

How do the models work?

Our biomathematical models use rosters to predict the fatigue levels and risk of different work schedules. The nature of the occupation dictates what data needs to be uploaded into the different models.

In its simplest form, an individual's work schedule comprises one or more duties over a given period. Each duty has characteristics, such as start and finish date-time and workload. A roster comprises one or more schedules. Roster data is uploaded into the model where our sophisticated, mature algorithms analyse the schedules, construct sleep patterns and predict levels of fatigue and risk, considering factors including workload, cumulative fatigue, circadian rhythm and time zone shifts.

The models output fatigue scores at regular intervals throughout the duty, allowing scheduling/rostering/crewing managers to assess rosters for fatigue risk. See [Fatigue Score](#) and [Duty Risk Metric](#) to learn more about the output.

Limitations of the model

A model isn't reality, and no user should expect a model alone to predict whether a schedule is safe or not for any individual. A model will predict likely fatigue for the average person in a homogeneous population (such as pilots or cabin crew) but won't assess the risk of the hazard to an airline.

To translate fatigue levels into a risk metric for the operation, the user must make judgements based on risk appetite, mitigations employed, experience and knowledge including relevant information from other sources. The strengths and weaknesses of the model and likely inherent errors from the modelling construction must also be known so the output can be put into context before a judgement on risk can be made.

The user shouldn't assume that the model output offers an absolute representation of the fatigue and risk of an individual. All models are based on the average of a given population. Consequently, the research conducted from which the algorithms are constructed must be that of the target audience that is study pilots to create a pilot model, and cabin crew to create a cabin crew model. Using studies of process workers for use by aircrew won't provide a good average for aircrew.

A model is a tool to help users make that decision of whether to accept the risk in its entirety, add in mitigations to reduce risk before permitting the duty to be implemented or to stop it and re-plan. A model isn't a decision-making system; it's a tool to help the user, or the Fatigue

Safety Action Group (FSAG), decide on which of the mitigations at their disposal is best to use within the performance capabilities of the model, when managing risk.

SAFE provides a score every 15 minutes and a Duty Risk for the whole duty. It does not judge whether the calculated scores are safe levels to work with or not, but it provides useful measures to help a user to make those judgment themselves. Humans decide the level of risk to take in concert with the risk appetite of the company and the constraints and mitigation options available at the time.

The output of any model requires a good understanding not only of the subject that it's modelling but also 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 or are they practical and fit with the resources available at the time?) before adopting any change.

The SAFE and CARE models are particularly suited to commercial aviation operations, because they are built specifically using the research studies implemented for each of these occupations and will identify potential hazards by calculating the likely fatigue score and Duty Risk of the average pilot or cabin crew, not specific individuals, that may be asked to perform the schedule of duties under scrutiny. It is the manager's responsibility to decide whether the hazard or portfolio of hazards identified is within the risk appetite of the airline or whether action needs to be taken to mitigate them or simply re-cast the schedule design.

Accordingly, it is vital that before SAFE or CARE are used in airline operations, training is taken. This should cover using the model, the assumptions that are built into the model, its limitation, how human performance is affected by fatigue and fatigue countermeasures. Rostering staff and managers will necessarily have to interpret the Fatigue and Duty Risk scores in their daily assessment of safety for the schedule under scrutiny.