine 1996 February; 53(2): 136–141.*

6. Akerstedt, T., Axelsson, J. and Kecklund, G. *Individual validation of model predictions of sleepiness and sleep hours. Somnologie, 2007, 11: 169-74.*

7. Akerstedt, T., Connor, J., Gray, A. and Kecklund, G. *Predicting road crashes from a mathematical model of alertness regulation-The Sleep/Wake Predictor. Accid. Analys. Prevent., 2008a, 40: 1480-5.*

8. Akerstedt, T., Ingre, M., Kecklund, G., Folkard, S. and Axelsson, J. *Accounting for partial sleep deprivation and cumulative sleepiness in the three-process model of alertness regulation. Chronobiol. Int., 2008b, 25: 309-19.*

9. ICAO. *Fatigue Risk Management Systems, Manual for Regulators, 2011 Edition*

10. ICAO. *Fatigue Risk Management Systems, Implementation Guide for Operators, July 2011*

11. www.jeppesen.com/crew