

How do the models work?

All versions of the Fatigue Risk Index(FRI) operate based on understanding and analysing duties of shift workers. Between the release of the first version of the Fatigue Index (FI) and the creation of FRI, there was an increase in information concerning trends in risk related to shift work.

An extensive review of this information was carried out, and this enabled an index to be constructed that was entirely related to risk, rather than to fatigue and performance. The new index, version 2.0, therefore consisted of two separate individual indices, one related to fatigue (the 'Fatigue Index') and one related to risk (the 'Risk Index'). Whilst the two indices are similar in many respects, they diverge in others. The main differences are due to the different time of day effect: the peak in risk occurs close to midnight, whereas the peak in fatigue tends to occur some five hours later, in the early morning.

In the spreadsheet that was produced for version 2.0 to incorporate the new Index, both the risk and the fatigue indices were expressed in terms of three individual components:

1. A cumulative component

This relates to the way in which individual duty periods or shifts are put together to form a complete schedule. The cumulative component associated with a particular shift depends on the pattern of work immediately preceding that shift.

2. A component associated with duty timing, i.e. the effect of start time, shift length and the time of day throughout the shift.

3. A job type / breaks component

This relates to the content of the shift, in terms of the activity being undertaken and the provision of breaks during the shift.

FRI⁺ and FRI^{PRO} contain similar algorithms with the same three components included in them along with more features. FRI^{PRO} has a change in architecture too to achieve consistency across all FRMSc models.

For version 2.0, a review was undertaken to determine whether a separate Fatigue Index was required to cater for permanent night workers. It was concluded that the current index could be applied directly to the majority of shift workers who, based on the phase of their melatonin rhythm, showed little adaptation to permanent night work. However, the index would need to be adapted for the 30% of workers who showed at least some adjustment.

The output of FRI

The model calculates a single predicted fatigue score for the whole of the duty assuming the shift worker arrives at work at the first duty of their schedule, fully rested. The cumulative fatigue of the duty is calculated and applied to the subsequent duties to present the user with a fatigue score for each duty throughout the schedule based on the user-defined schedule settings that describe commuting time, work breaks, and workload.

Any schedule or roster of schedules can be analysed. This includes schedules that have already been worked or schedules that are about to be worked.

The Fatigue and Risk Index consists of two separate indices, one of which is related to fatigue, and the other is related to risk. What is meant by "Fatigue" and "Risk" is explained by virtue of the scales used. This is explained below.

The outputs from the two indices are on two different scales, and both differ from the scale used in the previous version. The Fatigue Index is now expressed in terms of the average probability, multiplied by 100, of a high score (specifically a value of eight or nine on the Karolinska Sleepiness Scale (KSS)), and therefore takes a value between zero and 100. The KSS is a nine-point scale ranging from one (extremely alert) to nine (extremely sleepy – fighting sleep).

The output from the Risk Index represents the relative risk of the occurrence of an incident on a particular shift. As with the Fatigue Index, the risk is averaged over the entire shift. A level of one represents the average risk on a typical two-day, two-night, four-off schedule, involving 12-hour shifts starting at 08:00 and 20:00. A value of two represents a doubling of risk.

It should be noted that there are some large differences in the output from the two indices, and a shift with a high value on one index is not always assigned a high value on the other. This is an inevitable consequence of the different information from which the two indices have been constructed and, in particular, of the differential effect of time of day. Whereas both fatigue and risk are highest on the night shift, the risk of an incident occurring on the afternoon shift is higher than on the morning shift. This contrasts with fatigue which tends to be higher on a morning than on an afternoon shift.

In FRI^{PRO}, seven scales are offered including the more popular Samn-Perelli fatigue scale that, unlike the KSS scale, considers workload. The KSS scale is a sleepiness scale only. Other scales include some scales used in the research and construction of the models: vigilance degradation, response time degradation, and percentage of missed responses.

Both the Fatigue and the Risk Index are constructed from three separate components, namely:

1. A cumulative component

This relates to the way in which individual duty periods or shifts are put together to form a complete schedule. The cumulative component associated with a particular shift depends on the pattern of work immediately preceding that shift.

2. A component associated with duty timing, i.e., the effect of start time, shift length, and the time of day throughout the shift.

3. **A job type / breaks component**

This relates to the content of the shift, in terms of the activity being undertaken and the provision of breaks during the shift.

The output of the model is a database of fatigue scores that are used to compile a picture of fatigue throughout the schedule. The fatigue scores are available as an Excel download for further analysis.

The model also allows a degree of interactivity for the user to assess the effects of fatigue countermeasures on the schedule in order to minimise fatigue.

Limitations of the model

A model is not reality, and no user should expect a model alone to predict whether a schedule is safe or not. A model will predict likely fatigue and provide an indication of the likelihood of the shift worker reaching a very tired state but will not assess the actual risk of an incident happening to any organisation.

However, an alert worker may quickly identify and mitigate a safety hazard that would otherwise cause an incident.

To truly identify the likelihood and impact of any fatigue-related hazard requires the user to make a judgement based on risk appetite, experience, and knowledge including relevant information from other sources. Moreover, the strengths and weaknesses of the model must also be known so that the output can be put into context before a judgement on organisational risk can be made.

FRI provides a single fatigue score and a single risk score of the probability of the worker reaching a high fatigue level, for each duty. It does not judge whether that is safe or not, but it is a useful measure for a user to make that judgement themselves.

The output of any model requires a good understanding not only of the subject that it's modelling but also of the assumptions that are built into the model. A manager should not expect the model to make a judgement or decision; that is the preserve of the manager alone. The model merely provides one set of data from which the manager may formulate an opinion. The manager may use a model to evaluate the relative effects of alternative strategies (do they improve or degrade performance?) before adopting any change.