evinlo

EVINLO EXPLAINERS

Alertness and the daily check-in, explained

How Evinlo estimates how awake you may feel today from your sleep times, and why it shows nothing until your own answers follow the model.

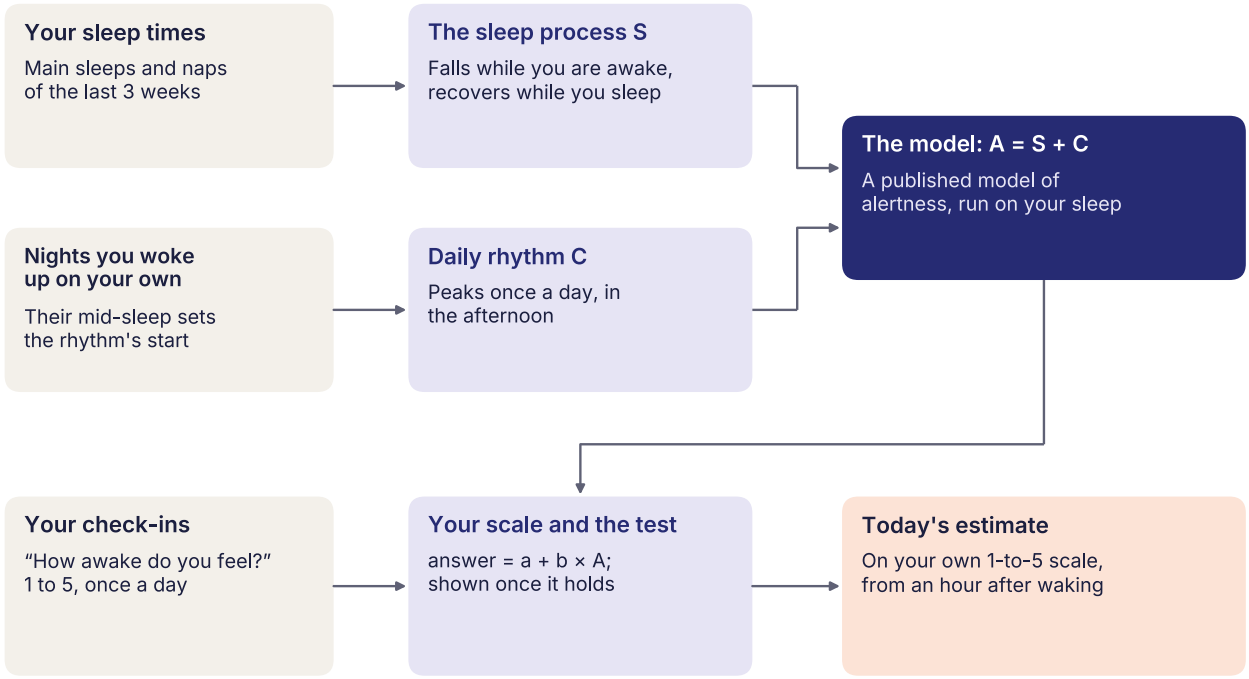


ALGORITHM
Alertness 1.3.0

WRITTEN
October 2026

The one-minute version

Once a day, Evinlo can ask you one question: **“How awake do you feel?”** Your answers teach it how a published model of alertness, run on your own sleep times, maps onto how you feel. When your answers clearly follow the model, the sleep page shows how awake you may feel through the rest of the day.



The alertness estimate in one picture. Your sleep times drive the model’s sleep process; nights you woke up on your own set where its daily rhythm starts. Your check-ins give the model your own scale, and the estimate appears only once a test shows your answers follow it.

IN SHORT

- **One tap a day.** “How awake do you feel?”, from 1 (Very sleepy) to 5 (Wide awake). Optional, and it stays on your iPhone.
- **A published model, run on your sleep.** How rested you are falls while you are awake and recovers while you sleep; a daily rhythm rises and falls on top of it.
- **Your scale, not someone else’s.** Evinlo fits a straight line from the model to your own answers.
- **Nothing until it fits.** No curve before at least 40 check-ins, and only when two tests in a row show that your answers follow the model.
- **“May” and “likely”.** It is an estimate for general wellness. It never says whether you are fit to drive or to work.

This explainer follows two synthetic people made by the app’s generator: **Maya**, who sleeps her need, and **Jonas**, who stayed up late on Monday night. Their check-ins are synthetic too: the generator draws them from the model itself, with day-to-day noise. They show how the rules work, not that the model fits real people.

What goes in

What	Where it comes from	What it is used for
Your check-ins	One tap on Home or the sleep page, at most one a day	Your scale, the test, your timing
Your sleep times	The main sleeps and naps of the last 3 weeks, as the sleep section found them	The sleep process, every 15 minutes of the day
Nights you woke up on your own	The free nights of the last 90 days, with your answers to “How did you wake up?”	Where your daily rhythm starts
Last night’s need	The sleep section (sleep need, catch-up, naps)	Whether last night was short
Tonight’s bedtime	Tonight’s sleep onset once recorded, else the suggested bedtime	The end of the waking day

A night without a recording is filled in from your usual times (the median of your last 14 usable nights, or 23:00 to 07:00 before there are any). When that night is last night, the curve says so.

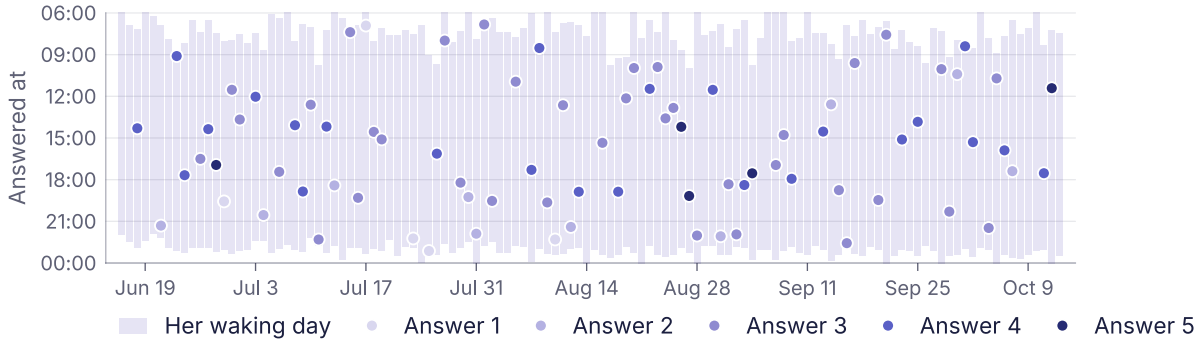
Your answers stay on this iPhone. Check-ins are kept in a file excluded from backups, never in iCloud and never synced, for 400 days. Unlike everything else Evinlo knows about you, they cannot be worked out again from Apple Health, so “Move to a new iPhone” carries them over in a file.

How it works, step by step

Step 1: the check-in

The question appears at most once a day, on Home and on the sleep page, from your final wake-up until 30 minutes before tonight’s bedtime. An answer can be undone right away, and “Not now” skips the day.

Answers at different times of day teach the model far more than the same answer every morning. When your latest 14 answers all fall in the same third of the waking day (at least 5 of them), the app says so. An optional reminder (offered with its switch on in the last step of the onboarding, and off until you turn it on everywhere else) asks at a different time each day: every block of three days gives each third of the waking day one reminder, never earlier than 60 minutes after you woke and never later than 90 minutes before bedtime.

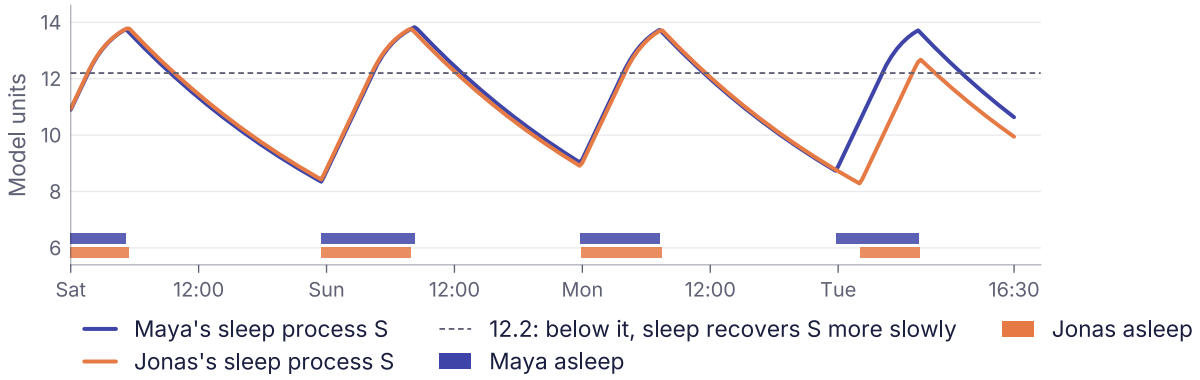


Maya's 81 answers over her 120 days. Each light bar is a waking day, from her final wake-up to her bedtime; each dot is one answer, darker for more awake. The synthetic answers fall at random times of the day, on most days.

Step 2: the sleep process

Evinlo uses the **three-process model of alertness**, published by Åkerstedt and Folkard and refined since (see the references). Its first process, **S**, is the sleep process: how rested the model thinks you are. It is high after a good night, falls through the day and recovers while you sleep.

- **Awake**, S falls toward 2.4: S = 2.4 + (S at waking − 2.4) × e^(−0.0353 × hours awake).
- **Asleep**, S climbs toward 14.3. From 12.2 it climbs quickly at first, then more slowly. Below 12.2 it climbs in a straight line, 0.80 an hour: this **brake** makes the model answer to repeated short nights.



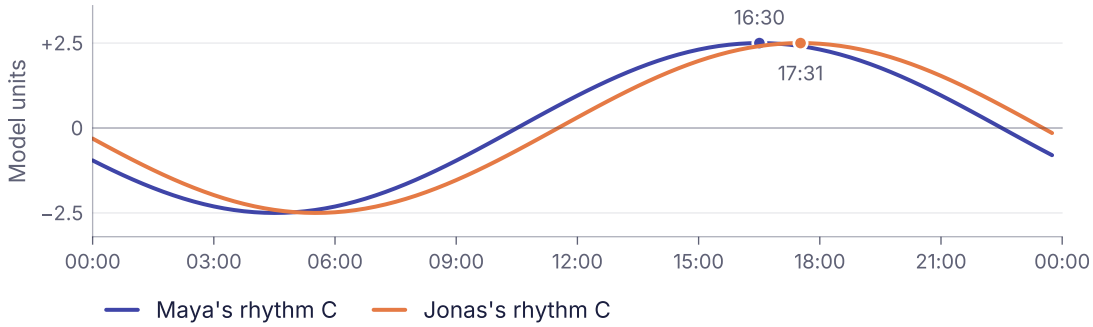
The sleep process S over four days, worked out every 15 minutes from the sleep the app recorded. Maya and Jonas run alike until Monday night: Jonas fell asleep after two, so on Tuesday he woke with less recovered.

On Tuesday morning Maya woke with S at about 13.7; Jonas, after his late night, at about 12.7. Both then lose about 3.5% of the distance to 2.4 every hour.

Step 3: the daily rhythm

The second process, **C**, is the daily rhythm: C = 2.5 × cos(2π × (clock − p) ÷ 24), a wave that peaks once a day at the hour p. By default p is 16.8 hours (16:48).

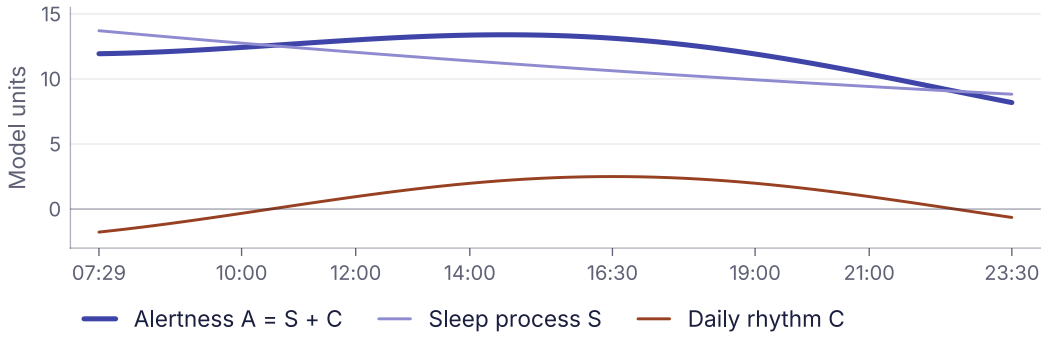
Evinlo moves p by your sleep times, the way chronotype questionnaires (which place you between a morning and an evening person) do: it takes the middle of your sleep on free nights of the last 90 days (at least 10 of them), corrects it for oversleeping, and shifts p by half its distance from 04:00, at most 1.5 hours either way. This mapping is an assumption, not a published rule, so Evinlo names no time of day from it.



The daily rhythm C on Tuesday. Maya's peaks at 16:30, from the middle of her free nights. Jonas's peaks at 17:31: his answers have taught Evinlo his own timing (Step 6).

Step 4: alertness, A = S + C

The model's alertness is the sum: **A = S + C**. Two parts that some versions of the model have are left out on purpose: the grogginess right after waking (in the validation Evinlo follows, its published values made predictions worse) and a 12-hour wave for an afternoon dip. So Evinlo does not predict an afternoon dip, and its curve starts only an hour after you woke.



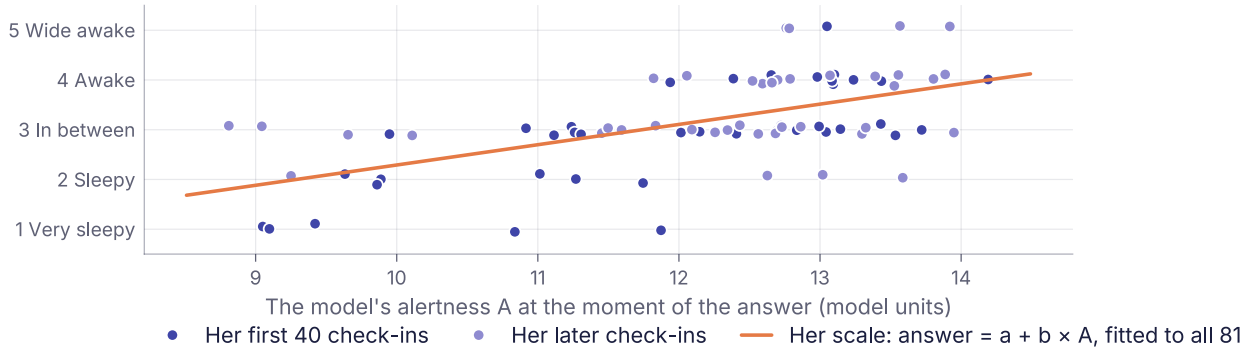
Maya's Tuesday in the model, from her wake-up at 07:29 to her bedtime at 23:30. S falls all day, the rhythm rises into the afternoon, and their sum is highest around 14:30.

Step 5: your scale

The model speaks in its own units. Your answers speak in 1 to 5. Evinlo connects them with a straight line, fitted to your check-ins:

> answer = a + b × A

Each check-in gives one point: the model's A at the moment you answered, and your answer. The line is fitted to the latest of them (at most 120, from the last 180 days).



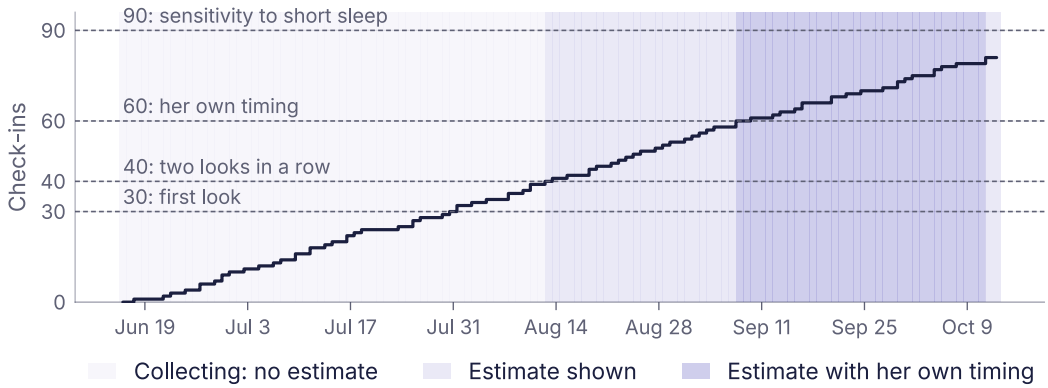
Maya's 81 check-ins against the model's A at the moment of each. The line is her scale: a = −1.79, b = 0.41. Her first 40 check-ins are the ones that passed the test.

Step 6: the test

A line can always be drawn. The question is whether your answers really **follow** the model. So the estimate is shown only when a test passes:

- at least 30 check-ins over at least 14 days, with at least 5 each in two of the three parts of the waking day, and at least two different answers;
- a slope b above zero, with a **t-value of 3.5 or more**, and a correlation of at least 0.2. The t-value says how many times larger the slope is than its uncertainty; the correlation how closely your answers follow the line, 1 being a perfect fit. The uncertainty is a Newey–West standard error, which allows for days that run together, such as a week of poor sleep;
- the test runs at 30 check-ins and then after every 10 more, never at every refresh, and it counts only once **two tests in a row** pass. A pattern that shows up by chance rarely still holds 10 check-ins later. So the estimate can appear from 40 check-ins on.

Once passed, the estimate stays while the latest 90 check-ins keep a positive slope with a t-value of at least 1.0.



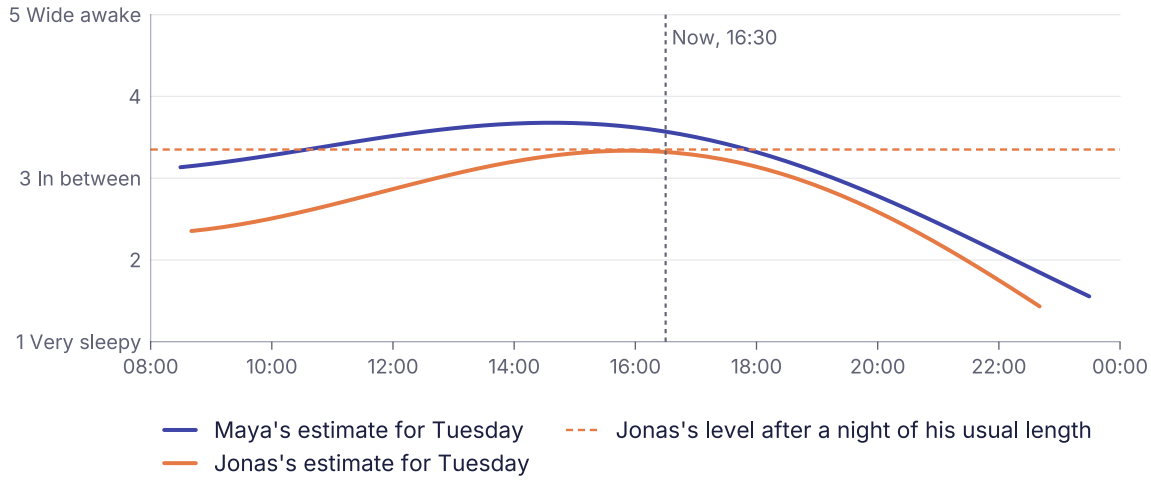
How the estimate came for Maya. The line counts her check-ins; the background shows what the app showed each day at 16:30.

Two more steps follow. **Your timing** (from 60 check-ins): the peak of the rhythm is fitted to your answers and blended with the starting point, which weighs as much as 60 check-ins. It is used only when it explains your answers at least 3% better; only then does the app name the time you may feel most awake. **Your sensitivity to short sleep** (from 90 check-ins, including 10 after nights at least an hour short of their need): how strongly the sleep process moves your answers, between half and twice the default.

Step 7: today's estimate

Once the test has passed, the sleep page shows today's curve on your own 1-to-5 scale, every 15 minutes from an hour after waking to tonight's bedtime, with the current time and today's answer marked. Two sentences may come with it:

- **The time you may feel most awake**, to the half hour, once your timing is your own and only while that time is still ahead.
- **"After a short night, you may feel less awake than usual today."** It compares like with like: the curve after last night against the same day after a night of your usual length ending at the same wake-up. It appears after a night at least an hour short of its need, when the day's level (from an hour after waking to an hour before bed) is at least 0.3 lower on your scale.



Tuesday's estimates on each person's own scale. Maya's curve after an ordinary night; Jonas's after his late night, below the level a night of his usual length would have given him.

A worked example

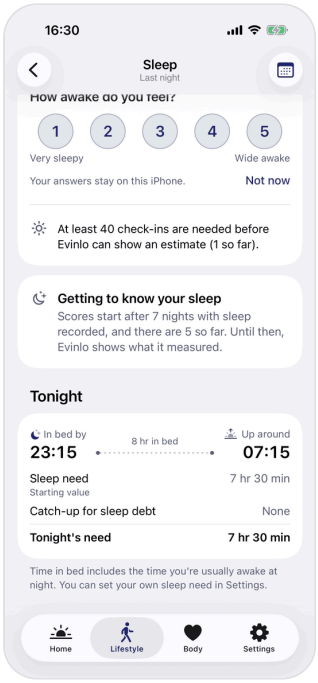
Maya answered 81 times in her 120 days; "3, In between" was her most common answer, 37 times. At her 40th check-in, on 13 August, Evinlo tested her first 30 check-ins and all 40 of them. Both tests passed (the second with t = 11.0 and a correlation of 0.74), and the estimate appeared. At 60 check-ins, on 8 September, her own timing explained her answers clearly better than the starting point, and it was used until 11 October. With 81 check-ins it fell short of the 3% it needs, so on Tuesday the app uses the starting point again (her rhythm peaks at 16:30) and names no time of day. That back and forth is the rule working: the timing is fitted afresh on every refresh.

On Tuesday she woke at 07:29 after a night of her usual length, and her bedtime is 23:30. Her scale is a = −1.79 and b = 0.41, and the latest 90 check-ins keep a t-value of 6.4. Her curve starts at 08:29. There is no sentence about a short night.

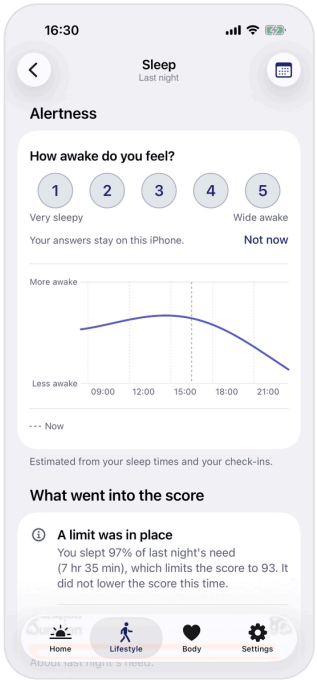
Jonas answered 81 times too. His scale is a little steeper than Maya's (b = 0.57), and his timing is his own: his rhythm peaks at 17:31, an hour after the starting point of 16:32 his sleep times gave. On Tuesday he woke at 07:40 after a night 2 hr 19 min short of its need. His curve starts at 08:40 and peaked at 15:55; at 16:30 that time has passed, so the app does not name it. Over the day his curve sits at 2.86 on average, about half a point below 3.35, the level a night of his usual length would have given, so the sleep page adds: "After a short night, you may feel less awake than usual today."

What you see in the app

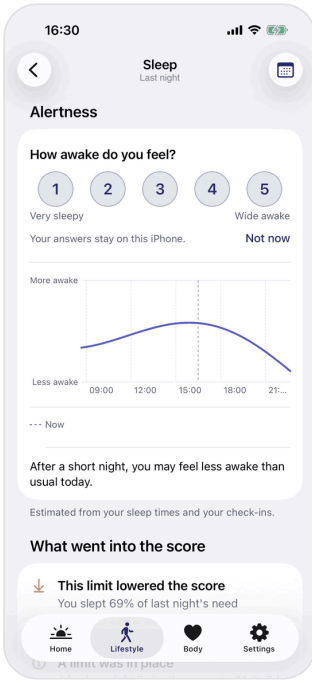
These screens were taken at 16:30 on Tuesday 13 October in the app's synthetic data mode.



A new person after their first night: "At least 40 check-ins are needed before Evinlo can show an estimate (1 so far)."



Maya's sleep page: today's check-in is still open, and her estimated curve, with "Estimated from your sleep times and your check-ins".



Jonas after his late night: "After a short night, you may feel less awake than usual today."

Honest limits

- **An estimate, not a measurement.** No device measures alertness. The curve comes from your sleep times and a model, scaled by your answers; it says "may" and "likely" and gives no advice about driving, work or health.
- **The synthetic answers fit by construction.** Maya's and Jonas's check-ins were generated from the model, so their passing tests show that the test finds the model when it is right. Whether real answers follow it can only be seen on real check-ins, which is why nothing is shown before they do.
- **Only sleep times go in.** Caffeine, light, illness, stress and a hard workout move how awake you feel; the model sees none of them.
- **No afternoon dip, no grogginess.** Evinlo leaves both parts out (Step 4), so it will not predict a post-lunch dip, and it shows nothing for the first hour after waking.
- **The starting point is an assumption.** Shifting the rhythm by your mid-sleep follows chronotype research in direction, but no published rule maps one onto the other.
- **Chance passes are rare, not impossible.** In simulations with answers that did not depend on the model at all, the test passed within a year in at most 3% of runs.

For the experts

Version and order of operations

This explainer describes `AlertnessModel.version` **1.3.0**; the curve carries it. The sleep section works alertness out for today only, after the night, the need and tonight's plan (`SleepEngine`, `AlertnessAnalysis`):

1. The waking windows of the last 180 days (`WakingDays`): from each night's final wake-up (else the usual one) to the next night's onset once final, else the suggested bedtime, else the usual onset; at least 3 hours. Each check-in belongs to the window it falls in; those of the last 180 days are used, at most the latest 120.
2. The starting phase from the free nights of 90 days (`StartingPhase`); for each used check-in, S from the sleep of the 21 days before it (`ThreeProcessModel`) and C on the clock of the zone it was given in.
3. The calibration (`AlertnessCalibrator`): the scale, the test at its looks, the personal phase, the sensitivity, the keep test. Once calibrated, the curve, the peak and "less awake than usual".

The model

With t the hours since waking or falling asleep, and the values restated by Ingre and colleagues (2014, equations 1.1–1.8):

$$S_{\text{awake}}(t) = 2.4 + (S_0 - 2.4)\, e^{-0.0353\, t}$$

$$S_{\text{asleep}}(t) = 14.3 - (14.3 - S_0)\, e^{g t}, \qquad g = \frac{\ln\big((14.3 - 14.0)/(14.3 - 7.96)\big)}{8} \approx -0.381$$

from $S_0 \geq 12.2$; below it S rises linearly by g × (12.2 − 14.3) ≈ 0.80 an hour until 12.2, then as above. The rhythm and the alertness, with the clock in hours and k_S the sensitivity to sleep loss (1 until it is fitted):

$$C = 2.5 \cos\!\left(\frac{2\pi\,(\text{clock} - p)}{24}\right), \qquad A = k_S\, S + C$$

The starting phase, with MSFsc the median mid-sleep of the free nights corrected for oversleeping (minus half the difference between the median sleep of free nights and of all nights, when free nights are longer):

$$p_1 = 16.8 + \text{clamp}\big(0.5\,(\text{MSFsc} - 4.0),\ -1.5,\ 1.5\big)$$

Your scale is the least-squares line answer = a + b × A. Maya's at 16:30 on Tuesday: a = −1.79, b = 0.41, so a model alertness of 13 reads as 3.5 on her scale.

The calibration in detail

- **The test** fits answer = a + b × A with the starting phase and $k_S = 1$ to the check-ins of each look, look m holding the check-ins among the first 10 × m stored ones. The Newey–West lag is $\lfloor 4\,(n \div 100)^{2/9} \rfloor$, at least 1. While the used check-ins are the whole history, the first look that ends a run of two passing looks counts; once older check-ins have rolled out, the latest look and the one before it must both pass.

- **Your timing** fits answer = $\beta_0 + \beta_1 S + \beta_2 \cos(2\pi \text{ clock} \div 24) + \beta_3 \sin(2\pi \text{ clock} \div 24)$. With $\beta_1 > 0$ and a residual sum of squares at least 3% below that of the starting fit, the fitted peak $\text{atan2}(\beta_3, \beta_2)$ is blended with the starting phase in a weighted circular mean (weights n and 60).
- **Sensitivity** fits answer = $\beta_0 + \beta_S S + \beta_C C$; $k_S = \beta_S \div \beta_C$, kept between 0.5 and 2.
- **States:** collecting (fewer than 40 check-ins), not tracking (with the next test), needs other times of day (the latest look covers fewer than two thirds), calibrated.

Edge cases

- **The first hour after waking** has no curve, because the model has no term for grogginess.
- **Last night not recorded:** the curve assumes your usual sleep times and says so.
- **Rolling windows:** when check-ins leave the 180 days or the latest 120, the tests keep their rhythm of one look every 10 check-ins, so the estimate never flickers at every refresh. Deleting check-ins older than 400 days can bring a test one check-in early.
- **Your timing can come and go.** It is fitted again on every refresh and used only while it is 3% better; Maya's came at 60 check-ins and left at 81.

Parameters

Every parameter is in `AlertnessParameters.v1` (`Packages/EvinloKit/Sources/EvinloAlgorithms/Contract/Parameters.swift`). Minutes unless stated.

Rule	Parameters
S: lower and upper asymptote, decay awake and recovery asleep per hour, the brake	`lowerAsymptote` 2.4, `upperAsymptote` 14.3, `wakeDecay` −0.0353, `sleepRecovery` ln(0.3 ÷ 6.34) ÷ 8, `brakeLevel` 12.2
C: mesor, amplitude, default peak (hours)	`circadianMesor` 0, `circadianAmplitude` 2.5, `defaultPhase` 16.8
Sleep inertia (not used: weight 0)	`inertiaAmplitude` −5.72, `inertiaDecay` −1.51, `inertiaWeight` 0
Starting phase: 90 days, at least 10 free nights; reference mid-sleep 4.0 h; 0.5 h of phase per hour; at most ±1.5 h; the fitted phase at most ±3 h from the default	`msfWindowDays` 90, `msfMinimumFreeNights` 10, `chronotypeReferenceMidSleep` 4.0, `chronotypeSlope` 0.5, `maxStartShift` 1.5, `maxPhaseShift` 3
Sleep periods: 21 days; missing nights from the last 14 usable nights, else 23:00 to 07:00	`periodHistoryDays` 21, `imputationNights` 14, `defaultOnsetHour` 23, `defaultWakeHour` 7
Warm-up before the first period: 7 days of 8 hours asleep and 16 awake, from S = 12	`warmUpDays` 7, `initialS` 12
Check-ins in use: 180 days, at most the latest 120; a scale from 10	`checkInWindowDays` 180, `rollingCheckIns` 120, `fitMinimumCheckIns` 10
The test: from 30 check-ins, every 10, over 14 days, 5 in each of 2 thirds, 2 different answers, t ≥ 3.5, r ≥ 0.2, 2 looks in a row	`gateCheckIns` 30, `gateEvery` 10, `gateDays` 14, `gateBinMinimum` 5, `gateBinsNeeded` 2, `gateDistinctRatings` 2, `gateT` 3.5, `gateR` 0.2, `gateConsecutivePasses` 2
Kept while the latest 90 have t ≥ 1.0	`keepWindow` 90, `keepT` 1.0
Your timing: from 60 check-ins, the starting phase weighs 60, at least 3% better	`phaseCheckIns` 60, `phasePriorWeight` 60, `phaseMinimumImprovement` 0.03

Rule	Parameters
Sensitivity: from 90 check-ins with 10 after nights 60 min short; between 0.5 and 2	`sleepLossCheckIns` 90, `sleepLossShortNights` 10, `sleepLossShortBy` 60, `sleepLossLowerBound` 0.5, `sleepLossUpperBound` 2
The curve: every 15 min from 60 min after waking; the peak at least 60 min before bed	`curveStep` 15, `curveStartAfterWake` 60, `peakMargin` 60
“Less awake than usual”: 0.3 lower, after a night 60 min short	`lessAwakeDrop` 0.3, `lessAwakeShortBy` 60
The check-in: a waking day of at least 180 min; closes 30 min before bed; the concentration hint over the latest 14, from 5; usual times from 14 nights	`minimumWindow` 180, `promptCloseBeforeBed` 30, `concentrationRecent` 14, `concentrationMinimum` 5, `usualTimesNights` 14

Where it lives in the code

All under `Packages/EvinloKit/Sources/EvinloAlgorithms/`, in `Alertness/`: the model (`ThreeProcessModel.swift`), the starting phase (`StartingPhase.swift`), waking days and sleep periods (`WakingDays.swift`), the calibration (`AlertnessCalibration.swift`), the day and the curve (`AlertnessAnalysis.swift`), the prompt and the thirds of the day (`CheckInBins.swift`); the parameters and the version in `Contract/Parameters.swift`.

What the validation checks

Tests in `Packages/EvinloKit/Tests/EvinloAlgorithmsTests` (`ThreeProcessModelTests`, `AlertnessCalibrationTests`, `AlertnessTodayTests`, `CheckInScheduleTests`) pin the model’s golden values and every rule of the test. `evinlo-eval simulate-checkins` runs synthetic answers through the same calibration on synthetic sleepers’ and one developer’s own sleep times. On the synthetic sleepers, answers that follow the model by half a point per model unit passed within a year for every simulated user, after a median of about 40 check-ins when they came at random times; answers that did not depend on the model passed in at most 3% of runs (the guardrail is 5%).

References

- Åkerstedt T., Folkard S., “The three-process model of alertness and its extension to performance, sleep latency, and sleep length”, *Chronobiology International* 14(2), 1997.
- Åkerstedt T. et al., “Accounting for partial sleep deprivation and cumulative sleepiness in the Three-Process Model of alertness regulation”, *Chronobiology International* 25(2), 2008.
- Ingre M. et al., “Validating and extending the three process model of alertness in airline operations”, *PLoS ONE* 9(10): e108679, 2014.
- Roenneberg T., Wirz-Justice A., Merrow M., “Life between clocks”, *Journal of Biological Rhythms* 18(1), 2003; Roenneberg T. et al., “Epidemiology of the human circadian clock”, *Sleep Medicine Reviews* 11(6), 2007.
- Jewett M. E. et al., “Time course of sleep inertia dissipation in human performance and alertness”, *Journal of Sleep Research* 8(1), 1999. Every rule with the simulations behind it is in docs/scoring.md, “Alertness and the daily check-in”.

Spotted something unclear or wrong, or want to improve a rule? Evinlo is open source, and every rule can be changed through a proposal and a replay. Start at **evinlo.com/contribute**.