zeitfarer

time shifts, you don't have to.

The Science Behind Zeitfarer

How sleep and the body clock are modeled, and how the model plans a pilot's trip

Prepared by Robin J McCall

McCall Aviation Consulting, LLC

September 17, 2026

Contents

How to read this document

Each chapter has the same shape. It opens with a plain-language page that any pilot can read in two minutes. The formal section that follows gives the equations, the constants, and their provenance. A short closing section says exactly how the Zeitfarer engine applies the science, and the sources for the chapter are listed at its end. The full reference list is at the back, followed by an appendix of every constant with its tier.

Three tiers are used for every number. Sourced: taken from a published measurement on a population, with the citation. Derived: computed from sourced numbers by arithmetic that is shown. Inferred: a judgment call, ours, named as such. A number from one pilot's own nights is never used to set a constant. One pilot can validate a model; one pilot cannot calibrate it.

1. Sleep, alertness, and the flight deck

Fatigue on the flight deck has two causes, and they are different. The first is how long you have been awake since your last real sleep. The second is what time it is on your body clock, whatever the local clock says. A crew member 20 hours awake is impaired at any hour. A crew member fully rested is still at a low point between 02:00 and 06:00 on their own clock, and lands there on most long-haul rotations. The two combine. Everything in this document follows from taking both seriously and refusing to guess at either one.

There is a third thing a pilot carries that is not a cause of fatigue but decides what can be done about it: the roster. A layover that is too short to adapt to, a report time that lands in the trough, a bunk break on the wrong side of the body night. The science says how sleep and the clock behave. The roster says which of the possible sleeps are actually available. The model has to hold both.

1.1 Two drivers, one impairment

The homeostatic drive is the pressure to sleep that builds with time awake and drains with sleep. The circadian drive is a roughly 24-hour rhythm in alertness set by a clock in the brain, entrained mainly by light. The regulator's own definitions describe the circadian trough plainly. FAA Advisory Circular 120-103A defines the window of circadian low as the period "roughly from 0200 to 0600, a time when physiological sleepiness is greatest and performance capabilities are lowest," and names a second, milder window "roughly from 1500 to 1700." Neither circular cites a study for the boundaries; they are a regulator's summary of the field, and this document treats them as such.

The homeostatic side has a number attached to it that pilots tend to remember. Dawson and Reid (1997) matched sustained wakefulness against blood alcohol concentration on the same performance test and reported three points: 17 hours awake performed like 0.05 percent, 22 hours like 0.08 percent, and 24 hours like 0.10 percent. The three points are not on a straight line, so the table is kept as a table rather than fitted.

Hours of sustained wakefulness	Performance-equivalent BAC	Source
17 h	0.05 %	Dawson and Reid, 1997
22 h	0.08 %	Dawson and Reid, 1997
24 h	0.10 %	Dawson and Reid, 1997

1.2 What the field studies found

The field evidence on airline crews is older and larger than most pilots expect. NASA's crew factors series in the early 1990s put actigraphs and logbooks on long-haul crews and measured where their sleep actually fell across layovers (Gander, Graeber, Connell, and Gregory, 1991), and then ran a controlled trial of planned cockpit rest on the same operation (Rosekind et al., 1994). In that trial a 40-minute rest opportunity yielded an average of 25.8 minutes of sleep with a sleep latency of 5.6 minutes, and 93 percent of the attempts produced sleep. The point for this document is not the nap itself, which Part 121 does not allow, but the demonstration that sleep in the operational environment can be measured and predicted.

Thirty years later the numbers are still being measured on the same kind of trip. Rempe et al. (2025) put actigraphy on 256 long-range and ultra-long-range pilots across 473 trips and nine routes and found an average of 7.2 hours (SD 1.7) of sleep per 24 hours of layover, with route means from 6 to 8 hours, and 83 percent of pre-trip sleep falling inside the 23:00 to 07:00 home-base night. When a layover began between 10:00 and 14:00 on the home-base clock rather than between 02:00 and 06:00, pilots slept 87 minutes longer on it. Where a layover starts on the body clock predicts how much sleep it yields. That single finding is why a model built on clock time alone cannot work.

1.3 The three clocks

A pilot on a rotation lives on three clocks at once. Local time is what the airport and the hotel run on. Home time is where the body clock started. The body clock itself is neither: it is home time plus whatever drift the light of the trip has produced, and it is the only one of the three that decides when sleep is possible. Every quantity in this document is stated on the body clock unless it says otherwise, and every conversion to local time is made at the last possible moment.

How Zeitfarer applies this

The engine models both drivers explicitly and never collapses them into a single fatigue score at the planning stage. Sleep pressure is tracked as a state variable that a short night leaves partly undischarged, so six hours awake after a bad night is a different pressure from six hours awake after a good one. The body clock is a second state variable moved by the light the route delivers. The Dawson and Reid table is used only to rank candidate plans against each other, with the 0.10 percent point as the zero of the scale, and the FAA's 0.04 percent limit for flight crew is reported as a flag, never used as a filter. Chapter 9 gives the detail.

Sources for this chapter

Dawson, D., & Reid, K. (1997). Fatigue, alcohol and performance impairment. Nature, 388(6639), 235.

Federal Aviation Administration. (2010). Basics of aviation fatigue (Advisory Circular 120-100).

Federal Aviation Administration. (2013). Fatigue risk management systems for aviation safety (Advisory Circular 120-103A).

Gander, P. H., Graeber, R. C., Connell, L. J., & Gregory, K. B. (1991). Crew factors in flight operations VIII: Factors influencing sleep timing and subjective sleep quality in commercial long-haul flight crews (NASA Technical Memorandum 103852).

Rempe, M. J., Rasmussen, I., Gregory, K., Johnson, C., Hsin, M., Flynn-Evans, E. E., Lamp, A., & Hilditch, C. J. (2025). Layover start timing predicts layover sleep quantity and timing on long-range and ultra-long-range trips. Sleep Advances, 6(1), zpaf002.

Rosekind, M. R., Graeber, R. C., Dinges, D. F., Connell, L. J., Rountree, M. S., Spinweber, C. L., & Gillen, K. A. (1994). Crew factors in flight operations IX: Effects of planned cockpit rest on crew performance and alertness in long-haul operations (NASA Technical Memorandum 108839).

2. The two-process model

Two things decide whether you can sleep right now: how long you have been awake, and where your body clock is. The first is a pressure that rises while you are up and drains while you sleep. The second is a wave that rises and falls once a day on your own clock. Sleep begins when the pressure climbs up to a line, and ends when it drains down to a lower line. Both lines ride the wave. That is the whole model, and it has stood since 1984.

The part that matters for a pilot is that there is one wave and two lines, not two curves. The gap between the lines is what keeps sleep in one block instead of switching on and off every few hours. The upper line is the one that moves: light, caffeine, a bunk, and stress all raise it. The lower line barely moves at all. Once you see the model that way, most of what a hard place to sleep does becomes one number.

2.1 Process S, the homeostatic pressure

Following Daan, Beersma, and Borbély (1984), sleep pressure S is dimensionless and runs between 0 and 1. It rises exponentially toward its ceiling while awake and falls exponentially toward its floor while asleep, each with its own time constant. The classic constants are 18.2 hours for the rise and 4.2 hours for the fall.

```
awake:    S(t) = mu_w + (S0 − mu_w) exp(−t / chi_w)        chi_w = 18.2 h, mu_w = 1.0
asleep:   S(t) = mu_s + (S0 − mu_s) exp(−t / chi_s)        chi_s =  4.2 h, mu_s = 0.0
```

The asymmetry is the physiology. Pressure builds slowly across a day and is discharged quickly across a night, which is why a full night after a long day is possible at all, and why the second half of the night contributes less than the first.

2.2 Process C, the circadian wave

The circadian component is not a sine wave. Daan, Beersma, and Borbély used a skewed waveform, and Achermann and Borbély (1992) gave it as a five-harmonic series whose coefficients have been reprinted by every group that works with the model since, including Rajaraman et al. (2009) and Skeldon, Dijk, and Derks (2014).

```
raw(h) = 0.97 sin(wh) + 0.22 sin(2wh) + 0.07 sin(3wh) + 0.03 sin(4wh) + 0.001 sin(5wh)      w = 2 pi / 24
C(h) = −raw(h − CBTmin + 3.99) / 1.003725987        so C runs from −1 at CBTmin to +1
```

The skew is the point. The wave takes the long way from trough to peak, about 16 hours, and the short way back, about 8 hours. The trough sits at the core body temperature minimum, CBTmin, and the peak lands in the early evening at what the field calls the wake maintenance zone, the stretch when falling asleep is hardest even when you are tired. Orient the wave the wrong way round and the evening zone lands at midday while every downstream number still looks plausible, which is why the orientation is checked against both landmarks rather than assumed.

One coefficient deserves a note for the nerds. Skeldon, Dijk, and Derks print the third harmonic as 0.007, in both the journal version and their preprint. The series originates with Achermann and Borbély (1992), and the two independent modeling groups that reprint it from the origin give 0.07 (Rajaraman et al., 2009; the Van Dongen group's US patent 11,241,194). Two shape arguments settle it: only with 0.07 do the coefficients decay monotonically, and only with 0.07 do the extrema land at exactly 4.00 and 20.00 hours after the phase reference, the 8 and 16 hour skew the waveform is known for. The difference is at most 0.063 on a peak-to-trough range of 2.0, so nothing a pilot sees turns on it, but the value used is 0.07 and the reason is recorded.

The 3.99 hour offset places the trough of the series at the phase reference, and the normalizer scales the series to a peak-to-trough range of exactly 2. CBTmin itself is placed two hours before the pilot's habitual wake time. That placement is widely repeated in the field and is carried here as an inferred number, not a sourced one, until a citation is recorded against it.

```
CBTmin = habitual wake − 2.0 h
```

2.3 Two thresholds on one wave

Sleep begins when S rises to the upper threshold and ends when S falls to the lower one. Both thresholds are the same wave at different heights.

```
H+(t) = H+0 + a C(phi(t))          H−(t) = H−0 + a C(phi(t))          a =
2.9 / 21.35 = 0.1358
```

The amplitude a is taken from Skeldon, Dijk, and Derks (2014), equation 15, where it was obtained by matching the two-process model to a neuronal model of the sleep switch rather than by fitting sleep records, which makes it independent of any dataset used here. Their threshold means are not taken, because they come with symmetric 45-hour time constants that are an artifact of that matching. The means are derived instead from the pilot's own habitual window.

The derivation has no fitted parameter. A habitual sleeper is at a steady state: the same S at bedtime every night and the same S at wake every morning. Solve that fixed point for the hours awake and the hours in bed at the pilot's sleep efficiency, then subtract the wave's value at each instant.

```
S_bed = the steady state of the awake and in-bed cycle at bedtime
S_wake = S after the night in bed at e
H+0 = S_bed − a C(bed hour)          H−0 = S_wake − a C(wake hour)
```

A 23:00 to 07:00 sleeper then reproduces 23:00 to 07:00 exactly, by construction, and two pilots with different habits get different thresholds from the same physics. The eight parameters of the model, listed by Skeldon et al. (2025), are the two time constants, the two asymptotes, the two threshold means, the amplitude, and the period. Six of them are sourced, and the two means are derived from the pilot.

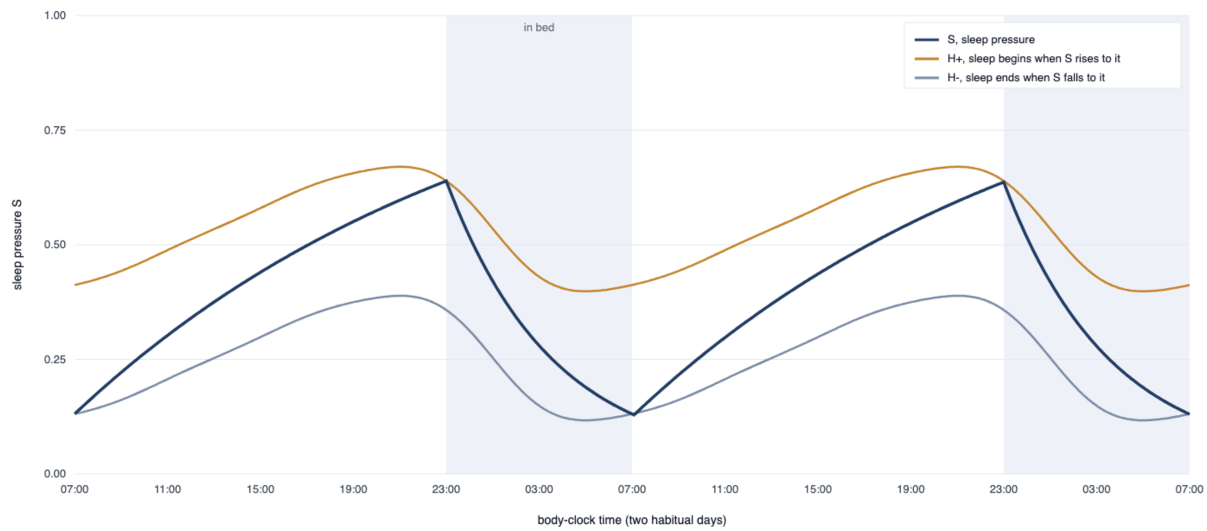

Figure 1. Sleep pressure S over two habitual days for a 23:00 to 07:00 sleeper at 90 percent sleep efficiency. S rises while awake and falls in bed. The upper threshold H+ (gold) and the lower threshold H- (steel) are the same circadian wave at two heights; here H+0 = 0.534 and H-0 = 0.253, both derived from the habitual window by the fixed point of section 2.3, not fitted. Sleep begins where S meets H+ and ends where it meets H-. Drawn from the equations of this chapter.

2.4 The labile upper threshold

Daan, Beersma, and Borbély argued that the upper threshold in particular is labile, responding to light, stimulants, and motivation, while the lower threshold is far more stable. Borbély, Daan, Wirz-Justice, and Deboer (2016) restate the point in their thirty-year reappraisal. It is the most useful single sentence in the literature for a pilot, because it says where a bunk, a cup of coffee, and a hard hotel night all act: on the onset threshold, and only there. Nothing about a bunk makes pressure drain faster once you are asleep. What the bunk does to sleep continuity, which is a different question, is taken up in chapter 6.

How Zeitfarer applies this

The engine is this model with no extra state. Every agent that makes sleep harder to start is added to the upper threshold, and none of them touches the lower one. Prior sleep debt is not a separate input: a short night leaves S incompletely discharged, so the next day starts from a higher pressure and every downstream sleep sees it. A ground sleep is entered as time in bed, not time asleep, or with a sleep tracker as the sleep it recorded, which the app converts to the time in bed it implies before the engine sees it, so S integrates as a mixture of the two exponentials at the pilot's efficiency e, with its own rate and asymptote. A pilot's own bedtime and wake time set the two threshold means, and the model reproduces that night exactly before it plans anything else.

in bed: rate = e / chi_s + (1 − e) / chi_w tau = 1 / rate mu = ((1 − e) / chi_w) / rate

S(t) = mu + (S0 − mu) exp(−t / tau) and h hours in bed are credited as h e hours of sleep

Sources for this chapter

Achermann, P., & Borbély, A. A. (1992). Combining different models of sleep regulation. Journal of Sleep Research, 1(2), 144-147.

Borbély, A. A., & Achermann, P. (1999). Sleep homeostasis and models of sleep regulation. Journal of Biological Rhythms, 14(6), 557-568.

Borbély, A. A., Daan, S., Wirz-Justice, A., & Deboer, T. (2016). The two-process model of sleep regulation: A reappraisal. Journal of Sleep Research, 25(2), 131-143.

Daan, S., Beersma, D. G. M., & Borbély, A. A. (1984). Timing of human sleep: Recovery process gated by a circadian pacemaker. American Journal of Physiology, 246(2), R161-R183.

Rajaraman, S., Gribok, A. V., Wesensten, N. J., Balkin, T. J., & Reifman, J. (2009). An improved methodology for individualized performance prediction of sleep-deprived individuals with the two-process model. Sleep, 32(10), 1377-1392.

Skeldon, A. C., Dijk, D.-J., & Derks, G. (2014). Mathematical models for sleep-wake dynamics: Comparison of the two-process model and a mutual inhibition neuronal model. PLoS ONE, 9(8), e103877.

Skeldon, A. C., et al. (2025). The complexity and commonness of the two-process model of sleep regulation from a mathematical perspective. npj Biological Timing and Sleep.

3. Falling asleep is a probability, staying asleep is a walk

The two-process model of chapter 2 has a switch in it: sleep starts the instant pressure reaches the upper line. Real people are not switches. A tired pilot lying down a little early may or may not drop off, and the odds depend on how tired and on what time it is on his clock. So the engine treats falling asleep as a probability that rises as pressure nears the line and falls in the evening zone where the body resists sleep. A probability can be integrated over a rest opportunity, which is how the plan says how much sleep a break is worth without pretending to know the minute.

Staying asleep is a different question with a different answer. Once asleep, a pilot does not stay asleep until pressure has drained: he sleeps in bouts of about half an hour, wakes briefly, and drops back. On the ground those wakes are so short he never remembers them. In a bunk they are longer and more frequent, and that is where the bunk's real cost sits. The engine models the night as that alternation of sleep and wake bouts, with one published bout length and nothing fitted.

3.1 Sleep propensity is analog

Bes (2025) argues that the deterministic gate in the two-process model obscures how the homeostatic and circadian processes interact, and proposes instead a propensity: the likelihood of falling asleep, or staying asleep, at a given time, with each process contributing a likelihood between 0 and 1. Lavie (1986) had measured the same thing thirty years earlier without naming it. Subjects on an ultrashort sleep-wake schedule who were told to try to sleep showed a decreased ability to do so in the two to three hours before habitual bedtime, the forbidden zone, whether they had been told to resist sleep or to attempt it. Decreased, not zero. A hard gate has always been a stronger claim than the evidence behind it.

The engine's two likelihoods are the natural reads of the quantities chapter 2 already has. The homeostatic term is how far pressure has climbed across the gap between the two thresholds; the circadian term is the wave itself, rescaled so that the trough of the night reads 1 and the evening peak reads 0.

```
P_h = clamp( (S − H−) / (H+ − H−), 0 to 1 )          P_c = (1 − C) / 2
```

3.2 How the two terms combine, and why it was decided by measurement

Bes proposes that the two likelihoods multiply. Paech, Ferguson, Sargent, Kennaway, and Roach (2012) measured the interaction directly, in 14 men across 13 days of forced desynchrony on a 28 hour day with 4.7 hour sleep opportunities, and found the opposite of what a product implies: in the presence of high homeostatic pressure the circadian modulation of sleep was reduced, and sleep efficiency stayed high across all circadian phases. The two claims are decidable by a derivative. Under a product the sensitivity of P to the circadian term is the homeostatic term, so the circadian effect grows as pressure rises. Under a union, the noisy-OR, that sensitivity is one minus the homeostatic term, and it vanishes as pressure rises, which is what Paech measured. The engine uses the union.

```
P = 1 − (1 − P_h)(1 − P_c)
```

A control settles it on the ground the model has to stand on. Integrated across a habitual 23:00 to 07:00 night, a model must return a night. The product returns 0.91 hours of sleep from 8 hours in bed, 11 percent. The union returns 6.94 hours, 87 percent, which is an ordinary sleep efficiency arrived at from a direction that had nothing to do with efficiency. Neither form has a free parameter, so the choice between them was a measurement, not a fit. The physiological claim of the union is plain: sleep comes if either the pressure or the clock favors it.

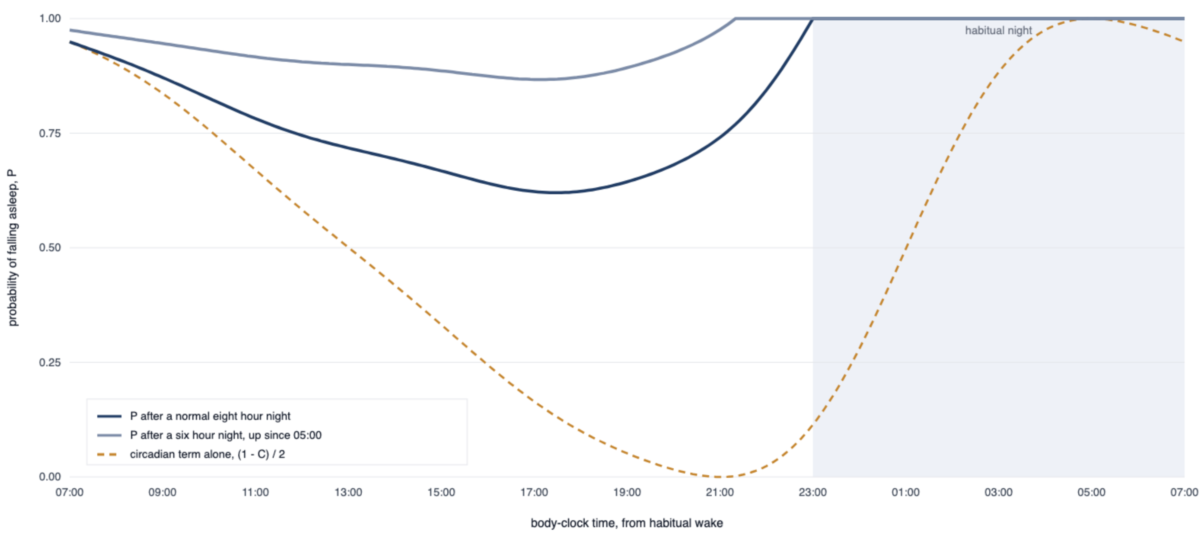

Figure 2. The probability of falling asleep across one body day for a 23:00 to 07:00 sleeper. The dashed line is the circadian term alone: highest in the early morning trough, lowest in the evening wake maintenance zone. The solid navy line is the union after a normal night, rising with pressure through the day and held down by the evening zone until bedtime. The steel line is the same man after a six hour night: the pressure term carries more of P, and the clock's modulation is much reduced, which is Paech's finding.

3.3 From a probability to an onset

A propensity is turned into a time by treating it as a hazard. The odds of the propensity are the rate at which sleep arrives, per hour, and the waiting time to onset is exponential with that rate. The deterministic reading of the same hazard is its mean: onset lands where the cumulative hazard reaches one. The hazard is capped so that even a man who is very sleepy at a good hour does not fall asleep instantly. The cap is a published sleep onset latency, 9.03 minutes, the unstressed control arm of Ackermann, Cordi, La Marca, Seifritz, and Rasch (2019), measured in a 90 minute nap opportunity.

```
k = P / (1 − P) per hour, capped so that onset latency is never under 9.03
min
```

The walk advances in steps of two and a half minutes. That resolution was chosen by measurement: once the cap bounded the hazard, the step became load-bearing, and five minute steps were shown to move onsets by more than the physiology justified.

3.4 Staying asleep is a walk of bouts

The propensity is a falling-asleep likelihood. Asked a staying-asleep question it over-answers, because the circadian term alone keeps it high after pressure has drained. In development that one property produced three separate defects, among them a plan that credited sixteen hours in bed. The resolution was not to invent a second function. The two-process walk already is one: once asleep, pressure falls until it meets the lower threshold, and the hysteresis between the thresholds is what makes a night one block. Onset from the propensity, duration from the walk. That division is the model.

Within the walk, sleep is not continuous. Arnardóttir, Þorsteinsson, and Karlsson (2010) measured sleep-wake cyclicity across the human lifespan and report, for adults aged 62 to 74, a mean sleep bout of 31.1 minutes (SD 4.7), with wake bouts in that band distributed exponentially. An exponential wake bout means a constant rate of dropping back to sleep, which is the simplest possible model and the one the data license.

```
tau_s = 31.1 min
```

The mean wake bout is not fitted. A two-state alternating process spends a fraction tau_s / (tau_s + tau_w) of its time asleep. Setting that fraction equal to the sleep efficiency of chapter 5 and inverting gives the wake bout from two numbers that already exist.

```
tau_w = tau_s (1 − e) / e          at e = 0.818 (age 64):  tau_w = 6.9 min
```

The referent is not exact, and the direction of the error is known. Efficiency is total sleep over time in bed, so it carries onset latency as well as wake after sleep onset, while the identity charges all of the non-sleep time to wake bouts. The wake bout is therefore a slight overestimate, and the model charges a little too much for each arousal rather than too little. The identity also carries a check: it holds exactly when the arousal rate of the simulation matches the bout it was derived from, 60 / 31.1, about 1.9 an hour, and that is the rate the engine's own age-matched arousal channel found by a different route.

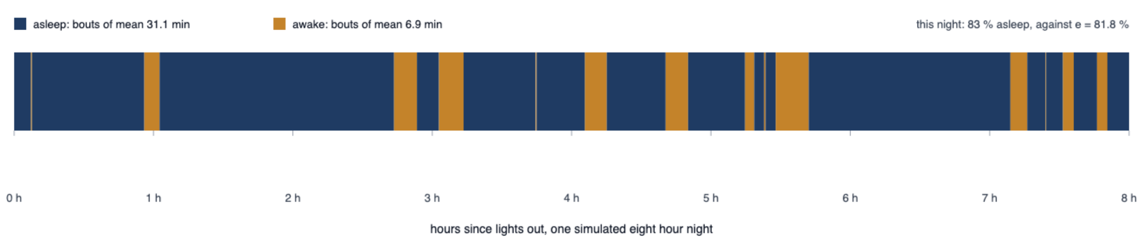

Figure 3. One simulated night of the alternating walk at the efficiency of a 64 year old: sleep bouts of mean 31.1 minutes (navy) and wake bouts of mean 6.9 minutes (gold), both drawn from exponential distributions. This particular night is 83 percent asleep; over many nights the fraction converges on the efficiency it was derived from. Most of the wakes are shorter than the sleeper would remember.

How Zeitfarer applies this

Every ground sleep on a card is a walk that begins where the propensity's hazard says and ends where pressure meets the lower threshold, and the credited sleep is the walk's own length at the pilot's efficiency. The expected sleep printed under a bunk break is an integral of the propensity over the break, with pressure falling as it is discharged so that the answer saturates instead of running away. The multiplicative form is kept in the code beside the union with a test that it fails the habitual-night control, so the decision cannot be quietly reversed by someone who has read only the newer paper. What the bunk adds to the walk, an arousal rate by facility class and by the ride, is chapter 6.

Sources for this chapter

Ackermann, S., Cordi, M., La Marca, R., Seifritz, E., & Rasch, B. (2019). Psychosocial stress before a nap increases sleep latency and decreases early slow-wave activity. Frontiers in Psychology, 10, 20.

Arnardóttir, H., Þorsteinsson, H., & Karlsson, K. Æ. (2010). Dynamics of sleep-wake cyclicity at night across the human lifespan. Frontiers in Neurology, 1, 156.

Bes, F. (2025). Treatise on an alternative perspective on the two-process model of sleep regulation. npj Biological Timing and Sleep, 2, 23.

Lavie, P. (1986). Ultrashort sleep-waking schedule. III. 'Gates' and 'forbidden zones' for sleep. Electroencephalography and Clinical Neurophysiology, 63(5), 414-425.

Paech, G. M., Ferguson, S. A., Sargent, C., Kennaway, D. J., & Roach, G. D. (2012). The relative contributions of the homeostatic and circadian processes to sleep regulation under conditions of severe sleep restriction. Sleep, 35(7), 941-948.

4. Where sleep may begin: the placement rule

A descriptive model says when sleep is possible. A plan has to say when to lie down, and those are not the same question. Left to the physiology alone, a planner offers sleep wherever the pressure crosses the line, and on a long layover that produces what one pilot called 70 minute sleeps five hours apart: correct on paper, and not what a human does. So the engine asks a prior question first. Where are you in your own day? Near your habitual night, take the night. In the middle of the day, take a short sleep. Nowhere else.

Then the roster has its say. Duty wins, because it is the one clock a pilot does not control. When the night would end long before the van, it is pulled later toward the duty, as far as the body allows. When the gap is too wide to close, the night stays where it belongs and a short sleep before duty follows it. And on nearly every layover the whole plan is drawn on home time, because the layover is too short to adapt and half-adapting is paid for on the way home.

4.1 The night, and how far it can be moved

The long sleep belongs in the habitual window on the body clock, the same window that set the thresholds in chapter 2. Dijk and Czeisler (1994) showed that sleep propensity peaks near the core body temperature minimum and falls steeply after it, a timing they called paradoxical because it places the strongest drive to sleep at the end of the night rather than the beginning. The consequence for a planner is a bound: a night can be delayed and still be a night up to the point where it would begin after the temperature minimum. An onset past that point is a day sleep, which the rising lower threshold will end early.

```
maximum night delay = CBTmin − habitual bedtime          6.0 h on a 23:00 to
07:00 profile
```

The bound is derived from the placement of the temperature minimum rather than picked, and it is carried as inferred. A delay is only worth what it buys. The gap is measured from the habitual bedtime to the start of the pre-duty band. A gap the night can close is closed exactly, so the night ends at get-ready and is the pre-duty sleep. A gap it cannot close is left alone: the night stays the pilot's own, and the short sleep before duty follows it. That rule was written after a first version delayed a night by the full bound on a 48 hour layover and lost two and a half hours of sleep for nothing.

4.2 The afternoon zone

The standard two-process model produces one peak of sleep propensity, at night, and one trough, in the evening. Bes (2025) names two more features of the measured 24 hour curve: a morning minimum at sleep offset and a secondary, smaller afternoon peak. The afternoon peak is the one pilots keep meeting in the air. Bes, Jobert, and Schulz (2009) modeled it and place it at 14:00 for a midnight to 08:00 sleeper, six hours after wake, with the evening trough near 20:00. Souabni et al. (2022), reviewing napping in older adults, found that all fifteen included studies started the nap between 13:00 and 14:00, six to seven hours after a typical wake.


```
afternoon zone center = habitual wake + 6.0 h          half-width 2.5 h
depth 1/3 of the circadian amplitude
```

The center is sourced and re-phases with the pilot's own wake time. The width and the depth are inferred: no published width exists, and Bes describes the peak as smaller without quantifying it, so one third of the circadian amplitude is a reading of smaller that is large enough to reach when a pilot carries real pressure and small enough that a rested pilot on his own schedule is not sent to bed at four in the afternoon. The zone lowers the onset threshold only. It is a peak in the likelihood of falling asleep, not a change in how long sleep lasts, so it never touches the offset threshold.

4.3 The evening zone

Lavie (1986) found the forbidden zone centered on 20:00 to 22:00 and the nocturnal sleep gate opening abruptly after it; Strogatz, Kronauer, and Czeisler (1987) placed minimal sleep tendency one to three hours before habitual bedtime. The five-harmonic wave of chapter 2 already produces this zone at its peak, so the engine carries no separate rule for it: the propensity of chapter 3 falls there on its own. What the evidence adds is that the zone reduces sleep and does not forbid it. Monterastelli, Adams, Eastman, and Crowley (2024), on an ultradian schedule of two hours awake and two hours dark that builds chronic pressure, comparing adolescents with adults, found that the adults given a two hour opportunity in the two hours before their melatonin onset fell asleep in 9.2 minutes and slept 69 of the 120 minutes; two hours after the onset they slept 90. De Zeeuw et al. (2018) found the zone still producing alertness at 37.5 hours awake. A tired pilot can sleep in the evening zone; he gets less.

4.4 The guards

Ingre et al. (2014) fitted a three-process model to 136 aircrew, 964 sleep logs, and 5,443 sleeps, and their sleep generator carries guards that any placer needs: no sleep within an hour of a duty period, no new sleep less than two hours after waking, and no waking in the first hour of a sleep, included, in their words, to prevent flip-flop behavior. The engine's guards do the same work in its own terms.

- After a sleep of four hours or more, nothing for six hours. The rule is the pilot's, the same six hours the afternoon zone is built on, and the reason given is that there is no pressure left to sleep on. The four hour line is inferred: three cycles of the shortest sleep.
- One short sleep in the afternoon zone, and the zone closes behind it.
- Two blocks closer than three hours belong to the same rest period. No aviation paper prints a merge gap; the nearest is Devine et al. (2020), who define a major sleep interval as longer than three hours. Inferred.
- The shortest sleep worth planning is one cycle, 70 minutes, the short end of the 70 to 120 minute range usually given for a sleep cycle. Below it the evidence is about naps: Rosekind et al. (1994) found a 40 minute cockpit opportunity yielding 25.8 minutes of sleep on 93 percent of attempts, and Dinges (1987) found a two hour opportunity six hours after waking yielding 80 minutes with benefits lasting a day. Those are the nap's case; the plan's unit is the cycle.

4.5 Why home time is held

Rempe et al. (2025) found 83 percent of pilots' pre-trip sleep inside the 23:00 to 07:00 home-base night, and that the timing of layover sleep was predicted by where the layover began on the home-base clock rather than on the local one. Gander et al. (2016), studying 39 Boeing 747-400 pilots on nine to thirteen day trips with one day layovers, cite an earlier nine day transpacific sequence in which the core body temperature rhythm ran with an average period of 24.6 hours: the clock drifts later across such a trip whatever the pilot intends, and the model has to carry that drift, which chapter 7 takes up. The sleep scientists at Washington State University who reviewed this model in August 2026 gave the same operational advice: anchor to home time and do not try to adapt, because almost all airline trips are too short for adaptation to complete. The engine's own measurement agreed. On a westbound trip with layovers of every length, an adapting plan produced no net gain in sleep over an anchored one at any layover length. Adaptation remains available as a switch on each layover for a genuinely long stay; it is off by default.

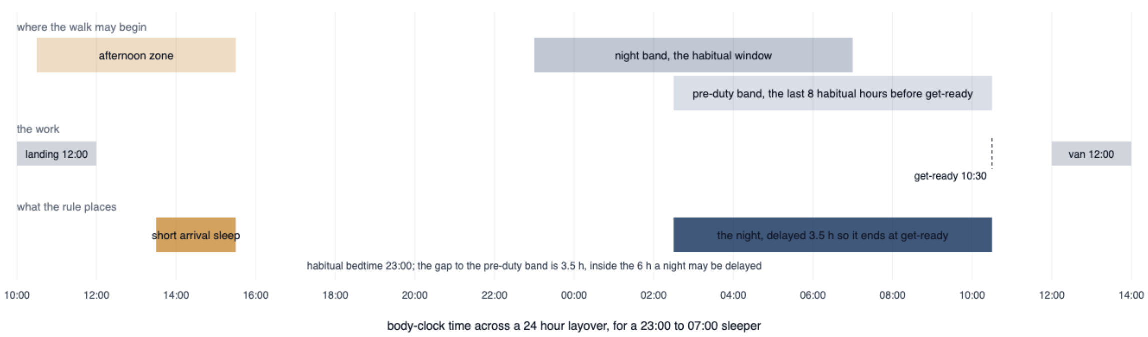

Figure 4. The placement rule on a 24 hour layover, drawn on the body clock of a 23:00 to 07:00 sleeper who lands at body noon and is collected at body noon the next day. The walk may begin only inside the three bands: the afternoon zone, habitual wake plus six hours with a half-width of 2.5 hours; the night band; and the pre-duty band. The short arrival sleep falls in the afternoon zone. The night is delayed 3.5 hours so that it ends at get-ready, because that gap is inside the 6 hour bound; a wider gap would have left the night in place with a short pre-duty sleep after it. Schematic, drawn from the rule rather than taken from a plan.

How Zeitfarer applies this

Every layover card, the pre-trip card, and the recovery card are built by the same rule: find the pilot's three bands on his body clock, let the walk of chapter 3 begin only inside them, let duty pull the night within its bound, and credit the walk's own length. The rule replaced a search that scored candidate plans against a rolling 24 hour criterion, which produced knife-edge answers, and it replaced a descriptive walk that offered sleep from the first free minute of a window. A first-break pilot on an augmented leg gets one piece of guidance the bands do not produce, a 90 minute stay-awake window before report so that he boards with pressure to spend in the bunk; that is an operating rule, attributed to the pilot who set it, not a model output. The hold on home time is the default on every layover card, with the switch to adapt beside it.

Sources for this chapter

Bes, F. (2025). Treatise on an alternative perspective on the two-process model of sleep regulation. npj Biological Timing and Sleep, 2, 23.

Bes, F., Jobert, M., & Schulz, H. (2009). Modeling napping, post-lunch dip, and other variations in human sleep propensity. Sleep, 32(3), 392-398.

de Zeeuw, J., Wisniewski, S., Papakonstantinou, A., Bes, F., Wahnschaffe, A., Zaleska, M., Kunz, D., & Münch, M. (2018). The alerting effect of the wake maintenance zone during 40 hours of sleep deprivation. Scientific Reports, 8, 11012.

Devine, J. K., Choynowski, J., Burke, T., Carlsson, K., Capaldi, V. F., McKeon, A. B., & Sowden, W. J. (2020). Practice parameters for the use of actigraphy in the military operational context: The Walter Reed Army Institute of Research Operational Research Kit-Actigraphy (WORK-A). Military Medical Research, 7(1), 31.

Dijk, D.-J., & Czeisler, C. A. (1994). Paradoxical timing of the circadian rhythm of sleep propensity serves to consolidate sleep and wakefulness in humans. Neuroscience Letters, 166(1), 63-68.

Dinges, D. F., Orne, M. T., Whitehouse, W. G., & Orne, E. C. (1987). Temporal placement of a nap for alertness: Contributions of circadian phase and prior wakefulness. Sleep, 10(4), 313-329.

Gander, P., Mulrine, H. M., van den Berg, M. J., Wu, L., Smith, A., Signal, L., & Mangie, J. (2016). Does the circadian clock drift when pilots fly multiple transpacific flights with 1- to 2-day layovers? Chronobiology International, 33(8), 982-994.

Ingre, M., Van Leeuwen, W., Klemets, T., Ullvetter, C., Hough, S., Kecklund, G., Karlsson, D., & Åkerstedt, T. (2014). Validating and extending the three process model of alertness in airline operations. PLoS ONE, 9(10), e108679.

Lavie, P. (1986). Ultrashort sleep-waking schedule. III. 'Gates' and 'forbidden zones' for sleep. Electroencephalography and Clinical Neurophysiology, 63(5), 414-425.

Monterastelli, A. J., Adams, J., Eastman, C. I., & Crowley, S. J. (2024). The forbidden zone for sleep is more robust in adolescents compared to adults. Frontiers in Sleep, 2, 1304647.

Rempe, M. J., Rasmussen, I., Gregory, K., Johnson, C., Hsin, M., Flynn-Evans, E. E., Lamp, A., & Hilditch, C. J. (2025). Layover start timing predicts layover sleep quantity and timing on long-range and ultra-long-range trips. Sleep Advances, 6(1), zpaf002.

Rosekind, M. R., Graeber, R. C., Dinges, D. F., Connell, L. J., Rountree, M. S., Spinweber, C. L., & Gillen, K. A. (1994). Crew factors in flight operations IX: Effects of planned cockpit rest on crew performance and alertness in long-haul operations (NASA Technical Memorandum 108839).

Souabni, M., Souabni, M. J., Hammouda, O., Romdhani, M., Trabelsi, K., Ammar, A., & Driss, T. (2022). Benefits and risks of napping in older adults: A systematic review. Frontiers in Aging Neuroscience, 14, 1000707.

Strogatz, S. H., Kronauer, R. E., & Czeisler, C. A. (1987). Circadian pacemaker interferes with sleep onset at specific times each day: Role in insomnia. American Journal of Physiology, 253(1), R172-R178.

5. Sleep efficiency, age, and what a report is worth

Nobody sleeps for the whole time they are in bed. The share that is sleep is called sleep efficiency, and it falls with age: about 89 percent for people in their twenties, about 83 percent in their fifties, and it keeps falling after sixty. That gap is the difference between an eight hour night that delivers seven hours of sleep and one that delivers six and a half. It is not something a pilot can feel, because most of the waking is too brief to remember.

So when the model asks how long you were in bed, it is asking the question you can answer, and doing the conversion with a published number for your age. A wearable that reports sleep directly skips the conversion. Either way, the plan is built on credited sleep, not on time lying down.

5.1 The measurement

Boulos et al. (2019) pooled 169 studies and 5,273 healthy adults recorded by overnight in-lab polysomnography and scored to the AASM standard. Per decade of age they found total sleep time falling by 10.1 minutes (95 percent CI 7.5 to 12.8), sleep efficiency by 2.1 percent (1.5 to 2.6), wake after sleep onset rising by 9.7 minutes (6.9 to 12.4), and sleep onset latency by 1.1 minutes (0.3 to 1.9). Their Table 2 gives the level by age band.

Age band	Sleep efficiency	95 % CI	Studies
18-34	89.0 %	88.0 to 90.0	65
35-49	85.4 %	83.7 to 87.1	35
50-64	83.2 %	81.0 to 85.4	27
65-79	77.5 %	73.0 to 81.9	16

Ohayon, Carskadon, Guilleminault, and Vitiello (2004) had earlier pooled 65 studies and 3,577 subjects and found that efficiency is the one sleep parameter that keeps declining after sixty, where most others level off. The two meta-analyses agree on the slope through midlife; Ohayon supplies the shape of the tail.

5.2 The curve, and why it is not re-pinned

The engine uses a straight line through age with a published slope and a chosen level, clamped to a physiological range.

$$e(\text{age}) = 0.90 - 0.021\,(\text{age} - 25)\,/\,10 \qquad \text{held within } 0.65 \text{ to } 0.95$$

The slope is the Boulos figure. The level, 90 percent at 25, was chosen from the secondary reporting of young adults at roughly 90 to 95 percent before the primary table had been opened, and is carried as inferred. When the table was read, the line landed inside the 95 percent interval at all four bands and within 0.1 points of the 50 to 64 mean: 89.8 against 89.0, 86.4 against 85.4, 83.3 against 83.2, 80.1 against 77.5 at the band midpoints. That is a validation, and the tier is left as inferred on purpose: a number that was guessed and then found to agree has been checked, not sourced.

The obvious improvement, pinning the line to the youngest band at 89 percent, is worse. The published decline is not linear: it runs 2.25 points per decade from 26 to 42, 1.47 from 42 to 57, and 3.80 from 57 to 72. A line pinned to the young band reads 81.0 percent at 64, the very bottom of the 50 to 64 interval. The level of 90 at 25 is what lets one straight line hold both ends of a curve that bends.

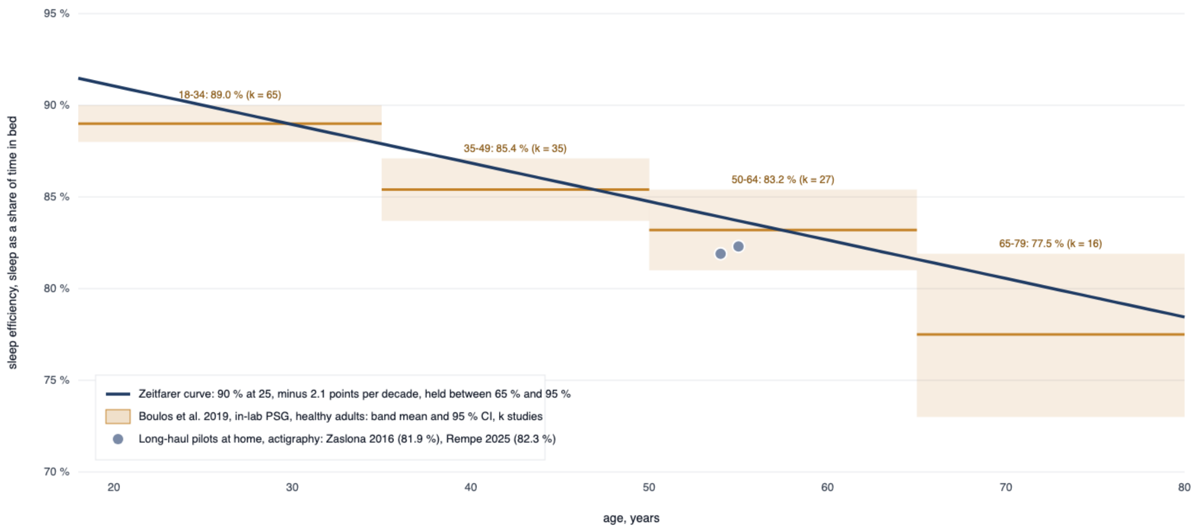

Figure 5. Sleep efficiency against age. The line is the curve the engine uses. The shaded spans are the Boulos et al. (2019) age bands with their 95 percent confidence intervals from in-lab polysomnography of healthy adults. The two points are long-haul pilots measured by actigraphy at home, from Zaslona (2016) and Rempe et al. (2025), placed at their samples' median and mean ages.

5.3 Pilots at home

The laboratory population is not the pilot population, so the curve was checked against the two actigraphic measurements of long-haul pilots sleeping at home. Zaslona (2016) recorded Delta Boeing 767 crews, median age 54, on two baseline days and found 81.8 (SD 8.0) and 82.0 (SD 7.3) percent. Rempe et al. (2025) found 82.3 percent in the 24 hours before the outbound duty across 256 pilots with a mean age near 55. Both sit about a point below the curve and inside the Boulos interval for their band. Actigraphy over a 24 hour interval and polysomnography over a night are not the same referent, so the agreement is read as support, not as a calibration. The level is kept at the laboratory anchor for three reasons: a line through the pilot point would fall below the youngest Boulos interval, actigraphy over 24 hours reads lower than a night of polysomnography for the same sleeper, and Boulos found efficiency 2.7 points higher per 10 percent more men in a sample, so pilot samples that are over 90 percent male should read above the laboratory, not below it. The gap is read as method, not population.

5.4 From time in bed to credited sleep

A report of h hours in bed is credited as h e hours of sleep, and the sleep pressure of chapter 2 integrates the block as a mixture of its awake and asleep exponentials at e, so the two views of the same night agree. A pilot who wears a sleep tracker reports what the device measured, and the model takes it as sleep: the equivalent time in bed is the measured sleep divided by e, and the pressure calculation runs on that. Reports already entered keep the meaning they had when entered. The bunk is the one place the age curve is not applied, because the in-flight model of chapter 6 accounts for its arousals explicitly and applying e as well would count the same disturbance twice.

How Zeitfarer applies this

Age is the one profile field the efficiency curve reads, and until it is set the plan assumes 40 and says so. Every ground sleep on every card is shown as time in bed alongside the sleep it is credited with, so the gap between the two is visible rather than silent. The sleep tracking device switch changes the question the cards ask, not the arithmetic behind it.

Sources for this chapter

Boulos, M. I., Jairam, T., Kendzerska, T., Im, J., Mekhael, A., & Murray, B. J. (2019). Normal polysomnography parameters in healthy adults: A systematic review and meta-analysis. The Lancet Respiratory Medicine, 7(6), 533-543.

Ohayon, M. M., Carskadon, M. A., Guilleminault, C., & Vitiello, M. V. (2004). Meta-analysis of quantitative sleep parameters from childhood to old age in healthy individuals: Developing normative sleep values across the human lifespan. Sleep, 27(7), 1255-1273.

Rempe, M. J., Rasmussen, I., Gregory, K., Johnson, C., Hsin, M., Flynn-Evans, E. E., Lamp, A., & Hilditch, C. J. (2025). Layover start timing predicts layover sleep quantity and timing on long-range and ultra-long-range trips. Sleep Advances, 6(1), zpaf002.

Zaslona, J. L. (2016). In-flight sleep as a pilot fatigue mitigation on long range and ultra-long range flights [Doctoral thesis, Massey University].

6. The bunk

The bunk is the hardest sleep on the trip to predict, and the one the plan most needs to get right. How much a break yields depends on where it falls on the body clock, on what the airplane provides to sleep in, on the ride, and on how long the break is, and pilots who answered the question honestly disagree with each other by hours. So the engine does not give a bunk break a number. It gives it a distribution, from many simulated attempts, and the plan is allowed to depend only on the low end of it.

What the polysomnography says is worth knowing on its own. In the same pilots, measured in the bunk and in their layover hotel, falling asleep took the same time in both places. The bunk's cost was in staying asleep: nearly twice the awakenings per hour and eighteen points less efficiency. That is why the engine models the bunk as a walk of sleep and wake bouts with an arousal rate, and why the ride you report after the break prices what you got.

6.1 What the measurements say

Signal, Gander, van den Berg, and Graeber (2013) recorded 21 male Boeing 777 pilots by polysomnography during a seven hour in-flight rest break and in their layover hotel. The pilots spent 4.7 of the 7 hours attempting to sleep and obtained 3.3 hours: 70 percent efficiency in the bunk against 88 percent in the hotel, with 7.7 awakenings an hour against 4.6, while latency to sleep onset did not vary by location. Slow wave sleep in flight had a median of 0.5 percent. A first-half break yielded 2.7 hours against 3.9 for a second-half break. Age was the most consistent factor: older crew had longer latencies, less sleep, and more awakenings.

Signal et al. (2024) put actigraphy on 72 Boeing 787-9 pilots flying Auckland to Chicago and back, 14.6 and 16.3 hours, in four-pilot crews, and reported sleep obtained as a percentage of the rest time available by quartile of the cruise on the domicile clock. The first quartile was the worst on both directions, 29 percent outbound and 17 percent inbound, and the best quartiles yielded 65 to 67 percent. Gander et al. (2013) found 3 hours 53 minutes of bunk sleep on ultra-long-range flights against 3 hours 15 minutes on long-range, about 12 minutes more per extra hour of flight, and Rempe, Basiarz, Rasmussen, Belenky, and Lamp (2022) reported means of 3.4 to 5.2 hours by direction, crew, and range across 231 pilots. Roach, Darwent, and Dawson (2010) estimated that bunk sleep gives about 70 percent of the recovery of a duration-matched bed sleep.

One more measurement matters for how the engine treats what pilots say. Signal, Gale, and Gander (2005) compared the same 21 pilots' subjective estimates of in-flight sleep with polysomnography and found them unbiased at the mean, with mean differences between minus 36 and plus 20 minutes, but noisy per person, with 95 percent limits near two hours. A large poll of pilots therefore averages the noise away and can be trusted for its mean. It cannot be trusted for any one pilot.

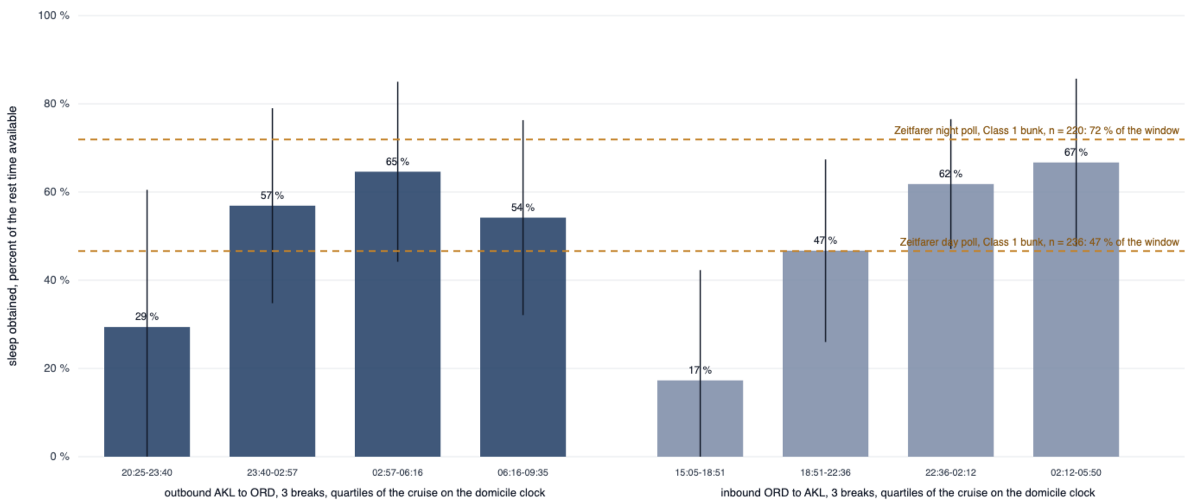

Figure 6. Bunk sleep as a share of the rest time available, by quartile of the cruise on the domicile clock, from Signal et al. (2024): 72 pilots, three-break schemes, means with one standard deviation. The dashed lines are the two Zeitfarer polls of 456 pilots, Class 1 bunk, night and day on the body clock. Sleep in a bunk is a function of where the break falls on the body clock before it is a function of anything else.

6.2 The four channels of a rest facility

The regulation carves a rest facility into exactly the properties the physiology needs. 14 CFR 117.3 defines a Class 1 facility as a bunk or other surface allowing a flat sleeping position, separate from the flight deck and the passenger cabin, temperature controlled, with control of light and isolation from noise and disturbance; a Class 2 facility as a seat allowing a flat or near-flat position, separated from passengers by at least a curtain, with some sound mitigation and reasonably free from disturbance; and a Class 3 facility as a seat reclining at least 40 degrees with leg and foot support. Flat against reclined is an onset question. Isolation, a curtain, or nothing is an arousal question. Control of light is a clock question. The engine carries each as its own channel.

Channel	Hotel	Class 1	Class 2	Class 3	Tier
onset penalty, added to H+	0.00	0.10	0.13	0.30	inferred, from the regulation's wording
settle-in overhead, minutes	15	20	10	5	operational, the pilot's own account
arousals per hour, a calm night	0.10	0.15	0.90	1.60	inferred, from the regulation's wording
arousals per hour, the prediction	0.10	1.2	1.9	scaled	sourced, the two polls

The overhead runs the other way from the comfort: the seat is cheaper to get into. On a 90 minute break the ten minutes saved outweigh the onset penalty. The onset penalty is an allowance added to the upper threshold of chapter 2 only, on the principle of section 2.4 that nothing about a facility makes pressure drain faster once a pilot is asleep. It is inferred from the regulation's wording; no study measures an onset cost by facility class, and the polysomnography above found none between bunk and hotel. It is kept beside the arousal channel by a ruling of September 15, 2026.

6.3 The prediction, and why it is a distribution

Before a break is flown, the engine simulates it 160 times at 30 second steps. Onset follows the propensity of chapter 3 at the bunk's raised threshold. Arousals arrive as a Poisson process at the facility's rate; each ejects the pilot and costs one wake bout, the derived bout of section 3.4, before the onset hazard can act again. The rest window itself opens 40 minutes after departure and closes 60 minutes before arrival, the crew's slots divide it, and the facility's settle-in and a 20 minute wake call are trimmed from each slot.

```
mean hours    the 20th percentile    probability of sleeping at all
expected number of episodes
```

Four outputs come back and only one may be planned on. The mean is a display number. The 20th percentile is the only figure a plan may depend on, because crediting the mean and getting nothing leaves a pilot short on approach, and the loss is not symmetric. A midday break's low quantile is close to zero, which is the model's way of saying what a pilot said unprompted during the beta: nineteen hours awake and then depending on sleeping well in the bunk is setting yourself up to be very tired.

The arousal rate the prediction uses is not the calm rate. It was set from two polls of pilots run in August 2026, 236 answering for a body-day break and 220 for a body-night break, asked what fraction of a Class 1 or Class 2 break they typically sleep. Class 1 came back at 72 percent of the window at night and 47 percent by day; Class 2 at 64 and 44. Inverted through the model's own population mean on standard breaks, one rate lands both phases: 1.2 arousals an hour reproduces 72 and 46 percent for Class 1, and 1.9 reproduces 64 and 43 for Class 2. Class 3 has no poll and is scaled by the calm-rate ratio. The polls are self-report and read high, so these rates are a ceiling on generosity rather than a floor. The wake-bout identity of chapter 3 provides an independent check: it holds exactly when the arousal rate matches the published sleep bout, about 1.9 an hour, which is the same order as the poll's answer and far from the calm 0.15 the prediction once used.

6.4 The ride

After the break the pilot reports two things: how long he was in the bunk, and how the ride was. The ride picker uses the FAA's own duration terms for turbulence, none, occasional, intermittent, and continuous, and every option names what happened to the sleep, because pilots gauge turbulence by how they slept. Each answer attaches an arousal rate: 0.15 an hour for a smooth night, the facility's calm rate; 0.50 for occasional bumps; 1.35 for intermittent bumps that kept waking him; and 2.15 for turbulence through most of the break. The rates are inferred, the pilot's judgment, with one anchor corroborated against a wearable on a real leg to within three minutes. They price the report, never the prediction.

```
ride efficiency = max( 0, e − (lambda_ride − lambda_calm) tau_w / 60 )
credited sleep  = min( reported time in bunk, window ) x ride efficiency
```

A smooth ride credits exactly the time-in-bed credit of chapter 5. Every rougher answer credits less, never more, which is the safe side of a wrong number. An unanswered ride reads as smooth. A pilot who wears a sleep tracker enters the sleep it recorded instead of the time in the bunk, and that number is credited as it is, capped at the break: the device has already seen the arousals, so the ride is kept as a record against the forecast rather than as a price. On those reports the forecast and the report share a referent for the first time, and their difference is the model's error.

6.5 Where the bunk's price is paid

Signal et al. (2013) put the bunk's cost in the stay-asleep channel with onset unchanged. In the engine that channel is the arousal walk, and the polls' rate of 1.2 an hour, the wake-bout identity's 1.9, and the polysomnography ratio of 7.7 to 4.6 awakenings between bunk and hotel all point the same way. The onset penalty stays as an allowance for what a reclining seat and a curtained cabin do to falling asleep, a question the polysomnography, run only in Class 1 bunks, did not test. What the engine does not model is the architecture: a median of half a percent of slow wave sleep in flight says bunk sleep is light sleep, and no term prices that. It is on the list in chapter 10.

How Zeitfarer applies this

The expected sleep printed under a bunk break is the low quantile of the ensemble, and a break at the body's daily peak reads as little to no sleep rather than as a number. The bunk is the one place the age efficiency curve is not applied, because the walk accounts for the arousals explicitly. A first-break pilot is told to stay awake in the run-up to report so that he boards with pressure to spend. After the break, the ride picker's answer is what the plan re-times from, and the plan never depends on a body-midday break for the sleep the next duty needs.

Sources for this chapter

Gander, P. H., Signal, T. L., van den Berg, M. J., Mulrine, H. M., Jay, S. M., & Mangie, J. (2013). In-flight sleep, pilot fatigue and Psychomotor Vigilance Task performance on ultra-long range versus long range flights. Journal of Sleep Research, 22(6), 697-706.

Rempe, M. J., Basiarz, E., Rasmussen, I., Belenky, G., & Lamp, A. (2022). Pilot in-flight sleep during long-range and ultra-long range commercial airline flights. Aerospace Medicine and Human Performance, 93(4), 368-375.

Roach, G. D., Darwent, D., & Dawson, D. (2010). How well do pilots sleep during long-haul flights? Ergonomics, 53(9), 1072-1075.

Signal, T. L., Gale, J., & Gander, P. H. (2005). Sleep measurement in flight crew: Comparing actigraphic and subjective estimates to polysomnography. Aviation, Space, and Environmental Medicine, 76(11), 1058-1063.

Signal, T. L., Gander, P. H., van den Berg, M. J., & Graeber, R. C. (2013). In-flight sleep of flight crew during a 7-hour rest break: Implications for research and flight safety. Sleep, 36(1), 109-115.

Signal, T. L., van den Berg, M. J., Zaslona, J. L., Wu, L., Hughes, M., Johnston, B., Dyer, C., Drane, M., & Glover, M. (2024). Managing the challenge of fatigue for pilots operating ultra-long range flights. Frontiers in Environmental Health, 2, 1329203.

Title 14, Code of Federal Regulations, section 117.3, Definitions.

7. Where the body clock actually is

Light sets the clock. Nothing else comes close. A pilot flying west through daylight for fourteen hours has given his clock a signal whether he meant to or not, and a pilot holding home time in a hotel room has given it a different one. The old way to model this was a flat rate: so many minutes a day of drift while anchored, so many while adapting. The engine now does what the physiology does. It takes the light the pilot actually sees, hour by hour, and runs it through a published model of the human pacemaker to find where the clock is on each night of the trip.

Melatonin is the one other tool that moves the clock by a useful amount, and only in the advance direction at a usable time of day. Exercise moves it a little. Food moves the stomach's clock and not the brain's. Caffeine does not move the clock at all; it hides the pressure for a while and hands the bill back later. The engine treats each of them at the strength the measurements give it, in that order.

7.1 The pacemaker and its period

Czeisler et al. (1999) measured the intrinsic period of the human circadian pacemaker at 24.18 hours under forced desynchrony, with far less scatter than earlier free-running studies had suggested. A clock that runs eleven minutes long every day needs a daily correction, and light supplies it. Khalsa, Jewett, Cajochen, and Czeisler (2003) measured the phase response curve to a single bright light pulse: light in the hours before the temperature minimum delays the clock, light after it advances the clock, and the crossover sits near the minimum itself. Zeitzer, Dijk, Kronauer, Brown, and Czeisler (2000) measured the dose response: half the maximal resetting at roughly 100 lux, and saturation at ordinary room brightness well short of daylight. The practical consequence is that the difference between a dark room and a lit one matters a great deal, and the difference between a lit room and the sun matters less than a pilot would guess.

7.2 The model

The engine uses the light-based model of the human pacemaker of Kronauer, Forger, and Jewett (1999), in the revised form restated by St Hilaire et al. (2007) and by Serkh and Forger (2014), with every parameter taken from those restatements. Process L converts light into a drive on the clock through a photoreceptor pool that activates and recovers; Process P is a van der Pol oscillator whose state marks the temperature minimum.

```
alpha(I) = alpha0 (I / I0)^p                              alpha0 = 0.05 per min,
I0 = 9500 lux,  p = 0.5
dn/dt  = 60 [ alpha (1 − n) − beta n ]                     beta = 0.0075 per min
B      = G alpha (1 − n)(1 − 0.4 x)(1 − 0.4 xc)            G = 33.75
dx/dt  = (pi/12) [ xc + mu (x/3 + 4x^3/3 − 256x^7/105) + B ]         mu =
0.13
dxc/dt = (pi/12) [ q B xc − x ((24 / (0.99729 tau_x))^2 + k B) ]     q =
1/3,  k = 0.55,  tau_x = 24.2 h
CBTmin = each minimum of x, plus 0.97 h
```

It is integrated at one minute steps. The population period of 24.2 hours is the same number Czeisler measured, and the 0.18 hour per day drift in dim light that the model produces is the measured one. The sign of the seventh order term is inferred from the requirement of a stable limit cycle and pinned by tests that the model entrains to a 24 hour day and free-runs at its own period in the dark.

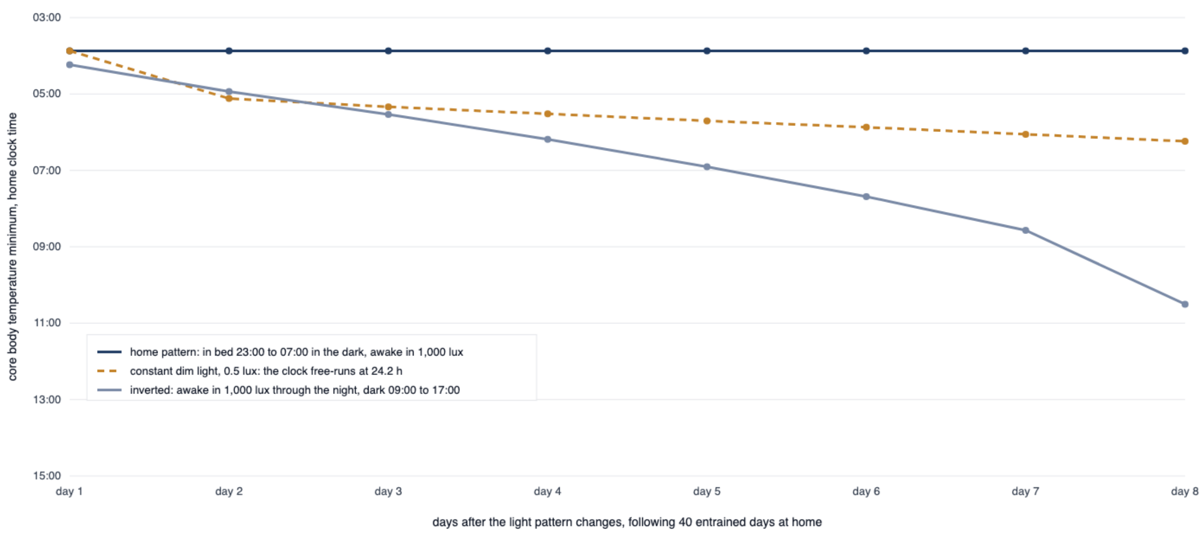

Figure 8. The pacemaker model run on three light patterns after 40 entrained days at home for a 23:00 to 07:00 sleeper, with the temperature minimum read on the home clock each day. On the home pattern it holds steady, here near 03:50, about an hour earlier than the 05:00 the engine pins from habitual wake for the wave of chapter 2, which is why the engine reads this model as a shift against its own home-held run rather than as a clock time. In constant dim light it drifts later by 0.18 hours a day after a first-day release. Awake in room light through the night and dark by day, it delays by 0.7 to 0.9 hours a day and accelerates. Every trip a pilot flies is some mixture of these.

7.3 The light a route delivers

The light history the model integrates is what the pilot does, hour by hour, in lux. In bed is dark, 0.5 lux. Awake indoors is 1,000 lux by local day and 100 by local night, or 100 inside an avoid-light window with the curtains closed. A meal or a seek-light window with the local sun up is taken outdoors at 10,000 lux, which the model treats as near saturation. A bunk break is at its facility's light: a Class 1 bunk is dark whatever the sun is doing outside, and a curtained seat is dimmed rather than dark. An awake hour on the flight deck follows the sun's elevation along the great circle, 10,000 lux by day and the dimmed-cabin value by night.

The route matters more than it looks. Computed along the great circles, a Los Angeles to Hong Kong flight of 14.7 hours departing in daylight never sees the sun set: the route climbs to 51 degrees north where the sun's westward ground track slows to the aircraft's speed, and the deck is lit for the whole flight outside the bunk. The return of 13.6 hours, flying east, compresses a full solar day into the flight: it departs in the dark, sees sunrise and sunset aloft, and lands in the dark, with 35 percent of the flight in darkness. Two flights on the same city pair deliver opposite light environments, and a model that priced them by elapsed time and direction alone put them within three minutes of each other. The lux values for the deck and the room are inferred within published bands and the model is not sensitive to them at the top end: 10,000 and 60,000 lux on the deck agree to a tenth of an hour over a trip.

7.4 Anchoring, and drift while anchored

The travel rates in the literature come from Aschoff, Hoffman, Pohl, and Wever (1975), as reported by Eastman, Gazda, Burgess, Crowley, and Fogg (2005): about 92 minutes a day of adaptation for delays after westward flights and 57 minutes a day for advances after eastward flights. The asymmetry follows from the period: a clock that runs long is already inclined to drift later. Clock offsets are also modular. A raw difference of 14 hours is a 10 hour delay, and the body takes the short way round, so on a transpacific flight the clock moves the opposite way from what the time zone arithmetic suggests. Gander et al. (2016), on 39 Boeing 747-400 pilots flying nine to thirteen day trips with one day layovers, cite an earlier transpacific sequence in which the core temperature rhythm ran at 24.6 hours across nine days: the clock drifts later across such a trip whatever the pilot intends.

Holding home time does not stop the drift; it reduces it. The engine's own measurement, on real rotations run through the light model, puts the anchored drift at roughly a quarter to a third of an hour a day on one day stays and a few hundredths on two and three day stays where the pilot sleeps on his own night in the dark. Where a plan holds home time and the pilot is awake through his body night under cabin light, the drift is larger, and the plan re-times what follows from the shifted clock rather than from the one he wishes he had.

7.5 Light windows

Anchored means advance on nearly every rotation, so the plan seeks light in the body morning and avoids it in the body evening, by Khalsa's curve: a two hour seek window beginning half an hour after the habitual wake, and a three hour avoid window before the habitual bedtime, both placed from the temperature minimum. No window is emitted when the clock is already within 0.45 hours of its target, because a window that cannot move the clock by more than its own reach is noise on a card. With age the eye admits less light at the wavelengths the clock reads: Turner and Mainster (2008) estimate that a 45 year old retains about half the circadian photoreception of early youth, and the illuminance needed to hold the clock rises with each decade. The engine's response is to lengthen the seek window and to tell an older pilot to get outside rather than sit by a window, since daylight and a lit room are two orders of magnitude apart.

7.6 Melatonin

Burgess, Revell, and Eastman (2008) measured the phase response curve to three daily pulses of 3 mg melatonin in 27 subjects on an ultradian dim-light schedule: a fitted maximum advance of 1.8 hours peaking 4.6 hours before the dim-light melatonin onset, a maximum delay of 1.3 hours peaking 11.1 hours after it, and a crossover about 9 hours before it. Burgess, Revell, Molina, and Eastman (2010) repeated the protocol at 0.5 mg and found shifts of the same size about two hours later. Timing beats dose. The engine carries the curve as a dual harmonic solved exactly through the five published conditions, and what falls out unforced is a crossover from advances to delays about an hour after the onset and a shallow dead zone through the first hours of sleep, both of which Breslow, Phillips, Huang, St Hilaire, and Klerman (2013) describe.

```
f(t) = a0 + a1 cos(wt) + b1 sin(wt) + a2 cos(2wt) + b2 sin(2wt),  t in hours
from the DLMO,  w = 2 pi / 24
f(-4.6) = +1.8,  f'(-4.6) = 0,  f(11.1) = -1.3,  f'(11.1) = 0,  f(-9) = 0
DLMO = habitual sleep midpoint - 7.0 h        one dose = one third of the
three-pulse curve        0.5 mg and 1 mg: the curve 2.0 h later
```

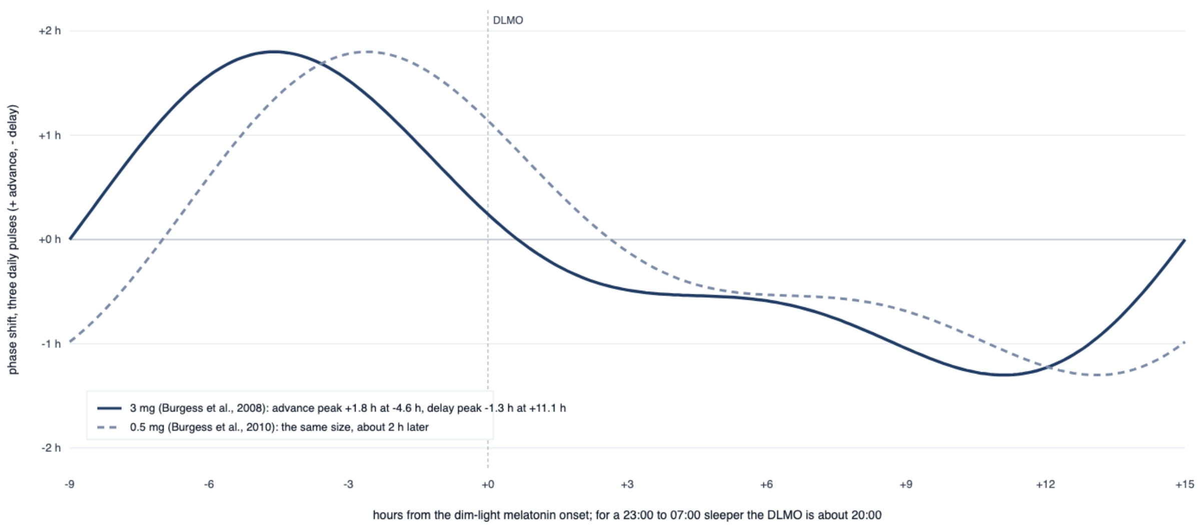

Figure 7. The melatonin phase response curve the engine uses, on Burgess's axis of hours from the dim-light melatonin onset. The 3 mg curve passes exactly through the published peaks and crossover; the 0.5 mg curve is the same size two hours later. One dose is credited one third of the three-pulse shift, an additivity assumption consistent with Lewy et al. (1998), whose four 0.5 mg doses produced at most 1.5 hours.

The dose is timed to the curve, not to sleep. A phase-timed dose is taken 5 hours before habitual bedtime, which lands it near the advance peak: Sletten et al. (2018) and Crowley and Eastman (2015) both dosed 0.5 mg at five hours before sleep in their protocols, and the pharmacopoeial advice of one to two hours before bed is the sleep-aid clock, not the phase clock. The sleep-aid effect is real but small: Ferracioli-Oda, Qawasmi, and Bloch (2013) pooled 19 trials and found onset latency shortened by 7.06 minutes, and Cruz-Sanabria et al. (2024) found the effect rising with dose to a plateau above 4 mg. The engine converts those minutes into a small lowering of the onset threshold and names the dose band rather than a milligram figure. Two rules are the pilot's, not the literature's: no dose on an operating flight, because a sedative has no place on a flight deck, and a dose that would land inside a flight or inside another planned sleep is dropped, never slid later, because a deferred dose keeps the sedation and spends the phase effect. A dose that would move the clock by less than one window's reach is withheld.

7.7 Exercise

Youngstedt, Elliott, and Kripke (2019) measured a phase response curve to one hour of moderate treadmill exercise in 101 adults, using urinary melatonin as the marker: advances of about 0.9 hours for exercise at 07:00 and 0.7 to 0.9 hours at 13:00 to 16:00, delays of 0.5 to 0.7 hours at 19:00 to 22:00, with the amplitude about a third of bright light in the same protocol. Leota et al. (2025), on 14,689 people and four million person-nights of wearable data, found no association between exercise and any sleep change when the exercise ended four or more hours before sleep onset, and delayed onset, shorter sleep, and higher resting heart rate inside that window, scaling with strain. The engine places one hour of exercise half an hour after waking, never within four hours of a night sleep, and never on an operating leg. Its phase response curve is not yet wired, by design: light is three times stronger and its curve comes first.

7.8 Meals

Wehrens et al. (2017) delayed meals by five hours for six days and found the plasma glucose rhythm delayed by 5.69 hours while melatonin and cortisol did not move. Food entrains the peripheral clocks and leaves the master clock alone, so the engine has no meal-to-phase term, and that zero is a measured one. What meals do to sleep is different. Chung, Bin, Cistulli, and Chow (2020), in 793 young adults, found eating within three hours of bedtime associated with nocturnal awakenings at an adjusted odds ratio of 1.43, and not with onset latency or duration. The engine places three meals on the body clock at 07:30, 12:30, and 19:00, keeps them three hours clear of sleep, and on an anchored trip keeps them on body time, because eating on local time drags the peripheral clocks toward a zone the master clock is deliberately resisting.

7.9 Caffeine

Wyatt, Cajochen, Ritz-De Cecco, Czeisler, and Dijk (2004) gave low hourly doses of caffeine across 29 days of forced desynchrony, the design that separates the homeostatic from the circadian contribution, and found that caffeine attenuated the wake-dependent decline in performance, was particularly effective at the circadian nadir, and extended the ability to stay awake, while increasing subjective sleepiness and holding subjects short of completing the transition to sleep. It masks pressure; it does not repay it. Gardiner et al. (2023) pooled the cost side: to avoid cutting total sleep, a cup of coffee should stop about 8.8 hours before bed and a double dose about 13.2 hours before. The engine tests every stretch the pilot is at the controls, to the bunk and after it, and plans a dose at the start of a stretch when any of three published lines is reached: the walk says pressure would cross the onset line before the stretch ends; any part of the stretch falls in the 02:00 to 06:00 window of circadian low that 14 CFR 117.3 defines, where Wyatt found caffeine particularly effective whether or not the pilot had slept; or the pressure, read as hours awake from rested through Dawson and Reid's scale, reaches the 0.04 percent flight-crew line of 14 CFR 91.17. One dose per stretch, never inside six hours of a planned sleep.

The six hour cutoff is a deliberate choice against the pooled finding, on the ordering that safety at the controls comes before the next sleep. The dose raises the onset threshold by a published 9 minutes of latency, decaying with a five hour half-life.

```
caffeine's effect on H+ = peak x 0.5^(hours / 5.0)
```

How Zeitfarer applies this

The plan is built twice. The first pass is drawn on a home-held clock to learn where the sleep blocks fall. The light history of that plan, in lux, hour by hour, runs the pacemaker model, and each night's temperature minimum is read off against the home-held clock. The second pass rebuilds the plan on the shifted clock. Melatonin enters the same model as a phase kick on the oscillator state at the instant of each dose. Light, exercise, meal, and caffeine windows are then placed on the body clock the model reports, and every conversion to local time is made last.

Sources for this chapter

Aschoff, J., Hoffmann, K., Pohl, H., & Wever, R. (1975). Re-entrainment of circadian rhythms after phase-shifts of the Zeitgeber. Chronobiologia, 2(1), 23-78.

Breslow, E. R., Phillips, A. J. K., Huang, J. M., St Hilaire, M. A., & Klerman, E. B. (2013). A mathematical model of the circadian phase-shifting effects of exogenous melatonin. Journal of Biological Rhythms, 28(1), 79-89.

Burgess, H. J., Revell, V. L., & Eastman, C. I. (2008). A three pulse phase response curve to three milligrams of melatonin in humans. The Journal of Physiology, 586(2), 639-647.

Burgess, H. J., Revell, V. L., Molina, T. A., & Eastman, C. I. (2010). Human phase response curves to three days of daily melatonin: 0.5 mg versus 3.0 mg. The Journal of Clinical Endocrinology and Metabolism, 95(7), 3325-3331.

Chung, N., Bin, Y. S., Cistulli, P. A., & Chow, C. M. (2020). Does the proximity of meals to bedtime influence the sleep of young adults? A cross-sectional survey of university students. International Journal of Environmental Research and Public Health, 17(8), 2677.

Crowley, S. J., & Eastman, C. I. (2015). Phase advancing human circadian rhythms with morning bright light, afternoon melatonin, and gradually shifted sleep: Can we reduce morning bright-light duration? Sleep Medicine, 16(2), 288-297.

Cruz-Sanabria, F., Bruno, S., Crippa, A., Frumento, P., Scarselli, M., Skene, D. J., & Faraguna, U. (2024). Optimizing the time and dose of melatonin as a sleep-promoting drug: A systematic review of randomized controlled trials and dose-response meta-analysis. Journal of Pineal Research, 76(5), e12985.

Czeisler, C. A., Duffy, J. F., Shanahan, T. L., Brown, E. N., Mitchell, J. F., Rimmer, D. W., Ronda, J. M., Silva, E. J., Allan, J. S., Emens, J. S., Dijk, D.-J., & Kronauer, R. E. (1999). Stability, precision, and near-24-hour period of the human circadian pacemaker. Science, 284(5423), 2177-2181.

Eastman, C. I., Gazda, C. J., Burgess, H. J., Crowley, S. J., & Fogg, L. F. (2005). Advancing circadian rhythms before eastward flight: A strategy to prevent or reduce jet lag. Sleep, 28(1), 33-44.

Ferracioli-Oda, E., Qawasmi, A., & Bloch, M. H. (2013). Meta-analysis: Melatonin for the treatment of primary sleep disorders. PLoS ONE, 8(5), e63773.

Gander, P., Mulrine, H. M., van den Berg, M. J., Wu, L., Smith, A., Signal, L., & Mangie, J. (2016). Does the circadian clock drift when pilots fly multiple transpacific flights with 1- to 2-day layovers? Chronobiology International, 33(8), 982-994.

Gardiner, C., Weakley, J., Burke, L. M., Roach, G. D., Sargent, C., Maniar, N., Townshend, A., & Halson, S. L. (2023). The effect of caffeine on subsequent sleep: A systematic review and meta-analysis. Sleep Medicine Reviews, 69, 101764.

Khalsa, S. B. S., Jewett, M. E., Cajochen, C., & Czeisler, C. A. (2003). A phase response curve to single bright light pulses in human subjects. The Journal of Physiology, 549(3), 945-952.

Kronauer, R. E., Forger, D. B., & Jewett, M. E. (1999). Quantifying human circadian pacemaker response to brief, extended, and repeated light stimuli over the phototopic range. Journal of Biological Rhythms, 14(6), 500-515.

Leota, J., Presby, D. M., Le, F., Czeisler, M. E., Mascaro, L., Capodilupo, E. R., Wiley, J. F., Drummond, S. P. A., Rajaratnam, S. M. W., & Facer-Childs, E. R. (2025). Dose-response relationship between evening exercise and sleep. Nature Communications, 16.

Lewy, A. J., Bauer, V. K., Ahmed, S., Thomas, K. H., Cutler, N. L., Singer, C. M., Moffit, M. T., & Sack, R. L. (1998). The human phase response curve (PRC) to melatonin is about 12 hours out of phase with the PRC to light. Chronobiology International, 15(1), 71-83.

Serkh, K., & Forger, D. B. (2014). Optimal schedules of light exposure for rapidly correcting circadian misalignment. PLoS Computational Biology, 10(4), e1003523.

Sletten, T. L., Magee, M., Murray, J. M., Gordon, C. J., Lovato, N., Kennaway, D. J., Gwini, S. M., Bartlett, D. J., Lockley, S. W., Lack, L. C., Grunstein, R. R., & Rajaratnam, S. M. W. (2018). Efficacy of melatonin with behavioural sleep-wake scheduling for delayed sleep-wake phase disorder: A double-blind, randomised clinical trial. PLoS Medicine, 15(6), e1002587.

St Hilaire, M. A., Klerman, E. B., Khalsa, S. B. S., Wright, K. P., Czeisler, C. A., & Kronauer, R. E. (2007). Addition of a non-photic component to a light-based mathematical model of the human circadian pacemaker. Journal of Theoretical Biology, 247(4), 583-599.

Turner, P. L., & Mainster, M. A. (2008). Circadian photoreception: Ageing and the eye's important role in systemic health. British Journal of Ophthalmology, 92(11), 1439-1444.

Wehrens, S. M. T., Christou, S., Isherwood, C., Middleton, B., Gibbs, M. A., Archer, S. N., Skene, D. J., & Johnston, J. D. (2017). Meal timing regulates the human circadian system. Current Biology, 27(12), 1768-1775.

Wyatt, J. K., Cajochen, C., Ritz-De Cecco, A., Czeisler, C. A., & Dijk, D.-J. (2004). Low-dose repeated caffeine administration for circadian-phase-dependent performance degradation during extended wakefulness. Sleep, 27(3), 374-381.

Youngstedt, S. D., Elliott, J. A., & Kripke, D. F. (2019). Human circadian phase-response curves for exercise. The Journal of Physiology, 597(8), 2253-2268.

Zeitzer, J. M., Dijk, D.-J., Kronauer, R., Brown, E., & Czeisler, C. (2000). Sensitivity of the human circadian pacemaker to nocturnal light: Melatonin phase resetting and suppression. The Journal of Physiology, 526(3), 695-702.

8. Cumulative loss and recovery

One bad night is a problem for one day. A trip is several days, and the cost of short sleep adds up in a way that a pilot cannot feel. In the best experiment on the question, people restricted to six hours in bed for two weeks performed as badly as people kept up all night once, and rated themselves only mildly sleepy while they did it. That is the reason the app carries a warning that runs across days, and the reason it does not trust how you feel to decide whether the warning is right.

Recovery is quicker than people fear and slower than they hope. The body sleeps long on the first night home and is back at its own baseline on the second. Nothing can be banked ahead: a night longer than you need before a trip does not buy you sleep on the trip, because a rested body simply sleeps less. The model counts hours awake beyond a daily limit, pays them back when a night runs under it, and never lets the total go below zero.

8.1 Excess wakefulness, not sleep debt

Van Dongen, Maislin, Mullington, and Dinges (2003) restricted healthy young adults to 4, 6, or 8 hours in bed for 14 nights and compared them with three nights of total sleep deprivation, across 830 subject-days in the laboratory. Fourteen nights at 4 hours produced lapses and working-memory deficits equivalent to two nights without sleep; fourteen at 6 hours, to one. Expressed as cumulative sleep debt, chronic restriction and total deprivation looked like different phenomena. Expressed as cumulative excess wakefulness, hours awake beyond a critical duration, the two collapsed onto a single near-linear function. The critical waking duration was 15.84 hours per day, with a standard error of 0.73 hours on the mean and a standard deviation of 3.58 hours between people. Subjective sleepiness saturated while performance kept falling; the subjects were, in the authors' words, largely unaware of the deficits. The paper also refutes, in its own words, the two-process prediction that homeostatic compensation would allow chronic reduction of sleep without cumulative cost, which is why the within-day model of chapters 2 and 3 is not asked to answer the multi-day question.

The referent is time out of bed. The experiment ran on strict time-in-bed schedules, so a person on 8 hours in bed was awake 16 hours against a limit of 15.84, which is why even that group accumulated a little. Feeding a wearable's measured sleep into the same limit would tell a normal sleeper he is accumulating a deficit in his own bed every night.

8.2 The chronic warning

The engine sums excess wakefulness across rolling 24 hour days from the pilot's home wake, in the paper's own form, with two rules the pilot set. A completed day contributes its hours awake minus 15.84 with its sign, so a day under the line pays back what an earlier day accumulated. The running total is floored at zero after each day, so a long night clears a deficit and never banks ahead. The unfinished day contributes only its positive excess so far. A nap is time in bed on its own day and cannot launder a long day. One night's worth of excess is 24 minus 15.84, 8.16 hours, and that is the line the warning speaks to.

```
day term = hours awake − 15.84, with its sign        running total = max(0,
running total + day term)
```

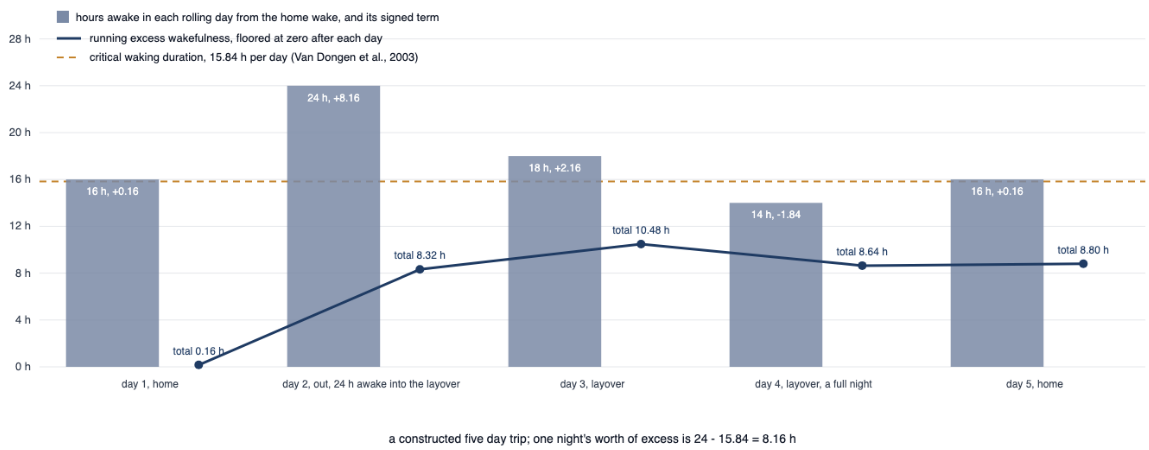

Figure 9. Cumulative excess wakefulness on a constructed five day trip. The bars are hours awake in each rolling day from the home wake; the line is the running total, each day's signed term added and the sum floored at zero. A 24 hour first duty day puts one night's worth of excess on the books, a full night on the layover pays some of it back, and nothing is ever banked ahead.

The one-for-one repayment is the paper's linear form, not a measured recovery rate. Banks, Van Dongen, Maislin, and Dinges (2010) restricted subjects to 4 hours for five nights and then gave a single recovery night of 0 to 10 hours in bed: every outcome improved with the dose, and 10 hours in bed still left lapses above baseline. Repayment is somewhat slower than one to one, and a rate would be an inference the engine does not make. The warning is read as a trend rather than a threshold, on a ruling made for a statistical reason: the individual spread in the critical duration is wide, an error in it compounds per day, and a message that says this trip is accumulating faster than your own baseline survives a wrong number where a message that says you have crossed three hours does not.

8.3 Age, need, and how much a body will sleep

Klerman and Dijk (2008) gave younger and older adults 16 hours a day of sleep opportunity for up to a week and measured the asymptote: 8.9 hours a day at ages 18 to 32, 7.4 hours at 60 to 76. The authors lean to reduced need rather than reduced ability, because alertness improved similarly in both groups. Two consequences follow for a pilot in his sixties. His critical waking duration is longer, not shorter, since sleeping less is not being deprived when the need is lower; and a recovery night at his age cannot return what a younger pilot's can, since the asymptote caps it. Van Dongen's 15.84 hours came from subjects in their late twenties, so applied to an older pilot it overstates his excess. Reading the warning as a trend against his own baseline cancels that offset, which is the second reason the threshold form was never shipped.

Kitamura et al. (2016) separated debt from need by giving fifteen young men a 12 hour opportunity for nine consecutive nights. Habitual sleep at home was 7.37 hours; the optimal duration, the asymptote, was 8.41; the difference, 1.04 hours of potential debt, bought 10.59 hours of sleep on the first night, and clearing one hour of debt took four days. The numbers are a young man's and are an upper bound for an older pilot, not an estimate. What they show is that the amount a body sleeps in a day is a function of debt, and the engine already carries that feedback: pressure discharges in proportion to itself, so eight hours asleep from a high pressure discharges three times what it does from a low one. What stays flat is the credit, the efficiency of chapter 5; how much debt, how fast it discharges, and how much of time in bed is sleep are three quantities and the engine keeps them apart.

8.4 Recovery after a trip, in pilots

Four datasets agree on the shape. Zaslona (2016), on Delta Boeing 767 crews returning from a Lagos rotation, found 473 minutes of sleep on the first post-trip day against a 401 minute baseline, 72 minutes more, and days two and three back at baseline with efficiency unchanged. Signal et al. (2024) found 8.4 hours in the first 24 hours home. Rempe et al. (2025) found post-trip sleep of 7.70 hours against 7.48 before the trip. The first night is extended and the second is ordinary, which is the pattern the recovery card of the app is built to: a long first night on the habitual wake, converging on the habitual bedtime over the following nights.

The rigorous model of the multi-day process exists and is not scheduled. McCauley et al. (2013) add a slow state variable that modulates the homeostatic equilibrium over days to weeks and produce a bifurcation: below a critical daily wakefulness, deficits converge; above it they escalate. They also find that the circadian timing of prior sleep, not only its duration, moves the equilibrium, which is the long-haul case exactly. The excess-wakefulness count is the implementable form of the same finding.

8.5 Why sleep cannot be banked

A pilot on a back-to-back roster asks whether an extra night before a trip can be stored. The model says no, and the physiology agrees on two grounds. The running total of section 8.2 floors at zero, so a surplus has nowhere to go. And a rested body sleeps less, not more: Klerman and Dijk attribute the extra sleep of the first extended nights to the presence of debt, and Kitamura's subjects slept 10.6 hours on night one and declined toward 8.4 as the debt cleared. What a good night before a trip buys is a low starting pressure, which the walk of chapter 3 carries forward as a lower S, and that is real but it is not a reserve.

8.6 Split sleep

Roach et al. (2017) compared healthy men living on consolidated and split sleep-wake schedules under forced desynchrony and found the split schedule not harmful and possibly beneficial: slightly more sleep per cycle, more slow wave sleep, and longer onset latency on every split sleep. Lo et al. (2019), in sleep-restricted adolescents, found a night plus an afternoon nap gave better vigilance, working memory, and processing speed than a continuous night of slightly more sleep, with a metabolic cost in glucose regulation. Two blocks are well studied and defensible. Three blocks in eighteen hours are not tested anywhere, and the one consistent cost of splitting, an onset latency paid on every block, is the mechanism behind the engine's settle-in overhead of chapter 6. The placement rule of chapter 4 produces two blocks on most layovers and never plans the third.

How Zeitfarer applies this

The chronic warning runs on the plan as well as on what was logged, so it can say what a pilot will have accumulated by the report on day four before he flies it. It is a note on the card, worded as a trend, and it never filters a plan. The recovery card plans the long first night on the habitual wake and the ordinary second night after it. Nothing in the engine offers to bank sleep, and the rolling window of chapter 9 is the only quantity that looks back.

Sources for this chapter

Banks, S., Van Dongen, H. P. A., Maislin, G., & Dinges, D. F. (2010). Neurobehavioral dynamics following chronic sleep restriction: Dose-response effects of one night for recovery. Sleep, 33(8), 1013-1026.

Kitamura, S., Katayose, Y., Nakazaki, K., Motomura, Y., Oba, K., Katsunuma, R., Terasawa, Y., Enomoto, M., Moriguchi, Y., Hida, A., & Mishima, K. (2016). Estimating individual optimal sleep duration and potential sleep debt. Scientific Reports, 6, 35812.

Klerman, E. B., & Dijk, D.-J. (2008). Age-related reduction in the maximal capacity for sleep: Implications for insomnia. Current Biology, 18(15), 1118-1123.

Lo, J. C., Twan, D. C. K., Karamchedu, S., Lee, X. K., Ong, J. L., Van Rijn, E., Gooley, J. J., & Chee, M. W. L. (2019). Differential effects of split and continuous sleep on neurobehavioral function and glucose tolerance in sleep-restricted adolescents. Sleep, 42(5), zsz037.

McCauley, P., Kalachev, L. V., Mollicone, D. J., Banks, S., Dinges, D. F., & Van Dongen, H. P. A. (2013). Dynamic circadian modulation in a biomathematical model for the effects of sleep and sleep loss on waking neurobehavioral performance. Sleep, 36(12), 1987-1997.

Rempe, M. J., Rasmussen, I., Gregory, K., Johnson, C., Hsin, M., Flynn-Evans, E. E., Lamp, A., & Hilditch, C. J. (2025). Layover start timing predicts layover sleep quantity and timing on long-range and ultra-long-range trips. Sleep Advances, 6(1), zpaf002.

Roach, G. D., Zhou, X., Darwent, D., Kosmadopoulos, A., Dawson, D., & Sargent, C. (2017). Are two halves better than one whole? A comparison of the amount and quality of sleep obtained by healthy adult males living on split and consolidated sleep-wake schedules. Accident Analysis and Prevention, 99(Pt B), 428-433.

Signal, T. L., van den Berg, M. J., Zaslona, J. L., Wu, L., Hughes, M., Johnston, B., Dyer, C., Drane, M., & Glover, M. (2024). Managing the challenge of fatigue for pilots operating ultra-long range flights. Frontiers in Environmental Health, 2, 1329203.

Van Dongen, H. P. A., Maislin, G., Mullington, J. M., & Dinges, D. F. (2003). The cumulative cost of additional wakefulness: Dose-response effects on neurobehavioral functions and sleep physiology from chronic sleep restriction and total sleep deprivation. Sleep, 26(2), 117-126.

Zaslona, J. L. (2016). In-flight sleep as a pilot fatigue mitigation on long range and ultra-long range flights [Doctoral thesis, Massey University].

9. Readiness, impairment, and the warnings

The cards show a pilot three kinds of number about readiness: how much rest he will have before a report, how the hours awake compare with a scale people recognize, and whether a leg crosses the hours when the body is least able to fly. Behind them the engine keeps a second, continuous alertness score that it never shows. Two scales on one card would be two answers to one question, so the one that maps to a published measurement is the one on the card and the other stays inside.

None of these numbers is a verdict. Each is a flag with a source, worded so a pilot knows what was measured and what was not. The line for flight crew is on an ordinary duty day, so it is reported and never used to filter a plan, and the chronic warning of chapter 8 is a trend against the pilot's own baseline.

9.1 Hours awake as blood alcohol

Dawson and Reid (1997) matched sustained wakefulness against blood alcohol concentration on the same performance test and published three points: 17 hours awake performed like 0.05 percent, 22 like 0.08, and 24 like 0.10. The points are not on a line, so the engine keeps them as a table joined by their own two slopes. The tails are the engine's: below 17 hours the first slope is extended down and clamped at zero, which puts the zero at 8.67 hours awake and stops a short day earning a bonus; above 24 hours the last slope continues so that worse keeps scoring worse. The zero of the fitness scale used to rank candidate plans is their topmost point, 0.10 percent, chosen after the more famous 0.05 percent was tried and found to zero every candidate on a long trip, so that nothing could be ranked. The flight-crew limit of 0.04 percent from 14 CFR 91.17 lands at 15.33 hours awake, an ordinary duty day, and is reported as a flag and never used as a filter. The comparison of fatigue to alcohol is Dawson and Reid's own framing; only the substitution of the aviation limit for their driving limit is the engine's.

```
17 h = 0.05 %      22 h = 0.08 %      24 h = 0.10 %        zero at 8.67 h
0.04 % at 15.33 h
```

Mid-trip, the pressure of chapter 2 is read back into equivalent hours awake against the pilot's own rested state, so the same scale applies whether the hours since sleep were one stretch or several.

```
equivalent hours awake = −chi_w ln( (1 − S) / (1 − S_rested) )
S_rested = the habitual steady state at wake
```

The scale has one known gap, and it is stated rather than filled: Dawson and Reid's axis is hours awake alone, and no published conversion from circadian phase to a performance decrement exists in this document's sources. Where a landing falls on the body clock is reported beside the fitness figure, not folded into it.

9.2 The windows of circadian low

FAA Advisory Circular 120-103A defines the window of circadian low as roughly 02:00 to 06:00 on the body clock and names a second, less severe window roughly 15:00 to 17:00; Advisory Circular 120-100 gives the same two on a twelve hour clock. Neither cites a study, and the engine treats them as the regulator's summary of the field. A leg whose flight time overlaps the night window by twenty minutes or more carries the warning on its card with the number of hours inside the trough. The afternoon window is not a warning: it is the afternoon zone of chapter 4, where the plan will put a short sleep if there is one to be had.

9.3 The rolling window

The Washington State University group that reviewed the model in August 2026 gave one criterion above the others: at any point, a pilot should be able to look back over the previous 24 hours and see seven to nine hours of actual sleep, not time in bed. The engine carries it as a continuous term, credited sleep in the trailing 24 hours sampled every 15 minutes against a floor of 7 hours, and reports it at 5, 7, 8, and 9 hours. It is the term that forbids a whole class of defect: a plan that skipped an entire body night because a candidate ten hours later scored one and a half percent better, which a search once produced, cannot survive a rolling window.

```
lookback(t) = credited sleep in [t − 24 h, t]        deficit(t) = max(0, 7.0 − lookback(t))
```

9.4 The internal score

The Sleep, Activity, Fatigue, and Task Effectiveness model, SAFTE, is the alertness model behind the FAST tool used in aviation and rail fatigue risk management (Hursh et al., 2004), and its core is disclosed in an expired United States patent, 6,579,233, which is public domain. The engine computes SAFTE effectiveness as an internal scorer and never renders it. A sleep reservoir of 2,880 units depletes at 0.5 units a minute awake; effectiveness is the reservoir ratio plus a circadian term whose amplitude grows with the deficit, 7 percent rested rising by 5 percent of the deficit, with a 24 hour component peaking 12 hours after the temperature minimum and a 12 hour component at half the size leading it by three hours. The identification with Process S is exact because the reservoir ratio is dimensionless: one minus S. Sleep inertia takes SAFTE's form too, a 5 percent decrement decaying at 0.04 per minute over the first two hours after waking.

```
E = 100 (1 − S) + a_p c + I        a_p = 7.0 + 5.0 S        I(t) = −5.0 exp(−0.04 t) for t < 120 min
```

The check that the identification is more than a guess: SAFTE's disclosed default acrophase is 17:00, and the engine's acrophase, derived from the pilot's habitual window rather than fixed, lands on exactly 17:00 for the default 23:00 to 07:00 profile. Two conventions agreeing on a number neither was fitted to is the argument.

9.5 The warnings, in one table

Warning	What it means	Source
Rest before duty, Short	credited sleep in the 24 hours before pushback is under 7 hours in bed	the rolling criterion
Crosses your WOCL	twenty minutes or more of the leg fall in the 02:00 to 06:00 body-clock trough	FAA AC 120-100, AC 120-103A
Fitness, hours awake	hours since rested at report, on the blood-alcohol equivalence	Dawson and Reid (1997)
Flight-crew line	equivalent hours awake past the 0.04 percent limit, 15.33 hours	14 CFR 91.17(a)(4)

Chronic warning	excess wakefulness accumulating across the trip, as a trend	Van Dongen et al. (2003)
Readiness caution	the full break cannot fix the landing; the bunk block carries the guidance	chapter 6
Age not set	the plan is built for a 40 year old until it is	chapter 5

Regulatory rest is read by a compliance check and not by the planner: 14 CFR 117.25(e) requires a rest period of at least 10 hours with an 8 hour sleep opportunity, and the pilot's own contract nine hours behind the door. Neither is a physiological number, and the engine does not treat them as one.

How Zeitfarer applies this

The card shows Dawson and Reid; the scorer uses SAFTE; the two never appear together. The flight-crew line reports and never filters. The rolling window and the chronic warning look back over what was logged and forward over what is planned, and both are worded as what was measured rather than as a judgment of the pilot.

Sources for this chapter

Dawson, D., & Reid, K. (1997). Fatigue, alcohol and performance impairment. Nature, 388(6639), 235.

Federal Aviation Administration. (2010). Basics of aviation fatigue (Advisory Circular 120-100).

Federal Aviation Administration. (2013). Fatigue risk management systems for aviation safety (Advisory Circular 120-103A).

Hursh, S. R., Redmond, D. P., Johnson, M. L., Thorne, D. R., Belenky, G., Balkin, T. J., Storm, W. F., Miller, J. C., & Eddy, D. R. (2004). Fatigue models for applied research in warfighting. Aviation, Space, and Environmental Medicine, 75(3 Suppl), A44-A53.

Hursh, S. R. (2003). U.S. Patent No. 6,579,233. Washington, DC: U.S. Patent and Trademark Office.

Title 14, Code of Federal Regulations, section 91.17, Alcohol or drugs.

Title 14, Code of Federal Regulations, section 117.25, Rest period.

Van Dongen, H. P. A., Maislin, G., Mullington, J. M., & Dinges, D. F. (2003). The cumulative cost of additional wakefulness: Dose-response effects on neurobehavioral functions and sleep physiology from chronic sleep restriction and total sleep deprivation. Sleep, 26(2), 117-126.

10. What the model does not know

A model that claims to know everything is not worth reading. This chapter lists what the engine does not know, what it does instead, and what would settle each question. Most of the gaps are gaps in the literature rather than in the engine, and several of them could be closed by data that already exists in a laboratory's files and has not been published in the form a model needs.

- The hotel settle time. No study measures the interval from the hotel door to sleep onset in any population. The only bound in print is a 1991 NASA diary figure of 3.97 hours from end of duty to first sleep on 24 hour layovers, which includes the ride and a meal. The engine carries a one hour default until the pilot enters his own unwind time, and that default is a placeholder by design, not a claim.
- The first-night effect. Ding, Chen, Dai, and Li (2022) pooled 53 polysomnographic studies and found every continuity measure worse on the first night in a new place, non-linearly with age, but the pooled sizes are behind a paywall and no aircrew dataset isolates the effect in a hotel. The engine models none of it.
- Bunk sleep architecture. Slow wave sleep in flight had a median of 0.5 percent in the polysomnography of chapter 6. The engine credits bunk sleep by duration and continuity and does not price its depth.
- The Class 3 seat. No poll and no in-flight study covers a reclining cabin seat for pilots; its rates are scaled from Class 2 by the regulation's wording. Laboratory nap data on seat angle exist and are not applied.
- The merge gap and the third block. No aviation paper prints a rule for when two sleeps are one rest period, and three sleeps in eighteen hours are untested in anything found. The engine's three hour gap is inferred and its placement rule avoids the third block rather than pricing it.
- In-seat rest under Part 121. The cockpit nap literature is good and the practice is not permitted to the pilots this app serves. The engine offers no in-seat sleep on an operating leg and only records it on a passenger leg.
- A pilot-specific efficiency level. The curve of chapter 5 is anchored to laboratory polysomnography; the pilot actigraphy runs about two points lower on a different referent. No polysomnographic efficiency by age within pilots exists, and the closed dataset of 332 long-haul pilots does not print it.
- Recovery at 50 to 65. The one study of recovery after restriction in that band, Stenuit and Kerkhofs (2005), prints its recovery table behind a paywall, and the engine's one-for-one repayment is the linear form of a younger population's data.
- The onset penalty by facility class, and stress. The penalty is inferred from the regulation's wording, and the polysomnography found no onset difference between a Class 1 bunk and a hotel. Pre-sleep stress lengthens onset by 8 to 16 minutes in the laboratory and the engine carries no stress term, on a ruling that the room-to-bed gap is process, not physiology.
- The afternoon zone's width and depth, the placement of the temperature minimum two hours before habitual wake, the light-window durations, the lux of a lit room and a flight deck, and the dose fraction of a single melatonin tablet are all inferred within published bands and say so in Appendix A.

- The exercise phase response curve is measured and not wired, by the ordering rule that light's curve comes first. The melatonin sleep-aid effect has a dose-response shape but no per-dose values in print, so the dose is a display band and not a term.
- Everything one pilot's nights cannot calibrate. The author's own trips, wearable, and polls of his fleet have validated the model at several points and calibrated it at none. A constant is a population number. One pilot can show a prediction wrong; he cannot make it right.

What would settle the most of these at once is the dataset the Washington State University group already holds: 256 pilots, 473 trips, and nine routes of actigraphy on the very rotations this model plans. The questions to put to it are written and waiting.

Sources for this chapter

Ding, L., Chen, B., Dai, Y., & Li, Y. (2022). A meta-analysis of the first-night effect in healthy individuals for the full age spectrum. Sleep Medicine, 89, 159-165.

Gander, P. H., Graeber, R. C., Connell, L. J., & Gregory, K. B. (1991). Crew factors in flight operations VIII: Factors influencing sleep timing and subjective sleep quality in commercial long-haul flight crews (NASA Technical Memorandum 103852).

Stenuit, P., & Kerkhofs, M. (2005). Age modulates the effects of sleep restriction in women. Sleep, 28(10), 1283-1288.

References

Every source cited in the chapters, in APA 7 form.

Achermann, P., & Borbély, A. A. (1992). Combining different models of sleep regulation. Journal of Sleep Research, 1(2), 144-147.

Ackermann, S., Cordi, M., La Marca, R., Seifritz, E., & Rasch, B. (2019). Psychosocial stress before a nap increases sleep latency and decreases early slow-wave activity. Frontiers in Psychology, 10, 20.

Arnardóttir, H., Þorsteinsson, H., & Karlsson, K. Æ. (2010). Dynamics of sleep-wake cyclicity at night across the human lifespan. Frontiers in Neurology, 1, 156.

Aschoff, J., Hoffmann, K., Pohl, H., & Wever, R. (1975). Re-entrainment of circadian rhythms after phase-shifts of the Zeitgeber. Chronobiologia, 2(1), 23-78.

Banks, S., Van Dongen, H. P. A., Maislin, G., & Dinges, D. F. (2010). Neurobehavioral dynamics following chronic sleep restriction: Dose-response effects of one night for recovery. Sleep, 33(8), 1013-1026.

Bes, F. (2025). Treatise on an alternative perspective on the two-process model of sleep regulation. npj Biological Timing and Sleep, 2, 23.

Bes, F., Jobert, M., & Schulz, H. (2009). Modeling napping, post-lunch dip, and other variations in human sleep propensity. Sleep, 32(3), 392-398.

Borbély, A. A., & Achermann, P. (1999). Sleep homeostasis and models of sleep regulation. Journal of Biological Rhythms, 14(6), 557-568.

Borbély, A. A., Daan, S., Wirz-Justice, A., & Deboer, T. (2016). The two-process model of sleep regulation: A reappraisal. Journal of Sleep Research, 25(2), 131-143.

Boulos, M. I., Jairam, T., Kendzerska, T., Im, J., Mekhael, A., & Murray, B. J. (2019). Normal polysomnography parameters in healthy adults: A systematic review and meta-analysis. The Lancet Respiratory Medicine, 7(6), 533-543.

Breslow, E. R., Phillips, A. J. K., Huang, J. M., St Hilaire, M. A., & Klerman, E. B. (2013). A mathematical model of the circadian phase-shifting effects of exogenous melatonin. Journal of Biological Rhythms, 28(1), 79-89.

Burgess, H. J., Revell, V. L., & Eastman, C. I. (2008). A three pulse phase response curve to three milligrams of melatonin in humans. The Journal of Physiology, 586(2), 639-647.

Burgess, H. J., Revell, V. L., Molina, T. A., & Eastman, C. I. (2010). Human phase response curves to three days of daily melatonin: 0.5 mg versus 3.0 mg. The Journal of Clinical Endocrinology and Metabolism, 95(7), 3325-3331.

Chung, N., Bin, Y. S., Cistulli, P. A., & Chow, C. M. (2020). Does the proximity of meals to bedtime influence the sleep of young adults? A cross-sectional survey of university students. International Journal of Environmental Research and Public Health, 17(8), 2677.

Crowley, S. J., & Eastman, C. I. (2015). Phase advancing human circadian rhythms with morning bright light, afternoon melatonin, and gradually shifted sleep: Can we reduce morning bright-light duration? Sleep Medicine, 16(2), 288-297.

Cruz-Sanabria, F., Bruno, S., Crippa, A., Frumento, P., Scarselli, M., Skene, D. J., & Faraguna, U. (2024). Optimizing the time and dose of melatonin as a sleep-promoting drug: A systematic review of randomized controlled trials and dose-response meta-analysis. Journal of Pineal Research, 76(5), e12985.

Czeisler, C. A., Duffy, J. F., Shanahan, T. L., Brown, E. N., Mitchell, J. F., Rimmer, D. W., Ronda, J. M., Silva, E. J., Allan, J. S., Emens, J. S., Dijk, D.-J., & Kronauer, R. E. (1999). Stability, precision, and near-24-hour period of the human circadian pacemaker. Science, 284(5423), 2177-2181.

Daan, S., Beersma, D. G. M., & Borbély, A. A. (1984). Timing of human sleep: Recovery process gated by a circadian pacemaker. American Journal of Physiology, 246(2), R161-R183.

Dawson, D., & Reid, K. (1997). Fatigue, alcohol and performance impairment. Nature, 388(6639), 235.

de Zeeuw, J., Wisniewski, S., Papakonstantinou, A., Bes, F., Wahnschaffe, A., Zaleska, M., Kunz, D., & Münch, M. (2018). The alerting effect of the wake maintenance zone during 40 hours of sleep deprivation. Scientific Reports, 8, 11012.

Devine, J. K., Choynowski, J., Burke, T., Carlsson, K., Capaldi, V. F., McKeon, A. B., & Sowden, W. J. (2020). Practice parameters for the use of actigraphy in the military operational context: The Walter Reed Army Institute of Research Operational Research Kit-Actigraphy (WORK-A). Military Medical Research, 7(1), 31.

Dijk, D.-J., & Czeisler, C. A. (1994). Paradoxical timing of the circadian rhythm of sleep propensity serves to consolidate sleep and wakefulness in humans. Neuroscience Letters, 166(1), 63-68.

Ding, L., Chen, B., Dai, Y., & Li, Y. (2022). A meta-analysis of the first-night effect in healthy individuals for the full age spectrum. Sleep Medicine, 89, 159-165.

Dinges, D. F., Orne, M. T., Whitehouse, W. G., & Orne, E. C. (1987). Temporal placement of a nap for alertness: Contributions of circadian phase and prior wakefulness. Sleep, 10(4), 313-329.

Eastman, C. I., Gazda, C. J., Burgess, H. J., Crowley, S. J., & Fogg, L. F. (2005). Advancing circadian rhythms before eastward flight: A strategy to prevent or reduce jet lag. Sleep, 28(1), 33-44.

Federal Aviation Administration. (2010). Basics of aviation fatigue (Advisory Circular 120-100).

Federal Aviation Administration. (2013). Fatigue risk management systems for aviation safety (Advisory Circular 120-103A).

Ferracioli-Oda, E., Qawasmi, A., & Bloch, M. H. (2013). Meta-analysis: Melatonin for the treatment of primary sleep disorders. PLoS ONE, 8(5), e63773.

Gander, P. H., Graeber, R. C., Connell, L. J., & Gregory, K. B. (1991). Crew factors in flight operations VIII: Factors influencing sleep timing and subjective sleep quality in commercial long-haul flight crews (NASA Technical Memorandum 103852).

Gander, P. H., Signal, T. L., van den Berg, M. J., Mulrine, H. M., Jay, S. M., & Mangie, J. (2013). In-flight sleep, pilot fatigue and Psychomotor Vigilance Task performance on ultra-long range versus long range flights. Journal of Sleep Research, 22(6), 697-706.

Gander, P., Mulrine, H. M., van den Berg, M. J., Wu, L., Smith, A., Signal, L., & Mangie, J. (2016). Does the circadian clock drift when pilots fly multiple transpacific flights with 1- to 2-day layovers? Chronobiology International, 33(8), 982-994.

Gardiner, C., Weakley, J., Burke, L. M., Roach, G. D., Sargent, C., Maniar, N., Townshend, A., & Halson, S. L. (2023). The effect of caffeine on subsequent sleep: A systematic review and meta-analysis. Sleep Medicine Reviews, 69, 101764.

Hursh, S. R. (2003). U.S. Patent No. 6,579,233. Washington, DC: U.S. Patent and Trademark Office.

Hursh, S. R., Redmond, D. P., Johnson, M. L., Thorne, D. R., Belenky, G., Balkin, T. J., Storm, W. F., Miller, J. C., & Eddy, D. R. (2004). Fatigue models for applied research in warfighting. Aviation, Space, and Environmental Medicine, 75(3 Suppl), A44-A53.

Ingre, M., Van Leeuwen, W., Klemets, T., Ullvetter, C., Hough, S., Kecklund, G., Karlsson, D., & Åkerstedt, T. (2014). Validating and extending the three process model of alertness in airline operations. PLoS ONE, 9(10), e108679.

Khalsa, S. B. S., Jewett, M. E., Cajochen, C., & Czeisler, C. A. (2003). A phase response curve to single bright light pulses in human subjects. The Journal of Physiology, 549(3), 945-952.

Kitamura, S., Katayose, Y., Nakazaki, K., Motomura, Y., Oba, K., Katsunuma, R., Terasawa, Y., Enomoto, M., Moriguchi, Y., Hida, A., & Mishima, K. (2016). Estimating individual optimal sleep duration and potential sleep debt. Scientific Reports, 6, 35812.

Klerman, E. B., & Dijk, D.-J. (2008). Age-related reduction in the maximal capacity for sleep: Implications for insomnia. Current Biology, 18(15), 1118-1123.

Kronauer, R. E., Forger, D. B., & Jewett, M. E. (1999). Quantifying human circadian pacemaker response to brief, extended, and repeated light stimuli over the phototopic range. Journal of Biological Rhythms, 14(6), 500-515.

Lavie, P. (1986). Ultrashort sleep-waking schedule. III. 'Gates' and 'forbidden zones' for sleep. Electroencephalography and Clinical Neurophysiology, 63(5), 414-425.

Leota, J., Presby, D. M., Le, F., Czeisler, M. E., Mascaro, L., Capodilupo, E. R., Wiley, J. F., Drummond, S. P. A., Rajaratnam, S. M. W., & Facer-Childs, E. R. (2025). Dose-response relationship between evening exercise and sleep. Nature Communications, 16.

Lewy, A. J., Bauer, V. K., Ahmed, S., Thomas, K. H., Cutler, N. L., Singer, C. M., Moffit, M. T., & Sack, R. L. (1998). The human phase response curve (PRC) to melatonin is about 12 hours out of phase with the PRC to light. Chronobiology International, 15(1), 71-83.

Lo, J. C., Twan, D. C. K., Karamchedu, S., Lee, X. K., Ong, J. L., Van Rijn, E., Gooley, J. J., & Chee, M. W. L. (2019). Differential effects of split and continuous sleep on neurobehavioral function and glucose tolerance in sleep-restricted adolescents. Sleep, 42(5), zsz037.

McCauley, P., Kalachev, L. V., Mollicone, D. J., Banks, S., Dinges, D. F., & Van Dongen, H. P. A. (2013). Dynamic circadian modulation in a biomathematical model for the effects of sleep and sleep loss on waking neurobehavioral performance. Sleep, 36(12), 1987-1997.

Monterastelli, A. J., Adams, J., Eastman, C. I., & Crowley, S. J. (2024). The forbidden zone for sleep is more robust in adolescents compared to adults. Frontiers in Sleep, 2, 1304647.

Ohayon, M. M., Carskadon, M. A., Guilleminault, C., & Vitiello, M. V. (2004). Meta-analysis of quantitative sleep parameters from childhood to old age in healthy individuals: Developing normative sleep values across the human lifespan. Sleep, 27(7), 1255-1273.

Paech, G. M., Ferguson, S. A., Sargent, C., Kennaway, D. J., & Roach, G. D. (2012). The relative contributions of the homeostatic and circadian processes to sleep regulation under conditions of severe sleep restriction. Sleep, 35(7), 941-948.

Rajaraman, S., Gribok, A. V., Wesensten, N. J., Balkin, T. J., & Reifman, J. (2009). An improved methodology for individualized performance prediction of sleep-deprived individuals with the two-process model. Sleep, 32(10), 1377-1392.

Rempe, M. J., Basiarz, E., Rasmussen, I., Belenky, G., & Lamp, A. (2022). Pilot in-flight sleep during long-range and ultra-long range commercial airline flights. Aerospace Medicine and Human Performance, 93(4), 368-375.

Rempe, M. J., Rasmussen, I., Gregory, K., Johnson, C., Hsin, M., Flynn-Evans, E. E., Lamp, A., & Hilditch, C. J. (2025). Layover start timing predicts layover sleep quantity and timing on long-range and ultra-long-range trips. Sleep Advances, 6(1), zpaf002.

Roach, G. D., Darwent, D., & Dawson, D. (2010). How well do pilots sleep during long-haul flights? Ergonomics, 53(9), 1072-1075.

Roach, G. D., Zhou, X., Darwent, D., Kosmadopoulos, A., Dawson, D., & Sargent, C. (2017). Are two halves better than one whole? A comparison of the amount and quality of sleep obtained by healthy adult males living on split and consolidated sleep-wake schedules. Accident Analysis and Prevention, 99(Pt B), 428-433.

Rosekind, M. R., Graeber, R. C., Dinges, D. F., Connell, L. J., Rountree, M. S., Spinweber, C. L., & Gillen, K. A. (1994). Crew factors in flight operations IX: Effects of planned cockpit rest on crew performance and alertness in long-haul operations (NASA Technical Memorandum 108839).

Serkh, K., & Forger, D. B. (2014). Optimal schedules of light exposure for rapidly correcting circadian misalignment. PLoS Computational Biology, 10(4), e1003523.

Signal, T. L., Gale, J., & Gander, P. H. (2005). Sleep measurement in flight crew: Comparing actigraphic and subjective estimates to polysomnography. Aviation, Space, and Environmental Medicine, 76(11), 1058-1063.

Signal, T. L., Gander, P. H., van den Berg, M. J., & Graeber, R. C. (2013). In-flight sleep of flight crew during a 7-hour rest break: Implications for research and flight safety. Sleep, 36(1), 109-115.

Signal, T. L., van den Berg, M. J., Zaslona, J. L., Wu, L., Hughes, M., Johnston, B., Dyer, C., Drane, M., & Glover, M. (2024). Managing the challenge of fatigue for pilots operating ultra-long range flights. Frontiers in Environmental Health, 2, 1329203.

Skeldon, A. C., Dijk, D.-J., & Derks, G. (2014). Mathematical models for sleep-wake dynamics: Comparison of the two-process model and a mutual inhibition neuronal model. PLoS ONE, 9(8), e103877.

Skeldon, A. C., et al. (2025). The complexity and commonness of the two-process model of sleep regulation from a mathematical perspective. npj Biological Timing and Sleep.

Sletten, T. L., Magee, M., Murray, J. M., Gordon, C. J., Lovato, N., Kennaway, D. J., Gwini, S. M., Bartlett, D. J., Lockley, S. W., Lack, L. C., Grunstein, R. R., & Rajaratnam, S. M. W. (2018). Efficacy of melatonin with behavioural sleep-wake scheduling for delayed sleep-wake phase disorder: A double-blind, randomised clinical trial. PLoS Medicine, 15(6), e1002587.

Souabni, M., Souabni, M. J., Hammouda, O., Romdhani, M., Trabelsi, K., Ammar, A., & Driss, T. (2022). Benefits and risks of napping in older adults: A systematic review. Frontiers in Aging Neuroscience, 14, 1000707.

St Hilaire, M. A., Klerman, E. B., Khalsa, S. B. S., Wright, K. P., Czeisler, C. A., & Kronauer, R. E. (2007). Addition of a non-photic component to a light-based mathematical model of the human circadian pacemaker. Journal of Theoretical Biology, 247(4), 583-599.

Stenuit, P., & Kerkhofs, M. (2005). Age modulates the effects of sleep restriction in women. Sleep, 28(10), 1283-1288.

Strogatz, S. H., Kronauer, R. E., & Czeisler, C. A. (1987). Circadian pacemaker interferes with sleep onset at specific times each day: Role in insomnia. American Journal of Physiology, 253(1), R172-R178.

Title 14, Code of Federal Regulations, section 117.25, Rest period.

Title 14, Code of Federal Regulations, section 117.3, Definitions.

Title 14, Code of Federal Regulations, section 91.17, Alcohol or drugs.

Turner, P. L., & Mainster, M. A. (2008). Circadian photoreception: Ageing and the eye's important role in systemic health. British Journal of Ophthalmology, 92(11), 1439-1444.

Van Dongen, H. P. A., Maislin, G., Mullington, J. M., & Dinges, D. F. (2003). The cumulative cost of additional wakefulness: Dose-response effects on neurobehavioral functions and sleep physiology from chronic sleep restriction and total sleep deprivation. Sleep, 26(2), 117-126.

Wehrens, S. M. T., Christou, S., Isherwood, C., Middleton, B., Gibbs, M. A., Archer, S. N., Skene, D. J., & Johnston, J. D. (2017). Meal timing regulates the human circadian system. Current Biology, 27(12), 1768-1775.

Wyatt, J. K., Cajochen, C., Ritz-De Cecco, A., Czeisler, C. A., & Dijk, D.-J. (2004). Low-dose repeated caffeine administration for circadian-phase-dependent performance degradation during extended wakefulness. Sleep, 27(3), 374-381.

Youngstedt, S. D., Elliott, J. A., & Kripke, D. F. (2019). Human circadian phase-response curves for exercise. The Journal of Physiology, 597(8), 2253-2268.

Zaslona, J. L. (2016). In-flight sleep as a pilot fatigue mitigation on long range and ultra-long range flights [Doctoral thesis, Massey University].

Zeitzer, J. M., Dijk, D.-J., Kronauer, R., Brown, E., & Czeisler, C. (2000). Sensitivity of the human circadian pacemaker to nocturnal light: Melatonin phase resetting and suppression. The Journal of Physiology, 526(3), 695-702.

Appendix A. Constants quoted in this document, with their tier

Constant	Value	Tier	Where it comes from
chi_w, chi_s	18.2 h, 4.2 h	Sourced	Daan, Beersma, and Borbély (1984); Skeldon, Dijk, and Derks (2014)
mu_w, mu_s	1.0, 0.0	Sourced	the same
a	2.9 / 21.35	Sourced	Skeldon, Dijk, and Derks (2014), equation 15
harmonic coefficients	0.97, 0.22, 0.07, 0.03, 0.001	Sourced via reprint	Achermann and Borbély (1992) as reprinted by Rajaraman et al. (2009)
wave offset and normalizer	3.99 h, 1.003725987	Derived	the trough and range of the published series
CBTmin before habitual wake	2.0 h	Inferred	widely repeated; no citation recorded yet
H+0, H-0	per pilot	Derived	the fixed point of the habitual window, section 2.3
minimum onset latency	9.03 min	Sourced	Ackermann et al. (2019), control arm
sleep bout, ages 62 to 74	31.1 min	Sourced	Arnardóttir et al. (2010)
wake bout	tau_s (1 - e) / e	Derived	the alternating-renewal identity, section 3.4
walk step	2.5 min	Numerical	resolution, chosen by measurement
maximum night delay	CBTmin minus bedtime, 6.0 h	Derived, inferred bound	section 4.1
afternoon zone center	habitual wake + 6.0 h	Sourced	Bes, Jobert, and Schulz (2009); Souabni et al. (2022)
afternoon zone width and depth	2.5 h either side; 1/3 of amplitude	Inferred	ours, section 4.2
wait after a long sleep	6 h after 4 h or more	Pilot's rule; 4 h inferred	section 4.4
merge gap	3.0 h	Inferred	nearest published: Devine et al. (2020)
shortest planned sleep	70 min	Sourced, secondary	one cycle, the short end of 70 to 120 min
stay-awake window	1.5 h	Operational, the pilot's	section 4.5
efficiency slope	2.1 points per decade	Sourced	Boulos et al. (2019)
efficiency level	90 % at age 25	Inferred, validated	inside the Boulos 95 % CI at all four bands; kept by ruling
efficiency bounds	65 % to 95 %	Inferred	a physiological clamp, ours
rest window edges	40 min after departure, 60 min before arrival	Operational	section 6.3
onset penalty by class	0.10, 0.13, 0.30	Inferred	14 CFR 117.3's wording; kept by ruling
settle-in by class	20, 10, 5 min; hotel 15	Operational, the pilot's	section 6.2
arousals per hour, calm	0.15, 0.90, 1.60; hotel 0.10	Inferred	14 CFR 117.3's wording
arousals per hour, predicted	1.2, 1.9	Sourced	the two polls of 456 pilots, inverted through the model
ride rates	0.15, 0.50, 1.35, 2.15	Inferred, the pilot's	one anchor corroborated on a wearable
ensemble and step	160 runs, 30 s	Numerical	section 6.3
pacemaker parameters	alpha0 0.05, beta 0.0075, G 33.75, p 0.5, I0 9500, mu 0.13, q 1/3, k 0.55, tau_x 24.2,	Sourced	St Hilaire et al. (2007); Serkh and Forger (2014)

	0.99729, 0.97 h		
light history lux	0.5 in bed, 100 room, 1,000 awake, 10,000 outdoors and deck by day	Inferred, within published bands	section 7.3
entrainment and recovery run	40 days, 7 days	Inferred	past any transient
melatonin PRC points	+1.8 h at -4.6; -1.3 h at +11.1; crossover -9	Sourced	Burgess, Revell, and Eastman (2008)
low-dose curve shift	2.0 h later	Sourced	Burgess et al. (2010)
DLMO from sleep midpoint	7.0 h before	Sourced	Burgess et al. (2010)
dose fraction	one third of three pulses	Inferred	additivity; consistent with Lewy et al. (1998)
phase-timed dose lead	5.0 h before bedtime	Sourced	Sletten et al. (2018); Crowley and Eastman (2015)
melatonin onset benefit	7.06 min	Sourced	Ferracioli-Oda et al. (2013)
light windows	seek 2.0 h, avoid 3.0 h, first seek 0.5 h after wake, hold within 0.45 h	Inferred	durations ours; direction from Khalsa et al. (2003)
exercise	1.0 h, 0.5 h after wake, 4.0 h clearance	Inferred; clearance sourced	Leota et al. (2025)
meals	07:30, 12:30, 19:00 body; 3.0 h clear of sleep	Inferred; clearance sourced	Chung et al. (2020)
caffeine	+9.0 min latency, 5.0 h half-life; cutoff 6 h	Derived; sourced; the pilot's	Gardiner et al. (2023)
critical waking duration	15.84 h per day	Sourced	Van Dongen et al. (2003)
rolling window floor	7.0 h of sleep in any 24 h, sampled every 15 min	Sourced	the Washington State University review, August 2026
wakefulness to BAC	17 h 0.05 %, 22 h 0.08 %, 24 h 0.10 %	Sourced	Dawson and Reid (1997)
fitness scale zero	0.10 % BAC	Sourced	Dawson and Reid (1997), their topmost point
flight crew flag	0.04 % BAC	Sourced	14 CFR 91.17(a)(4)
sleep inertia	-5 % decaying at 0.04 per minute for 120 min	Sourced	SAFTE, US Patent 6,579,233
SAFTE effectiveness	2880 units, 0.5 per min, 7 % + 5 % amplitude, 12 h ratio 0.5 leading 3 h, peak 12 h after CBTmin	Sourced	US Patent 6,579,233
regulatory rest	10 h with an 8 h sleep opportunity; 9 h behind the door	Sourced; contractual	14 CFR 117.25(e); the pilot's contract

Appendix B. Notation

S is sleep pressure, C the circadian wave, phi the body-clock phase, H+ and H- the onset and offset thresholds, e the sleep efficiency, chi the time constants, mu the asymptotes.