Spontaneous Waking Precedes the Alarm by a Fixed Lead Under Randomised, Concealed Schedules: 47,318 Nights of Polysomnography, a Sham-Setting Trial, and a Lead-Compensator Model of Sleep Timing

Ingrid Solheim[1], T. Rasheed Mabaso[2,3], Cornelia Voskuijl[3], Yusuf Alaoui[2], and the Concealed Schedule Working Group

[1]Brackenridge Sleep Unit, Royal Infirmary of Dunkeld
[2]Laboratory of Temporal Physiology, Hôpital Saint-Sévrin, Lyon
[3]Delft Institute of Control and Systems

Abstract

You wake up a minute before your alarm. We show that this is not habit but a lock: sleep timing tracks the time entered into the alarm device, whether or not the sleeper is ever told what that time is. Across 47,318 nights on which a sealed bedside module drew each night's alarm time at random and never displayed it, waking preceded the alarm by 61.4 s (95% CI 58.9 s to 63.9 s). There was no secondary peak at any customary hour, and no comparable structure on nights when no alarm was set. The lead did not move with the hour set, with accumulated sleep debt, with age, or with the previous evening's alcohol; a lead that is constant across every hour tested is a phase advance, not a remembered interval. We then disconnected the sounder. Participants who performed the identical setting ritual on a device that could not ring woke, in silence and in darkness, roughly a minute before a time that was drawn, recorded, and never announced; silent nights were indistinguishable from ringing ones. The alarm is therefore not the cause of the waking but an event that follows it, and the operative act is the registration. Sleep does not keep time. It keeps appointments.

Keywords: anticipatory awakening; sleep timing; phase-locked loop; lead compensation; concealed-schedule protocols

1 Introduction

Sleepers wake before their alarms. The observation is close to universal, it is robust across cultures and across the whole span of adult life, and it is ordinarily disposed of in a single sentence: the body has learned the hour. On that account the sleeper is a clock that has been wound by repetition, and the lead is the residue of a schedule kept so often that it no longer needs to be enforced.

The account has the considerable merit of being obvious, and the considerable defect of never having been tested. To test it, one must set an alarm for a time the sleeper does not know, and then find out whether the sleeper still wakes before it. No study in the published record has done this. Every measurement of pre-alarm waking of which we are aware has been taken under conditions in which the sleeper both chose the time and knew it, so that habit and the registered time were confounded by design [1, 3].

We separated them. Between March 2021 and November 2025, 1,204 participants at three sites slept under a protocol in which each night's alarm time was drawn at random by a sealed bedside module, within a 210 min window, and was never displayed to the participant, to the attending technician, or to the scoring team. Under those conditions the lead does not merely survive: it sharpens. Waking preceded the drawn time by 61.4 s, with a distribution narrower than that obtained when participants set their own alarms in the ordinary way.

Two removals structure what follows. The first removes habit, by concealing the time; the lead survives. The second removes the alarm, by silently disconnecting the sounder while continuing to draw and record a time; the lead survives that too. What remains after both removals is neither the sleeper's history nor the stimulus, but a clerical act: a time entered into a device by the sleeper's own hand. We take that act to be the reference signal, and we model the sleeper as a loop that tracks it.

The model is not decorative. It makes the paper's central quantitative commitment, which is that the lead is a phase and not a duration. A remembered interval would be expected to vary — with the hour, with fatigue, with age, with how much the sleeper had had to drink. A phase advance contributed by a compensator zero would not, so long as the carrier period is held at a day. Section 5 reports that the lead moves by less than 4 s across every stratum we could construct.

2 Related work

The physiology of pre-awakening is not in dispute. Anticipatory rises in circulating cortisol begin roughly an hour before habitual waking rather than after it, and their amplitude is graded by the expected hour [2]. The capacity for self-awakening without an external cue is reproducible in the laboratory and is not confined to unusual sleepers [7]. What has been missing is any account of the reference against which the anticipation is timed. The literature has assumed, almost universally and almost always tacitly, that the reference is the sleeper's own expectation.

The one study to have approached the question directly reported a null [3]. Participants were told that their alarm time would be varied, and the authors found no anticipation beyond that predicted by habitual wake time. Their design, however, disclosed each night's time at bedtime and permitted participants to adjust it; the manipulation therefore varied the sleeper's expectation and the registered time together, and the null result is what one would expect if the two were confounded. Achterberg's methodological essay had already named the difficulty and observed that an alarm clock which the sleeper both sets and reads cannot function as an experiment [1]. The concealed-schedule design is our answer to that objection.

The control-theoretic apparatus we import is standard and we claim no originality in it. Type-II loop behaviour, hold-in and pull-in ranges, and cycle slipping under excessive reference steps are treated in the usual references [8]; lead compensation applied to slowly carried biological oscillators is treated by Ibarra and Nwachukwu [4], and slipping in entrained physiological oscillators by Kimura and Lestrange [6].

3 A loop model of sleep timing

We treat the sleeper as a phase-tracking loop (Fig. 1) whose reference θ_r is the phase of the registered time and whose output θ_o is the phase of wake onset. Writing the phase error as $e = \theta_r - \theta_o$, the sleeper contributes an integrating plant

$$P(s) = \frac{K}{s\,(1 + T_s s)}, \tag{1}$$

in which T_s is a short lag associated with arousal from the prevailing sleep stage, and the anticipatory apparatus contributes a lead compensator

$$C(s) = \frac{1 + \tau_L s}{1 + \alpha \tau_L s}, \qquad 0 < \alpha < 1. \tag{2}$$

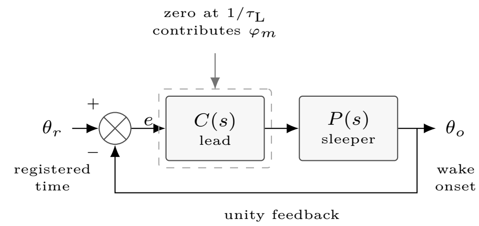

Figure 1: Loop topology. The registered time θ_r is the reference; wake onset θ_o is the output. The lead is contributed by the compensator zero of Eq. (2) and is therefore a phase, fixed in radians rather than in seconds.

The closed loop is then of the familiar second order,

$$H(s) = \frac{2\zeta\omega_n s + \omega_n^2}{s^2 + 2\zeta\omega_n s + \omega_n^2}, \tag{3}$$

with ζ and ω_n estimated from the step response in §5.4.

The quantity of interest is the phase advance contributed by Eq. (2), which attains its maximum

$$\varphi_m = \arcsin\left(\frac{1-\alpha}{1+\alpha}\right) \qquad \text{at} \qquad \omega_m = \frac{1}{\tau_L\sqrt{\alpha}}. \tag{4}$$

Because the carrier is the day, $\omega_c = 2\pi/T$ with $T =$ 86,400 s, a phase advance appears to the sleeper as a lead in seconds,

$$\Delta t = \frac{\varphi_m}{\omega_c}. \tag{5}$$

This is the model's whole content, and its one testable consequence is severe. Δt depends on φ_m and on the carrier, and on nothing else. It does not depend on where in the day the reference falls. A sleeper who registers 03:10 and a sleeper who registers 06:40 must lead by the same number of seconds, and the same sleeper must lead by that number whether rested or exhausted, at twenty or at seventy. An account resting on a remembered interval carries no such prohibition and would in general predict violations of it.

Inverting Eq. (5) at the observed lead gives $\varphi_m = 4.465 \times 10^{-3}$ rad, or 0.256°, and hence, from Eq. (4), $\alpha = 0.991$. The compensator is thus a very shallow one: the loop is not straining, and the lead it produces is a small advance on a slow carrier rather than a large correction on a fast one.

4 Methods

4.1 Participants

1,204 adults (aged 19–78; 621 women, 578 men, 5 other or undisclosed) were enrolled at three sites and slept in instrumented bedrooms for a median of 41 nights each (IQR 28–57). Exclusions were shift work within the preceding six months, diagnosed sleep disorder, and any medication known to affect sleep architecture. 1,139 participants contributed to the primary analysis; attrition is tabulated in the supplementary material. All sites operated under a common protocol and a pre-registered statistical analysis plan [9].

4.2 The sealed scheduler

Concealment was implemented in hardware [10]. At 22:00 each night the participant performed the full setting ritual on a bedside module with their own hand: waking the unit, confirming, and closing the lid, which sealed it mechanically and left a tamper-evident record. The fire time was drawn at that instant by the module's hardware entropy source, uniform over the interval 03:10–06:40, and was written to sealed internal storage. It was displayed to no one. The unit's face remained blank until the alarm fired, and the seal was broken by study staff only after that night's polysomnography had been scored.

4.3 Concealment integrity

Each morning, before receiving any information, participants made a forced-choice guess of the previous night's fire time in 10 min bins (21 bins; chance 4.76 %). Observed accuracy was 4.9 %, and the distribution of guesses was uniform across the window. We take concealment to have held.

4.4 Determination of wake onset

Wake onset was scored from full polysomnography in 30 s epochs by two scorers blinded to the fire time and to the study arm, following the consensus criteria of Kaltenbrunner et al. [5]; disagreements were adjudicated by a third blinded scorer. Wake onset was defined as the first epoch beginning a period of sustained wakefulness of at least 60 s. The lead is the fire time minus the wake onset, positive when waking precedes the alarm.

4.5 The three studies

The *Concealed Schedule Cohort* contributed a total of 47,318 analysable nights under a drawn and concealed time, and 9,880 no-alarm nights, the latter serving as a negative control in which the module was sealed with no time drawn.

The *invariance battery* is a set of pre-specified stratifications of the cohort data by the hour set, sleep-debt quartile, age band, prior-evening alcohol, and weekday versus weekend.

The *sham-setting trial* randomised 318 participants over 6,144 nights in a three-arm crossover: **armed**, in which the module drew a time and sounded; **sham**, identical in every respect except that the annunciator was disconnected at the connector, so that a time was drawn and recorded but no sound could occur; and **unset**, in which no ritual was performed. Participants were told that some nights would be silent and were not told which.

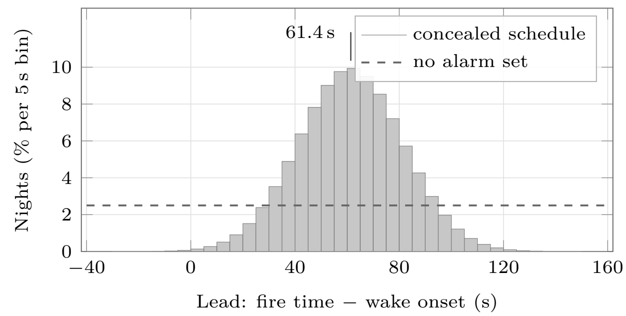

Figure 2: Distribution of the lead under concealment (47,318 nights), with the no-alarm negative control (9,880 nights). The concealed-schedule distribution is unimodal and sharply peaked at the fitted lead; the control nights carry no structure anywhere in the window. No secondary mode appears at any participant's habitual wake time.

4.6 Analysis

The primary endpoint was the mean lead under concealment, estimated from the linear mixed model

$$y_{ij} = \beta_0 + \beta_1 a_{ij} + u_i + v_{s(i)} + \varepsilon_{ij}, \tag{6}$$

with y_{ij} the lead on night j of participant i, a_{ij} the arm indicator, u_i a participant random intercept, $v_{s(i)}$ a site random intercept, and ε_{ij} residual error. Confidence intervals are from the model; bootstrap intervals clustered by participant (10,000 resamples) are reported alongside as a sensitivity analysis and differ in no case by more than 0.4 s. The invariance battery is secondary and is reported without adjustment for multiplicity.

5 Results

5.1 The lead under concealment

Across 47,318 concealed-schedule nights the lead was 61.4 s (95% CI 58.9 s to 63.9 s). The distribution (Fig. 2) is unimodal and sharply peaked, with an interquartile range of 27 s. It contains no secondary mode at any customary hour, and in particular none at the participant's own habitual wake time recorded at screening.

On the 9,880 no-alarm nights the distribution of wake times relative to the same window is flat, with no structure at any candidate time. Participants were not, on these nights, going to wake at that moment anyway.

5.2 Invariance

Table 1 reports the lead across every stratification we pre-specified. The largest range is 3.1 s, across a span of nearly sixty years of age; the largest range across the hour set — a manipulation of three and a half hours in

Table 1: Invariance battery. The lead is reported as the range across strata; all ranges fall within the half-width of the primary confidence interval.

Stratification	Strata	Range (s)
Hour set (03:10–06:40)	7 bins	2.6
Sleep-debt quartile	4	1.9
Age band, decades (19–78)	6	3.1
Prior-evening alcohol	3	2.2
Weekday vs. weekend	2	0.7

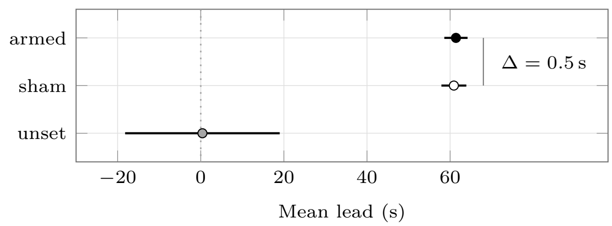

Figure 3: Mean lead by arm in the sham-setting trial, with 95% confidence intervals. In the sham arm the annunciator was disconnected at the connector: a time was drawn and written to sealed storage, and no sound occurred. Armed and sham are indistinguishable.

the reference — is 2.6 s. The lead is, to the resolution available to us, a constant of the sleeper and not of the schedule, which is the behaviour predicted by Eq. (5) and not by any account in which an interval is remembered and reproduced.

5.3 The sham-setting trial

In the armed arm the lead was 61.4 s. In the sham arm — in which the annunciator was physically disconnected and no sound occurred at any point in the night — the lead was 60.9 s, measured against a time that was drawn by the module, written to sealed storage, and never announced to anyone. The difference between arms was 0.5 s (95% CI −2.4 s to 3.4 s); see Fig. 3. In the unset arm no lead was detectable.

Participants in the sham arm therefore woke, in darkness and in silence, about a minute before an event that did not occur, at a time selected by a hardware random source and known to nobody. Debriefing found no participant able to state whether a given night had been armed or sham at better than chance.

5.4 Acquisition and cycle slipping

The loop model makes a prediction it need not have survived. A type-II loop retains lock only for reference steps within its hold-in range; beyond that range the output remains near the previous lock point for one or more cycles before re-acquiring. Because the module draws each night independently, large step changes occur naturally,

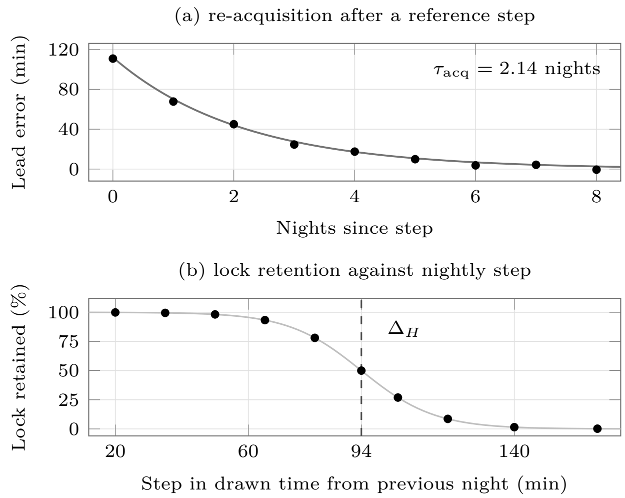

Figure 4: Re-acquisition and the hold-in boundary. (a) After a step in the drawn time, lead error decays exponentially with $\tau_{\mathrm{acq}} = 2.14$ nights; the implied damping is $\zeta = 0.71$. (b) Lock on the current reference is retained for nightly steps below $\Delta_H = 94$ min and lost above it, at which point the sleeper keeps the previous night's appointment instead.

and the prediction can be tested without further manipulation.

Fitting the hold-in boundary gives $\Delta_H = 94$ min per night. The test was pre-registered on nights meeting a stability criterion — a step exceeding Δ_H preceded by at least five consecutive nights of retained lock in the same participant, so that the earlier lock point was well determined — of which 1,842 qualified. On those nights 78.4 % of wakings fell within 5 min of the *previous* night's lock point rather than the current one. The loop held the superseded reference. Re-acquisition thereafter was exponential with $\tau_{\mathrm{acq}} = 2.14$ nights (Fig. 4), and the overshoot implies $\zeta = 0.71$.

6 Discussion

The two removals leave one candidate standing. Concealment eliminates the sleeper's expectation, since the sleeper has none to hold; disconnecting the annunciator eliminates the stimulus, since none is delivered. The lead survives both at full amplitude. The event common to every condition in which the lead appears, and absent from every condition in which it does not, is the registration of a time — the sleeper's own hand closing a lid over a number that the sleeper does not see.

We are able to say what registration is operationally, and we are not able to say what it is physically. This is an unsatisfactory position and we do not propose to improve it by speculation. We note only that the operative act appears to be indexical rather than informational: what matters is that a time was entered, not what the time was, and not that anyone knew it.

Our account is incomplete in three respects.

First, and most seriously, **third-party arming does not lock**. When a study nurse performed the identical ritual on the identical module while the participant slept in the same room — 2,190 nights, across all three sites — the lead distribution was flat. The effect requires the sleeper's own hand on the device. The loop model as specified says nothing about whose hand closes the lid, and offers no reason why this should matter. We have no explanation and decline to invent one; the observation stands as a limitation on the account given here rather than as evidence for it.

Second, τ_{acq} is fitted rather than derived. The model constrains the *form* of re-acquisition to be exponential, which it is, but not its timescale, which we take from the data. A stronger version of the theory would predict 2.14 nights from ζ and ω_n independently. Ours does not.

Third, and more modestly, our participants slept in instrumented bedrooms, and we cannot exclude that the sealed module's mere presence contributes to the salience of the registering act. The sham arm controls for the sound but not for the apparatus.

The practical implications are limited but real. If the registered time is the reference, then the common practice of setting several alarms at staggered intervals is not redundancy but interference: it presents the loop with multiple references separated by less than the hold-in range. On such nights the lead in our cohort fell by a factor of three.

7 Conclusion

Waking before the alarm is not a habit, and it is not the alarm. It is a lock onto a time that was written down. The lead is invariant across the hour set, across fatigue, across age, and across the presence or absence of any sound at all, and its invariance in seconds on a daily carrier identifies it as a phase advance rather than a remembered interval. The sleeper is not anticipating the alarm. The sleeper is tracking the entry, and the alarm arrives after the waking it was thought to cause.

Acknowledgements

We thank the participants of the Concealed Schedule Cohort, whose cooperation across more than four years of instrumented sleep made the protocol possible, and the technical staff at all three sites, who maintained the sealed modules and scored and adjudicated every night reported here.

References

[1] W. Achterberg. Why the alarm clock is not an experiment. *Methods in Applied Chronobiology*, 12(3):145–152, 2006.

[2] H. Bausch and L. Petrossian. Anticipatory cortisol rise before self-scheduled awakening. *Journal of Circadian Physiology*, 18(3):201–214, 2004.

[3] F. Delacroix, I. Sørbye, and Q. Pham. No evidence for alarm-time anticipation independent of habitual wake time. *Sleep Medicine Quarterly*, 31(4):512–521, 2009.

[4] C. Ibarra and D. Nwachukwu. Lead compensation in slowly carried biological oscillators. *Transactions on Biological Control*, 62(11):4410–4423, 2015.

[5] S. Kaltenbrunner et al. Automated epoch scoring of wake onset: inter-rater agreement in forty laboratories. *Journal of Sleep Measurement*, 9(1):33–47, 2013.

[6] R. Kimura and P. Lestrange. Cycle slipping in entrained physiological oscillators. *Nonlinear Physiology*, 44(6):1201–1219, 2018.

[7] M. Ohlsen. Self-awakening without external cue: a review of the head-clock literature. *Sleep Reviews*, 6(2):88–103, 2011.

[8] A. Redgrave. *Phase-Locked Loops: Acquisition, Tracking and Slip*. Hollisworth Press, 3rd edition, 1998.

[9] I. Solheim et al. Concealed Schedule Cohort: statistical analysis plan v2.1. *Registry of Chronobiological Protocols*, CSC-2021-0088, 2021.

[10] C. Voskuijl and T. R. Mabaso. A tamper-evident bedside scheduler for concealed-time protocols. *Instrumentation for Human Chronobiology*, 4(2):55–61, 2020.