Additional file 1: Supplementary methods and results for radial drawdown

Haobo Liu and Ke Li

S1 Scope of this supplement

This supplement provides the discrete algorithms, proofs, pre-specified benchmark settings, animal-data preprocessing and numerical contrasts that support the main Methodology article. The sections follow the dependency structure of the evidence rather than the chronology of the analyses. All values are taken from archived analysis tables; rendering the figures and compiling the manuscript do not rerun parameter selection.

S2 Discrete radial-drawdown algorithms

S2.1 Shared recursion

For a sampled track $\{(t_k, \mathbf{x}_k)\}_{k=0}^{n}$ and scale δ, the recursion begins with the origin and running maximum at $k = 0$. The scan updates radial distance from the current origin and stores the earliest index attaining the largest observed radius. Once the selected confirmation rule is met, it returns the current origin, the stored maximum as endpoint, the confirmation location and time as trigger, and the endpoint–trigger lag. The endpoint then becomes the next origin. If the deployment ends before confirmation, the final interval is marked as censored rather than emitted as a completed event.

In the primary RD implementation, confirmation requires $M_k - r_k \geq \delta$ at a sampled point. A crossing time may be interpolated along the final segment after this inequality is met, but interpolation neither creates the event nor inserts a new point into the recursion. The supplementary RD-L sensitivity instead treats the observed trajectory as the piecewise-linear path joining consecutive fixes. Confirmation occurs when the minimum distance from the current origin to the newest segment is at most $M_k - \delta$. Because each implementation recurses from its

own endpoints, one event table cannot be obtained by filtering the other.

S2.2 Discrete correctness, termination and cost

For a current origin index o, define

$$r_k^{(o)} = \|\mathbf{x}_k - \mathbf{x}_o\|, \qquad M_k^{(o)} = \max_{o \le i \le k} r_i^{(o)}, \qquad D_k^{(o)} = M_k^{(o)} - r_k^{(o)}. \tag{S1}$$

If the set $\{k > o : D_k^{(o)} \ge \delta\}$ is non-empty, let c be its smallest member and let e be the smallest index attaining $M_c^{(o)}$. The discrete RD event is the unique triple (o, e, c), after which recursion sets $o \leftarrow e$. If the set is empty, the final excursion is right-censored.

The literal scan is sound and complete under this definition. It cannot emit an event before c because every earlier drawdown is below δ. At c, the threshold is met and the stored earliest running maximum is exactly e. Moreover, $D_c^{(o)} \ge \delta > 0$ implies $r_c^{(o)} < M_c^{(o)}$, so $e < c$, while $M_c^{(o)} > 0 = r_o^{(o)}$ implies $o < e$. Every confirmed event therefore satisfies

$$o < e < c. \tag{S2}$$

The next origin index increases strictly, which proves termination on a finite sequence and bounds the number of confirmed events by $n - 1$.

The reference implementation stores the input, K event rows and a constant amount of scan state, requiring $O(n + K)$ storage. Restarting at an endpoint can revisit fixes between that endpoint and its trigger. The direct implementation therefore has the conservative time bound $O(nK) \le O(n^2)$; it constructs no pairwise distance matrix. We do not assume a linear worst-case bound.

The original fixes attached to consecutive endpoints are nested within the corresponding step at a fixed δ. Event tables across different values of δ, however, are a scale-indexed family rather than a nested hierarchy because recursion re-anchors the origin. For example, the one-dimensional path $(0, -1, 0, -2, 0)$ has endpoint index 1 at $\delta = 1$ and endpoint index 3 at $\delta = 1.5$. Event count is likewise not required to be monotone in δ. Each scale is therefore recomputed independently and reported as a response curve.

S2.3 Segment distance used by RD-L

Let $\mathbf{a}$ and $\mathbf{b}$ be adjacent position vectors relative to the current origin and let $\mathbf{v} = \mathbf{b} - \mathbf{a}$. The closest position on the segment is

$$\mathbf{z} = \mathbf{a} + q\mathbf{v}, \qquad q = \min\left(1, \max\left(0, -\frac{\mathbf{a}\cdot\mathbf{v}}{\mathbf{v}\cdot\mathbf{v}}\right)\right). \tag{S3}$$

The segment enters the threshold circle when $\|\mathbf{z}\| \le M_k - \delta$. Degenerate zero-length segments use $\|\mathbf{a}\|$. Ties in the running maximum resolve to the earliest fix so that plateaus have a deterministic endpoint.

Figure S1 shows how this alternative observation model can confirm a threshold crossing between two sampled fixes.

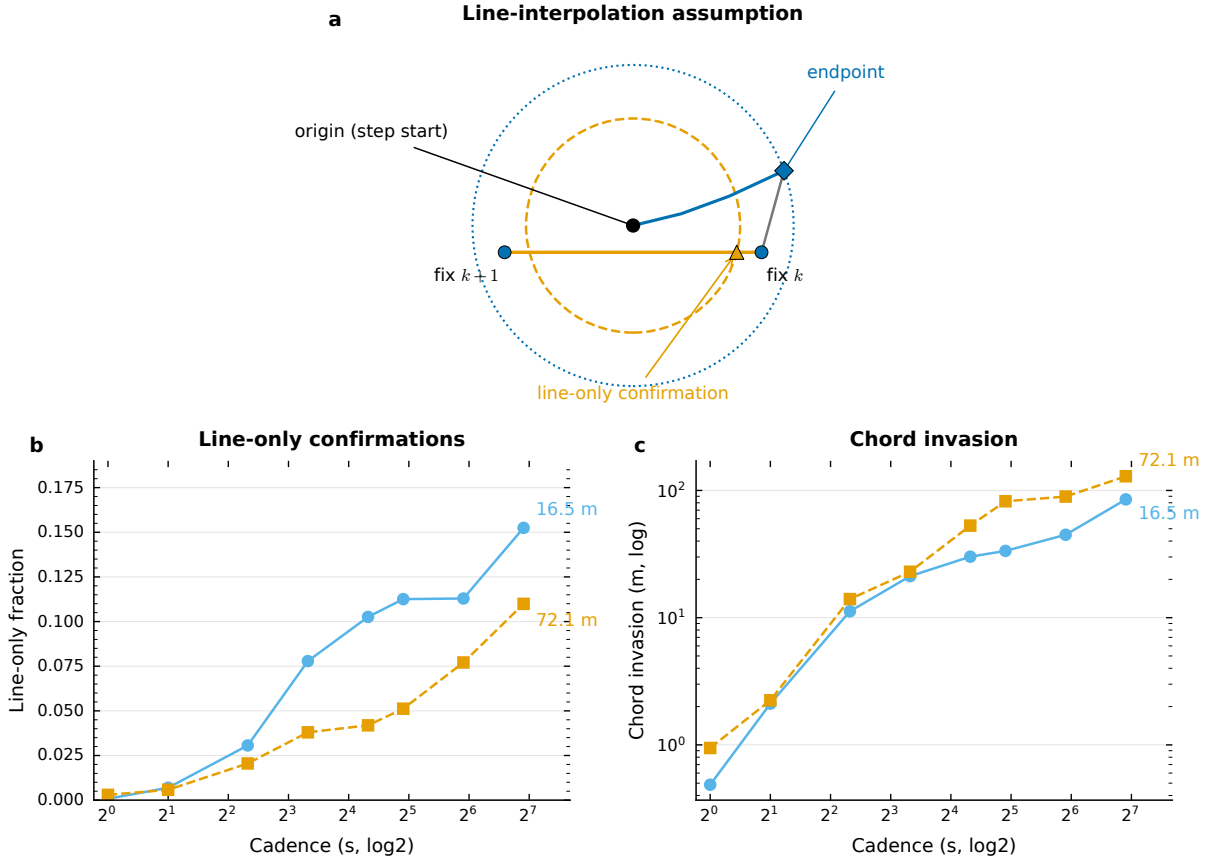

Figure S1: Piecewise-linear interpolation sensitivity. **a**, the path begins at the circle centre and reaches its endpoint on the running-maximum circle M. Two later sampled fixes remain outside $M - \delta$, but the chord joining them enters the confirmation circle. RD-L therefore confirms at the chord crossing, whereas sampled-fix RD does not confirm at either bounding fix. **b**, fraction of RD-L confirmations produced by a chord but not by a sampled point across cadences. **c**, median radial penetration of these line-only confirmations. Circles with solid lines denote 16.5 m; squares with dashed lines denote 72.1 m. Panels **b**–**c** use the Crozet wandering-albatross multisensor archive [1, 2].

S3 Observation-theory derivations

S3.1 Peak-loss distribution

For the local radial peak

$$r(s) = \begin{cases} M + ws, & s \le 0, \\ M - us, & s \ge 0, \end{cases} \tag{S4}$$

the nearest samples occur at $-A$ and $\Delta - A$, where $A \sim \text{Uniform}(0, \Delta)$. The observed maximum is the larger of $M - wA$ and $M - u(\Delta - A)$, so

$$Q_\Delta = M - M_{\text{sample}} = \min\{wA, u(\Delta - A)\}. \tag{S5}$$

For $0 \le q \le wu\Delta/(w+u)$,

$$\Pr(Q_\Delta > q) = \Pr\left(A > q/w,\ A < \Delta - q/u\right) \tag{S6}$$

$$= 1 - \frac{q}{\Delta}\left(\frac{1}{w} + \frac{1}{u}\right). \tag{S7}$$

Hence Q_Δ is uniform with

$$Q_\Delta \sim \text{Uniform}\left(0, \frac{wu\Delta}{w+u}\right), \tag{S8}$$

$$\mathbb{E}[Q_\Delta] = \frac{wu\Delta}{2(w+u)}, \tag{S9}$$

$$\text{Var}(Q_\Delta) = \frac{w^2u^2\Delta^2}{12(w+u)^2}. \tag{S10}$$

S3.2 Endpoint-time error

The pre-peak sample is selected when $wA \le u(\Delta - A)$, which occurs with probability $u/(w+u)$. Otherwise the post-peak sample is selected. Integration over these two phase ranges gives

$$\mathbb{E}[E_\Delta] = \frac{\Delta(w-u)}{2(w+u)}, \qquad \mathbb{E}[|E_\Delta|] = \frac{\Delta(w^2+u^2)}{2(w+u)^2}. \tag{S11}$$

For $w = u$, signed bias is zero but expected absolute error is $\Delta/4$.

S3.3 Exact trigger time and effective threshold

After the true peak, RD observes retreat $us - Q_\Delta$. Samples occur at $-A + j\Delta$, so the first sampled trigger is

$$T_{\Delta,\delta}(A) = -A + \Delta \left\lceil \frac{A + [\delta + Q_\Delta(A)]/u}{\Delta} \right\rceil. \tag{S12}$$

A discrete threshold δ therefore requires $us \geq \delta + Q_\Delta$, followed by the wait $B_\Delta \in [0, \Delta)$ to the next fix. The true retreat at confirmation is

$$\Theta_{\text{eff}} = \delta + Q_\Delta + uB_\Delta, \tag{S13}$$

and obeys

$$\delta \leq \Theta_{\text{eff}} < \delta + \frac{wu\Delta}{w+u} + u\Delta. \tag{S14}$$

Equivalently,

$$0 \leq T_{\Delta,\delta}(A) - \frac{\delta}{u} < \Delta \frac{2w+u}{w+u}. \tag{S15}$$

S3.4 Closed phase-retention law

An isolated event remains separately observable only if its trigger precedes the next competing update at horizon H, that is, if $T_{\Delta,\delta}(A) < H$. Thus

$$p_{\text{keep}} = \frac{1}{\Delta} \int_0^\Delta \mathbf{1}\{T_{\Delta,\delta}(A) < H\} \, \mathrm{d}A. \tag{S16}$$

This integral has an exact interval-length form. Write $H = m\Delta + h$, where $m = \lfloor H/\Delta \rfloor$ and $0 \leq h < \Delta$, and define

$$a_* = \frac{u\Delta}{w+u}, \qquad c_j = \frac{uj\Delta - \delta}{u+w}, \qquad I_m = [0, \Delta - h), \quad I_{m+1} = [\Delta - h, \Delta). \tag{S17}$$

When $h = 0$, I_{m+1} is empty. For an interval J, let $|J|$ denote its length and treat intersections with reversed endpoints as empty. Then

$$\begin{aligned} p_{\text{keep}} = \frac{1}{\Delta} \sum_{j \in \{m, m+1\}} \Big(& \big| I_j \cap [0, a_*] \cap [0, c_j] \big| \\ & + \mathbf{1}\{u(j-1)\Delta \geq \delta\} \big| I_j \cap [a_*, \Delta) \big| \Big). \end{aligned} \tag{S18}$$

The first term covers phases for which the pre-peak sample is the observed maximum. The second covers phases for which the post-peak sample is the maximum and phase cancels from the threshold inequality. Strict-before and by-horizon conventions differ only on a finite, zero-measure set of phases. The law depends on three dimensionless coordinates:

$$\beta = w/u, \qquad \kappa = \delta/(u\Delta), \qquad \eta = H/\Delta. \tag{S19}$$

It gives $p_{\text{keep}} = 0$ for $H \le \delta/u$ and $p_{\text{keep}} = 1$ for $H > \delta/u + \Delta(2w+u)/(w+u)$. Recursive retention must be evaluated on the complete path because omitting an event changes the next origin.

S3.5 Single-cell identifiability bound

Assume that speed is bounded by V and that only the endpoints of one sampling interval are observed. Path A remains at the origin, whereas path B moves a distance a along one ray and returns within the interval. The two paths produce identical fixes whenever $2a \le V\Delta$, although only path B contains a radial drawdown of amplitude a. Therefore, for any $\theta < V\Delta/2$, choosing $a \in (\theta, V\Delta/2]$ makes the presence of a θ event undecidable from the fixes. The condition $\Delta < 2\theta/V$ is necessary, but not sufficient, for identifying a continuous event of scale θ.

S3.6 Sampling-scale equivariance

RD is spatially scale equivariant: multiplying every coordinate and δ by $b > 0$ leaves event indices unchanged and multiplies all spatial outputs by b. Suppose the latent process is self-similar,

$$\mathbf{X}(ct) \overset{d}{=} c^H \mathbf{X}(t). \tag{S20}$$

A path sampled at interval $c\Delta$ and divided by c^H then has the same distribution as a path sampled at Δ. By spatial equivariance, its RD event law depends on the threshold only through $\delta/(c\Delta)^H$. Equal detector statistics at the two cadences therefore require

$$\frac{\delta_{\text{det}}(c\Delta)}{(c\Delta)^H} = \frac{\delta_{\text{det}}(\Delta)}{\Delta^H}, \qquad \delta_{\text{det}}(c\Delta) = c^H \delta_{\text{det}}(\Delta). \tag{S21}$$

This scaling relation is a property of the sampled detector. A biological retreat threshold is instead a fixed physical-scale hypothesis that must be evaluated against the identifiability bound

and an independent behavioural channel.

S3.7 Fixed-origin stability

If $\|r - \hat{r}\|_\infty \leq \epsilon_r$, then their running maxima differ by at most ϵ_r. Consequently the two drawdown functions differ by at most $2\epsilon_r$. If position and origin errors are each bounded by ϵ_x, radial error is bounded by $2\epsilon_x$ and drawdown error by $4\epsilon_x$.

For a confirmed discrete event (o, e, c) with a unique endpoint, define the endpoint and trigger margins

$$\gamma_E = r_e^{(o)} - \max_{o \leq i \leq c,\ i \neq e} r_i^{(o)}, \tag{S22}$$

$$\gamma_D = \min\left\{\delta - \max_{o<k<c} D_k^{(o)},\ D_c^{(o)} - \delta\right\}. \tag{S23}$$

If $\gamma_E > 4\epsilon_x$ and $\gamma_D > 4\epsilon_x$, every perturbed fix sequence within the error bound has the same endpoint e and first trigger c. If the corresponding margins hold at every successive common origin, induction gives the same complete recursive event sequence. For a censored final excursion, $\max_{o<k\leq n} D_k^{(o)} < \delta - 4\epsilon_x$ preserves the absence of a trigger. These are deterministic sufficient conditions, not distributional error assumptions.

For a continuous threshold crossing, uniqueness and transversality are also required to bound timing error. If the local drawdown slope is bounded below by $\nu > 0$, crossing-time error is at most the drawdown error divided by ν. When an event margin fails and the recursive origin changes, all subsequent radial coordinates can change; position error alone then provides no unconditional whole-path stability guarantee.

S4 Comparative stability and movement-process bridge

S4.1 Coordinate and traversal invariance

Let Q be orthogonal, $a \in \mathbb{R}^2$ and $b > 0$, and define $\widetilde{X}(t) = a + bQX(t)$ with $\widetilde{\delta} = b\delta$. Every radial distance and running maximum is multiplied by b, so the drawdown inequality and all recursive endpoint indices remain unchanged. Translation, rotation and reflection are therefore exact invariances, whereas spatial scaling is an equivariance.

Let h be a strictly increasing continuous bijection and $\widetilde{X}(t) = X(h(t))$. The ordered path locations and all suprema over corresponding prefixes remain unchanged. Spatial endpoints and

event order are identical, while endpoint and trigger times are mapped by h^{-1}. Continuous RD geometry is therefore invariant to traversal speed. Fixed-time boundaries do not share this property. Discrete RD retains the coordinate invariances but not arbitrary time-reparameterization invariance because reparameterization changes the sampled path.

S4.2 Amplitude stability and local-angle impossibility

Suppose $X(t) = Y(t) + Z(t)$ with $\sup_t \|Z(t)\| \le b$. At a common origin o,

$$\big|\|X(t) - X(o)\| - \|Y(t) - Y(o)\|\big| \le \|Z(t) - Z(o)\| \tag{S24}$$

$$\le 2b. \tag{S25}$$

The supremum operator is 1-Lipschitz in the sup norm, so the two running maxima differ by at most $2b$ and the drawdowns by at most $4b$. Endpoint and threshold margins greater than $4b$ preserve a fixed-origin event; successive common-origin margins give recursive stability by induction.

Local heading has no corresponding amplitude-only bound. Set

$$Y(t) = (vt, 0), \qquad Z(t) = (0, b\sin\omega t). \tag{S26}$$

The perturbation remains within a strip of half-width b, yet its tangent is $(v, b\omega\cos\omega t)$. For any fixed turn threshold below π, increasing ω creates arbitrarily many alternating heading changes above that threshold. Stabilizing an angle-based method therefore requires additional assumptions about frequency, curvature, smoothing or minimum displacement that are absent from the amplitude bound.

For sampled position error, the observed displacement differs from a true displacement of length ℓ_i by a vector of norm at most $2\epsilon_x$. If $\ell_i > 2\epsilon_x$, the error ball does not contain the origin and its tangent geometry gives

$$|\widehat{\psi}_i - \psi_i| \le \arcsin(2\epsilon_x/\ell_i). \tag{S27}$$

Turn-angle error is at most the sum of the two adjacent heading-error bounds. No non-trivial bound remains when $\ell_i \le 2\epsilon_x$. This contrasts with the $4\epsilon_x$ RD drawdown bound, which is expressed against an accumulated spatial threshold.

S4.3 Geometric visibility of a directional update

A biological update and its geometric visibility to RD are distinct. Let a candidate update occur at the current radial maximum $M > 0$, with $0 < \delta \leq M$. Suppose that the first post-update displacement has length $L > 0$ and forms an angle θ with the outward radial direction. The law of cosines gives

$$r_1^2 = M^2 + L^2 + 2ML\cos\theta. \tag{S28}$$

The first post-update fix confirms a drawdown of at least δ if and only if $r_1 \leq M - \delta$, or

$$\cos\theta \leq c(M, L, \delta) = \frac{\delta^2 - 2M\delta - L^2}{2ML}. \tag{S29}$$

For an isotropic post-update direction, θ is uniform modulo 2π and the conditional visibility probability is

$$p_{\text{vis}}(M, L, \delta) = \begin{cases} 0, & c < -1, \\ 1 - \pi^{-1}\arccos(c), & -1 \leq c \leq 1, \\ 1, & c > 1. \end{cases} \tag{S30}$$

For a movement model with joint density $f(L, \theta \mid M, \text{update})$, visibility is the integral of the same indicator against f. For a sequence of post-update fixes before the next update, replace r_1 by the minimum sampled radius in that interval. The complement of p_{vis} includes genuine updates that continue outward or whose redirection is too small or too brief to produce the declared retreat. It is therefore a sensitivity term for the RD estimand, not a false-negative rate for all possible state changes.

S4.4 Event recovery and retained information

Let $Y = (e_1, \ldots, e_K)$ denote an externally defined sequence of navigation-update indices on a sampled latent path. The label is specified by the movement generator or an outcome-isolated observable and is not defined by RD. For each epoch beginning at origin $o_k = e_{k-1}$, assume (i) e_k is the unique radial maximum before the next update, (ii) its radial lead over every competing sampled point exceeds $4b$, (iii) drawdown remains below $\delta - 4b$ before the designated confirmation and reaches at least $\delta + 4b$ there, and (iv) confirmation occurs before the next update. If the realized path differs from the latent navigation path by at most b in position, the amplitude-stability bound preserves the endpoint and first trigger in that epoch. Because the

recovered endpoint becomes the next common origin, induction gives

$$S_{\mathrm{RD}}(X) = Y. \tag{S31}$$

Thus RD is a lossless representation of this event sequence within the stated margin class and retains all label information,

$$I(Y; S_{\mathrm{RD}}) = H(Y). \tag{S32}$$

This result establishes information-theoretic optimality for the declared radial-renewal target, not universal optimality across biological targets.

More generally, suppose an estimator $\widehat{Y} = g(S_{\mathrm{RD}})$ has sequence-classification error at most $\eta \leq 1 - K_Y^{-1}$ over K_Y possible labels. Fano's inequality gives

$$I(Y; S_{\mathrm{RD}}) \geq H(Y) - h_2(\eta) - \eta \log(K_Y - 1), \tag{S33}$$

where h_2 is binary entropy. The bound converts localization error measured on independent truth into a lower bound on retained event information.

The scope of this comparison is exact. A boundary-time variable generated by a fixed schedule independent of Y has $I(Y; S_{\mathrm{time}}) = 0$ for event phase. A thresholded local-angle detector has no uniform recovery guarantee under the same amplitude assumption, as shown by two finite sampled constructions. First, for any angle threshold $\alpha \in (0, \pi)$, take

$$X_j = \{R\sin(jq),\, R[1 - \cos(jq)]\}, \qquad j = 0, \ldots, 2m, \qquad q = \pi/m < \alpha. \tag{S34}$$

The path reaches radial distance $2R$ at $j = m$ and returns to its origin at $j = 2m$, so RD detects a radial renewal of any scale $\delta \leq 2R$; every interior sampled turn is only $q < \alpha$, so the angle rule detects none. Second, let $Y_i = (ia, 0)$ and $X_i = (ia, b(-1)^i)$. The realized path remains within b of a monotone skeleton, yet successive headings alternate by $2\arctan(2b/a)$. Choosing $a < 2b/\tan(\alpha/2)$ creates an arbitrary number of angle-threshold crossings as the sequence length grows. Because the latent skeleton has zero drawdown and the amplitude bound gives $\|D^X - D^Y\|_\infty \leq 4b$, RD creates none at any $\delta > 4b$. Thus a local-angle rule can both miss a smooth macro reversal and fragment bounded local motion, whereas RD separates these two constructions at the declared spatial scale. These results do not apply to an unrestricted history-dependent model. An HMM or change-point model supplied with radial drawdown or an

equivalent path-history feature can emulate RD; without restrictions on the available features, no data-processing argument can rank all possible segmentation rules.

For a binary external event B and a non-differential binary boundary indicator $\widetilde{B}$, let sensitivity and specificity be Se and Sp. Since

$$\mathbb{E}[\widetilde{B} \mid B] = (1 - Sp) + (Se + Sp - 1)B, \tag{S35}$$

the event–outcome covariance is attenuated by the Youden factor $J = Se + Sp - 1$. Under equal outcome effects and approximately independent effective events, the leading non-centrality scale is proportional to $N_{\text{eff}} J^2$. We therefore compare methods on held-out external truth using localization risk, mutual information or proper scoring loss together with $N_{\text{eff}} J^2$; a smaller ecological p value is not itself a calibration criterion.

S4.5 Self-similar first-drawdown scale

For an H-self-similar process $X(at) \overset{d}{=} a^H X(t)$, radial distance, its running maximum and drawdown share the same scaling. Therefore

$$\tau_{a^H \delta} \overset{d}{=} a\tau_\delta, \qquad \tau_\delta \overset{d}{=} \left(\frac{\delta}{\sigma}\right)^{1/H} \tau_1, \tag{S36}$$

where σ denotes the spatial amplitude in the normalized representation. If a local bout lasts T_L and the tolerated probability of an internal RD trigger is q, then

$$\Pr(\tau_\delta \le T_L) \le q, \tag{S37}$$

$$\delta \ge \sigma \left[\frac{T_L}{F_{\tau_1}^{-1}(q)}\right]^H. \tag{S38}$$

For Brownian local motion, $H = 1/2$ and the process amplitude is proportional to $\sqrt{D}$, giving $\tau_\delta \propto \delta^2/D$ and $\delta_{\text{local}} \propto \sqrt{DT_L}$.

S4.6 Biological feasible band

Following the movement-ecology distinction among navigation, motion, internal state and external factors, write the observed trajectory during goal epoch k as

$$X(t) = \Gamma(t; G_k) + \xi_k(t) + \varepsilon(t), \qquad C_k \le t < C_{k+1}. \tag{S39}$$

G_k denotes the navigation goal and C_k the time at which it changes; Γ is the macro-scale goal-directed path, ξ_k is local movement and ε is observation error. RD is informative about a goal update only when the update creates a radial apex followed by sufficient retreat.

Let δ_{noise} enforce the predeclared observation-error budget, $V\Delta/2$ be the single-cell lower bound and $\delta_{\text{local}}(q)$ suppress internal local-bout triggers. Let A_{rev} be available macro retreat after the goal update, m_{micro} its stability margin, and H_c the time to the next competing update. Combining the retention and stability inequalities gives

$$\delta_{\min} = \max\{\delta_{\text{noise}}, V\Delta/2, \delta_{\text{local}}(q)\}, \tag{S40}$$

$$\delta_{\max} = \min\left\{A_{\text{rev}} - m_{\text{micro}}, u\left[H_c - \Delta\frac{2w+u}{w+u}\right]\right\}. \tag{S41}$$

A biologically interpretable macro-RD scale requires $\delta_{\min} < \delta_{\max}$. This condition separates three limitations: inadequate observation, insufficient separation between local and macro-scale motion, and a response that does not produce enough retreat before the next update.

Each quantity in this band has an operational source. Static periods, replicate positions, parallel observables or an explicit location-error model can estimate δ_{noise}. A high-frequency pilot provides a conservative speed envelope V, rise and retreat speeds w, u, and empirical distributions of τ_δ and H_c. Complete local bouts or a fitted self-similar local model can estimate $\delta_{\text{local}}(q)$ and m_{micro}. An independent activity clock, externally generated model truth or held-out annotated update provides A_{rev}, event recall and merge information. An unavailable term cannot be replaced by an outcome-maximizing scan; the supported claim is then restricted to the quantities that were estimated.

S4.7 Detection-theoretic scale selection

Let $F(\delta)$ denote the false-event or local-fragmentation rate on a noise/local null, $R(\delta)$ recall on detector-independent event truth, $M_g(\delta)$ the proportion of adjacent true events merged, $P_{\text{keep}}(\delta)$ confirmation before the next competing update and $N(\delta)$ usable event or individual coverage. Before ecological outcomes are read, specify false-event, miss and merge budgets α, β, μ, a confirmation target π and minimum coverage n_0. With upper and lower uncertainty bounds U and L, define

$$\mathcal{D} = \left\{\delta : U_F(\delta) \le \alpha, \quad L_R(\delta) \ge 1-\beta, \quad U_{M_g}(\delta) \le \mu, \quad L_{P_{\text{keep}}}(\delta) \ge \pi, \quad L_N(\delta) \ge n_0\right\}. \tag{S42}$$

The admissible scale set is $\mathcal{D} \cap [\delta_{\min}, \delta_{\max}]$. If the objective is maximum spatial resolution and event retention, choose the smallest point inside a continuous admissible component after excluding candidates whose confidence intervals touch a failure boundary. If robustness to both edges is the stated objective, choose the log-δ point that maximizes the minimum standardized margin to all active constraints. Both rules are fixed before outcome inspection and the complete admissible band is reported.

The procedure has three valid outcomes. With both a null and independent truth, it can select an operational scale under explicit error budgets. With a null but no event truth, it can identify a lower observation or coarse-graining bound but cannot estimate recall or merging of biological updates. With neither, RD remains a predeclared multiscale description, and biological outcomes can test the complete scale band with simultaneous inference; the same records cannot both select a significant δ and validate it. A cadence-invariant θ_{bio} requires a separate shared-threshold model and outer held-out animals or deployments.

S5 Benchmark design

S5.1 Known-truth mathematical verification

The information-limit experiment generated 100,000 pairs of continuous paths with identical sampled fixes. One path in each pair remained event-free, whereas the other completed an outward–return excursion within a sampling cell under the same speed bound. Fix-wise mutual information about event presence was therefore zero, and the equal-prior Bayes error was 0.5.

The isolated-peak experiment crossed three event scales, three peak asymmetries and eight values of cadence relative to event duration, yielding 72 configurations. Each configuration used 100,000 deterministic phases. Pre-specified checks compared the empirical Q_Δ distribution, normalized mean, endpoint-side probability, trigger identity, delay bound and dimensionless scale collapse with their analytic values.

The recursive-competition experiment used triangular radial renewals with 64 known endpoints per path. It crossed three absolute scales, two base cadences, six values of ρ, six values of λ and 128 phases, retaining the 156 feasible absolute parameter combinations. Equal (ρ, λ) pairs tested scale collapse; increasing λ at fixed ρ tested the prediction that competition before confirmation merges events.

The independent training and evaluation partitions are listed in Table S1.

Table S1: Independent training and evaluation partitions

Stage	Seeds or paths	Use
Training	11000–11039	Comparator parameters only
Validation	12000–12019	Directional gates; no retuning
Native evaluation	21000–21099	100 held-out paths per process
Observation evaluation	21000–21039	40 paths per process under 160 conditions
Invariance	100 radial-renewal paths	800 variant–transform comparisons

The observation grid combined eight cadences (1–128 s in powers of two), five Gaussian-noise levels (0–0.20δ) and four sampling phases. Four task families and 40 paths yielded 25,600 path–condition observations. Potts parameters were $W = 8$ and 15°; Gaussian change-point minimum segment was 32 s; FPT radius was 32 m; HMM optimization used three initializations, at most 50 EM iterations and relative tolerance 10^{-5}.

Construct-matched implementation controls for each declared estimand are reported in Table S2.

Table S2: Construct-matched native-resolution positive controls

Task	Target object	Object-matched score
Radial renewal	Endpoint	RD/RD-L F1 1.000
Distribution change	Persistent change	Gaussian CP F1 1.000
Local heading turn	Turn location	Potts F1 0.999
Two-state process	Hidden-state transition	HMM F1 1.000
Two-state process	Residence intensity	FPT AUROC 0.960
Straight path	Radial-renewal absence	RD/RD-L 0 h^{-1}

Each scripted process was paired with its declared estimator. These scores verify the implementations and distinguish their estimands; they do not provide a detector-independent comparison of biological event recovery.

S5.2 Detector-independent mechanistic applicability benchmark

Four published movement-model classes generated trajectories and model-native event truths without using RD or any comparator. The Fleming maximum-entropy central-place model used the exact linear-SDE transition of an underdamped random oscillator; its target event was the fixed-centre change in radial velocity from outward to inward motion [3]. The Duchesne biased correlated random walk combined directional persistence with attraction towards a fixed target; its event was first entry into the published target neighbourhood [4]. The Pirotta model class combined macro states for outward travel, return travel and movement towards a moving

attractor with micro states for transit and residence; its events came directly from the generating state sequence [5]. The Preisler quadratic-potential model used the exact Ornstein–Uhlenbeck transition towards a point attractor; its event was first entry into the target neighbourhood [6]. A predeclared grid of trip duration, emission separation and transition intensity represented the Pirotta model class across 18 formal parameter cells.

The comparison included RD, Potts state-change step detection, local turn peaks, fixed-time boundaries, projection extrema, FPT, Gaussian change points and a two-state HMM. A development partition of 1,380 trajectories selected all method parameters and the strongest non-RD reference for each model–estimand pair. These choices were frozen before evaluating 8,400 trajectories across 39 parameter cells. Training and evaluation seeds were disjoint. The formal benchmark produced 182,400 path–estimand–method records and satisfied all 17 pre-specified integrity checks.

Events were matched one-to-one within model-specific tolerances. Cell-macro means gave every predeclared parameter cell equal weight. For each RD–reference contrast, paths were bootstrapped within cells and cell means were then averaged, preserving the intended generalization across the mechanism grid.

Across the independent mechanisms, RD's comparative advantage was specific to the central-place apex and return onset. RD–reference contrasts are reported in Table S3, and the complete estimator matrix in Table S4.

Table S3: RD relative to the training-frozen strongest non-RD reference on model-native event truths. Intervals are 95% path-bootstrap intervals stratified within parameter cell.

Model-native truth	Frozen reference	RD–reference F1	95% interval	Positive cells
Fleming central-place apex	Potts	0.0268	0.0226–0.0309	9/9
Pirotta return onset	Projection extrema	0.2460	0.2370–0.2552	18/18
Pirotta all macro switches	Potts	0.0386	0.0296–0.0471	10/18
Pirotta enter moving attractor	Potts	−0.0152	−0.0196–−0.0108	3/18
Pirotta leave moving attractor	Potts	−0.0236	−0.0295–−0.0174	1/18
Pirotta micro-state switch	HMM	−0.1050	−0.1109–−0.0990	4/18
Duchesne target arrival	HMM	−0.1479	−0.1771–−0.1187	0/3
Preisler target arrival	FPT	−0.0691	−0.0780–−0.0600	0/9

Table S4: Cell-macro mean F1 for all methods in the formal mechanistic benchmark. Bold values are the largest held-out mean within each model-native truth; they did not affect training-frozen reference selection.

Model-native truth	RD	Potts	Turn	Time	Projection	FPT	Gaussian CP	HMM
Fleming central-place apex	**0.830**	0.803	0.716	0.595	0.754	0.725	0.073	0.436
Pirotta return onset	**0.395**	0.128	0.057	0.054	0.149	0.052	0.030	0.054
Pirotta all macro switches	**0.338**	0.299	0.267	0.255	0.299	0.259	0.049	0.217
Pirotta enter moving attractor	0.165	**0.180**	0.140	0.149	0.144	0.150	0.049	0.130
Pirotta leave moving attractor	0.140	**0.164**	0.131	0.145	0.139	0.147	0.041	0.123
Pirotta micro-state switch	0.294	0.308	0.391	0.339	0.295	0.360	0.060	**0.399**
Duchesne target arrival	0.133	0.114	0.092	0.113	0.099	0.260	0.248	**0.281**
Preisler target arrival	0.032	0.058	0.069	0.077	0.033	**0.101**	0.016	0.017

S6 Observation-phase results

Analytic and numerical results agreed in all 72 isolated-peak configurations. The maximum Kolmogorov–Smirnov distance from the analytic peak-loss distribution was 1.0×10^{-5}, the maximum normalized mean error was 2.5×10^{-11}, and the maximum trigger-identity error was 5.68×10^{-14}; no delay-bound violation occurred. An independent implementation of the closed retention law matched a 200,000-phase midpoint enumeration across 90 parameter settings, with a maximum absolute discrepancy of 2.86×10^{-6}.

All 156 recursive conditions satisfied the pre-specified checks for geometry, scale collapse, competition direction and event cardinality. The maximum normalized spread among conditions with the same ρ and λ was zero. At $\rho = 1$, recall decreased from 1.000 at $\lambda = 0.5$ to 0.500 at $\lambda = 1$; at $\rho = 2$ and 4, it decreased from 1.000 at $\lambda = 0.5$ to zero at $\lambda = 2$. These tests connect

the isolated observation kernel to recursive event loss without using animal labels.

Within $\rho \leq 0.25$ and $\lambda \leq 0.25$, both implementations had a median F1 of 1.000 with path-cluster interval $[1.000, 1.000]$. For primary RD, Spearman correlations were -0.517 with ρ and -0.526 with λ; the corresponding fractional-logit coefficients per doubling were -1.347 and -1.417. At noise level 0.20δ, mean and median F1 within this region were 0.920 and 0.960. At 128 s without added noise, F1 was 0.957, the count ratio was 0.917 and median timing error was 30.4 s. All transformation comparisons returned zero event-time error.

In a separate five-scale Crozet holdout containing 18,008 native sampled-fix events and 468,208 event–configuration observations, a frozen ρ model improved log loss over a constant model by 0.1611 (95% bird-bootstrap CI 0.1121–0.2101). Recall at $\rho \leq 0.5$ exceeded recall at $\rho \geq 2$ by 0.5732 (0.5115–0.6322). This result establishes a cross-scale ordering by ρ; absolute recovery is reported separately for each scale and cadence.

S7 Real-data scale ledger

Uesaka et al. describe the Crozet wandering-albatross data archived in Dryad [1, 2]; Gunner et al. describe the San Lorenzo Magellanic-penguin data archived in Figshare [7, 8].

S7.1 Crozet observation scales

The 5-Hz micro-scale experiment treated GNSS position and horizontal Doppler velocity from the same logger as parallel observables. Records required finite latitude, longitude, north and east Doppler velocity, horizontal accuracy in $(0, 20]$ m, and intervals of 0.15–0.25 s within continuous 5-Hz blocks. Candidate scales were 1, 2, 4, 8, 16, 32 and 64 m. On 18 training birds, the selected scale was the smallest candidate with position–Doppler recall and endpoint F1 both at least 0.80, a 95% upper bound on reference-null extra position events no greater than 10^{-3} s^{-1}, and sufficient exposure and individual coverage. One metre was the first scale to meet all criteria. On five held-out birds, it gave recall 0.946, F1 0.920 and an extra-event upper bound of 8.09×10^{-4} s^{-1} over 42.79 h.

The 1-s cadence experiment used 16.5 and 72.1 m as complementary observation scales. A movement-only stability and event-retention analysis identified 16.5 m as a fine-scale candidate; across 5,000 bird bootstraps, 16.5 and 33 m were selected in 64.14% and 34.86% of resamples. This data-specific analysis defines 16.5 m as an observation scale. The 72.1-m value came

from an independently frozen movement-only analysis of the public 120-s wandering-albatross trajectories released with Goto et al. [9]; environmental response columns were excluded from scale selection. It therefore enters the Crozet analysis as a transferred observation scale. The event-clock band of 50, 100, 200 and 400 m was written and hashed before the Doppler transition table was opened; simultaneous inference across all four scales replaced single-scale selection.

S7.2 GNSS-Doppler flight-to-low-speed event clock

The event clock labelled `landing_gps_conservative` in the archived code was constructed from raw horizontal Doppler speed, independently of pressure, wet/dry state and the position coordinates used by RD. Source timestamps were converted from the logger time basis to UTC. After applying the same ≤ 20 m horizontal-accuracy filter, records were split at gaps greater than 1 s. A 20-fix forward mean provided a nominal 4-s speed series. An onset required smoothed speed to remain at least 4 m s^{-1} throughout the preceding 150 fixes and below 4 m s^{-1} throughout the following 150 fixes; the corresponding median raw speeds had to exceed 5 m s^{-1} before the onset and remain below 2.5 m s^{-1} afterwards. Each side had to span no more than 31 s, and starts separated by less than 10 s were consolidated. These rules yielded 401 conservative putative flight-to-water transitions from 21 birds and 24 deployments.

RD endpoints and transition times were matched one-to-one within each predeclared interval. The primary interval was 0–120 s after the endpoint; $\pm$120 s and 0–600 s provided timing sensitivities. Segment-common circular shifts preserved the event series and its temporal dependence while breaking alignment with the movement boundaries. Bird-level estimates received equal weight; 2,000 common shifts generated the null, and 5,000 common bird resamples generated four-scale max-$|T|$ simultaneous intervals.

S7.3 Official BCPA comparison

The Crozet event-clock benchmark also used the official CRAN `bcpa` 1.3.2 implementation [10]. Movement-only tracks were sampled at 5 s and analysed using persistence velocity, a primary 30-fix window, a two-fix window step, $K = 2$, $\tau =$ TRUE, range 0.6 and a 15-s change-point clustering width. Windows of 50 and 80 fixes and clustering widths of 5 and 30 s were frozen as sensitivity settings. Within every segment and RD scale, BCPA points ranked by official window support and RD endpoints were deterministically reduced to the same post-edge count. Both methods used the same circular shifts and bird-bootstrap resamples.

S7.4 San Lorenzo movement and dive-effort preprocessing

The source file `Penguins.csv` contained 3,072,835 rows at a nominal cadence of 1 Hz. The movement pass read only `Timestamp`, `ID`, `DR_longitude` and `DR_latitude`. Dead-reckoned WGS84 coordinates were projected to UTM zone 20S (EPSG:32720), and tracks were split at every non-1-s timestamp interval or non-finite coordinate. Fragments shorter than three fixes were discarded. The resulting movement data contained 3,072,830 rows from 27 birds in 49 continuous segments. Non-terminal sampled-fix RD endpoints were computed at 20, 50 and 100 m; the focal 50-m scale and its neighbours were frozen before activity columns were opened.

The behaviour pass then read `Behaviour`, `Depth_sm` and `VeDBA_sm` from the unchanged source file and verified the segment structure and file hash. Dive state was defined by `Behaviour == ''dive''`. For each endpoint, absolute dive-effort change compared dive fractions in equal pre- and post-boundary windows of 5, 10 or 20 min. Each side required at least 80% non-missing behaviour, and terminal or censored boundaries were excluded for every method. Segment-common circular shifts preserved local autocorrelation and endpoint configuration. Inference used equal bird weights, 2,000 shifts and 2,000 bird-bootstrap resamples.

Scale provenance and supported interpretation for both animal datasets are summarized in Table S5.

Table S5: Detection scales, selection sources and supported interpretations in the two animal datasets [1, 2, 7, 8]

Dataset	Cadence	Scale	Selection source	Interpretation
Crozet albatross	5 Hz	1 m	Position–Doppler error budgets; 18 training and 5 holdout birds	GNSS observation floor
Crozet albatross	1 s	16.5 m	Behaviour-blind movement stability–retention candidate	Fine event-timescale experiment
Crozet albatross	1 s	72.1 m	Frozen movement-only analysis of public 120-s Goto tracks [9]	Transferred observation-scale test

Dataset	Cadence	Scale	Selection source	Interpretation
Crozet albatross	1 s	50–400 m	Frozen doubling grid before Doppler clock opened	Flight-to-low-speed localization band
San Lorenzo penguin	1 s	50 m	Position quality, 30-s displacement and coverage before activity columns	Focal dive-effort response
San Lorenzo penguin	1 s	20/100 m	Predeclared neighbouring scales	Behavioural scale-band sensitivity

S8 Full real-data contrasts

Penguin dive-effort contrasts across the predeclared scale band are reported in Table S6.

Table S6: San Lorenzo Magellanic-penguin 10-min dive-effort contrasts [7, 8]

Variant	δ (m)	Qualified events	Effect (95% CI)	Positive birds
RD	20	1093	0.0249 (0.0056, 0.0442)	20/27
RD	50	469	0.0492 (0.0226, 0.0764)	22/27
RD	100	232	0.0744 (0.0379, 0.1073)	19/27
RD-L	50	470	0.0491 (0.0235, 0.0766)	22/27

The primary equal-budget Crozet comparison with official BCPA is reported in Table S7.

Table S7: Equal-budget Crozet GNSS-Doppler event-clock comparison with official BCPA. Entries are differences between RD and BCPA observed-minus-circular-null F1; intervals are four-scale max-$|T|$ simultaneous 95% bird-bootstrap intervals.

δ (m)	RD–BCPA F1	Simultaneous 95% interval
50	0.0593	0.0301–0.0884
100	0.0725	0.0405–0.1044
200	0.0873	0.0488–0.1257
400	0.0734	0.0219–0.1250

The pre-specified BCPA window and clustering-width sensitivities are reported in Table S8.

Table S8: Pre-specified BCPA window and clustering-width sensitivity. Each setting has a separate four-scale max-$|T|$ correction; all 20 intervals exclude zero.

BCPA setting	δ (m)	RD–BCPA F1	Simultaneous 95% interval
W30/CW15	50	0.0593	0.0301–0.0884
W30/CW15	100	0.0725	0.0405–0.1044
W30/CW15	200	0.0873	0.0488–0.1257
W30/CW15	400	0.0734	0.0219–0.1250
W50/CW15	50	0.0618	0.0382–0.0853
W50/CW15	100	0.0714	0.0405–0.1024
W50/CW15	200	0.1022	0.0647–0.1398
W50/CW15	400	0.0918	0.0441–0.1395
W80/CW15	50	0.0539	0.0251–0.0828
W80/CW15	100	0.0654	0.0285–0.1022
W80/CW15	200	0.0861	0.0423–0.1300
W80/CW15	400	0.0856	0.0395–0.1316
W30/CW5	50	0.0391	0.0023–0.0758
W30/CW5	100	0.0727	0.0428–0.1025
W30/CW5	200	0.1028	0.0608–0.1448
W30/CW5	400	0.0863	0.0274–0.1451
W30/CW30	50	0.0615	0.0359–0.0870
W30/CW30	100	0.0797	0.0495–0.1100
W30/CW30	200	0.0883	0.0531–0.1236
W30/CW30	400	0.0757	0.0225–0.1289

The scale-wise Crozet alignment above the circular null is reported in Table S9.

Table S9: Crozet wandering-albatross flight-to-low-speed alignment above the circular null. Intervals are multiscale max-T 95% bird-bootstrap intervals.

δ (m)	Matches/401	Endpoint–transition lag (s)	RD–null F1
50	237	23.0	0.0566 (0.0335, 0.0797)
100	173	29.4	0.0810 (0.0524, 0.1097)
200	119	44.8	0.1073 (0.0726, 0.1419)
400	64	39.3	0.0985 (0.0505, 0.1465)

Equal-budget fixed-time and local-turn comparisons are reported in Table S10.

At 20 and 100 m, RD and RD-L endpoints matched exactly. At 50 m, the exact endpoint F1 was

Table S10: Equal-budget Crozet flight-to-low-speed comparisons. Entries are RD-minus-comparator F1 with multiscale max-T 95% bird-bootstrap intervals; boundary counts are exactly equal within each segment and evaluation window.

δ (m)	RD–fixed-time F1	RD–local-turn F1
50	0.0492 (0.0250, 0.0734)	0.0532 (0.0334, 0.0729)
100	0.0737 (0.0436, 0.1038)	0.0784 (0.0506, 0.1061)
200	0.1113 (0.0732, 0.1494)	0.1019 (0.0655, 0.1384)
400	0.1017 (0.0525, 0.1510)	0.1022 (0.0529, 0.1516)

0.996872; within 1 s, precision was 1.000 and recall 0.997917. The direct RD-minus-RD-L difference in the 50-m, 10-min behaviour effect was 0.000177 (95% CI 0–0.000530).

The paired 50-m penguin method comparisons are reported in Table S11.

Table S11: Paired bird-level comparison of 50-m RD with rate-matched or geometric boundaries for 10-min dive effort

Comparator	RD minus comparator	95% bird CI	Birds with larger RD effect
Rate-matched fixed time	0.058806	0.031764–0.085837	21/27
Equal-count turn angle	0.018764	−0.026895–0.060841	14/25
Projection extrema	−0.018892	−0.058404–0.021198	13/27

The fixed-time comparison tests whether event abundance and temporal coverage alone can generate the depth association. Turn-angle and projection extrema are reported as complementary geometric objects rather than as interchangeable estimators of radial renewal.

Native-event recovery across Crozet sampling cadences is reported in Table S12.

Table S12: Crozet wandering-albatross native-event recovery across cadence [1, 2]

Cadence (s)	Median recall	Median compression
1	1.000	0.000
2	0.953	0.024
4	0.861	0.089
8	0.689	0.227
16	0.521	0.395
32	0.376	0.569
64	0.267	0.701
128	0.190	0.794

The 72.1-m cadence and observation-model sensitivity for the Crozet flight-to-low-speed clock is reported in Table S13.

Table S13: Crozet wandering-albatross flight-to-low-speed contrasts at 72.1 m [1, 2]

Variant	Cadence and window	Post–pre transition probability	95% bird CI
RD	1 s, 2 min	0.0679	0.0366–0.1095
RD-L	1 s, 2 min	0.0695	0.0387–0.1105
RD	1 s, 10 min	0.1853	0.1240–0.2619
RD-L	1 s, 10 min	0.1850	0.1235–0.2622
RD	120 s, 2 min	-0.0086	-0.0340–0.0167
RD-L	120 s, 2 min	0.0057	-0.0151–0.0263
RD	120 s, 10 min	0.0559	0.0000–0.1164
RD-L	120 s, 10 min	0.0836	0.0252–0.1462

S9 Reproducibility and data access

The archived analysis environment used Python 3.11.5 with NumPy 1.25.2, pandas 2.1.2, SciPy 1.11.4 and pyproj 3.6.1. Figures were generated with matplotlib 3.7.3 and seaborn 0.13.2. The official BCPA analysis used R 4.2.0 and CRAN `bcpa` 1.3.2; the clean source tarball and package-version checks were stored with the batch manifest. Random seeds for the formal simulation partitions are listed in Table S1. Circular-shift and bootstrap seeds were deterministic hashes of the dataset, method, scale and window identifiers. The file-to-result map is provided in Table S14.

Table S14: File-to-result reproducibility map. Paths are relative to the versioned analysis repository.

Manuscript element	Primary code	Frozen result inputs
Discrete RD and sampling theorems	`experiments/rd_validity_theory_v1/src/`	`runs/discrete_audit_v1`, `runs/sampling_kernel_closed_form_v1`
Figures 2; Tables S3–S4	`experiments/recognized_model_benchmark_v1/src/`	`formal/run_v1/`
Figure 3	`experiments/unified_task_benchmark_v1/src/`	`formal/observation_phase_v1_20260818/`
Figure 4; Table S12	Crozet cadence scripts in the archived analysis workflow	`UESAKA_RD_CADENCE_MECHANISM_V1/`
Figure 5a; Tables S7–S10	`external_sensor_truth_benchmark_v1/src/`; `landing_step_method_benchmark_v2/src/`	`formal/event_clock_v3/`; `formal/analysis_v2/`
Figure 5b–d; Tables S6, S11	`penguin_rd_p_l_behavior_v1/src/`	`runs/formal_20260818_173419/`
All figures	`scripts/build_manuscript_figures.py`	Frozen CSV and compressed event-table inputs listed above

The canonical sampled-fix RD implementation, exact sampling utilities, unit tests and worked example are openly available in the versioned Radial Drawdown software repository (v0.1.0). The release was verified against the frozen manuscript implementation on 1,000 random tracks and 90 analytic sampling-kernel cases. The raw Crozet wandering-albatross multisensor data are available from Dryad [2], and the San Lorenzo Magellanic-penguin biologging data are available from Figshare [8].

References

[1] Uesaka L, Goto Y, Naruoka M, Weimerskirch H, Sato K, Sakamoto KQ. Wandering albatrosses exert high take-off effort only when both wind and waves are gentle. eLife. 2023;12:RP87016. https://doi.org/10.7554/eLife.87016.3.

[2] Uesaka L, Goto Y, Naruoka M, Weimerskirch H, Sato K, Sakamoto KQ. Behavioral datasets of wandering albatrosses collected at Possession Island, Crozet, France, in 2019 and 2020. Dryad. 2023;p. Dataset. https://doi.org/10.5061/dryad.tx95x6b2j.

[3] Fleming CH, Subaşı Y, Calabrese JM. A maximum-entropy description of animal movement. Physical Review E. 2015;91(3):032107. https://doi.org/10.1103/PhysRevE.91.032107.

[4] Duchesne T, Fortin D, Rivest LP. Equivalence between step selection functions and biased correlated random walks for statistical inference on animal movement. PLOS ONE. 2015;10(4):e0122947. https://doi.org/10.1371/journal.pone.0122947.

[5] Pirotta E, Edwards EWJ, New L, Thompson PM. Central place foragers and moving stimuli: a hidden-state model to discriminate the processes affecting movement. Journal of Animal Ecology. 2018;87(4):1116–1125. https://doi.org/10.1111/1365-2656.12830.

[6] Preisler HK, Ager AA, Wisdom MJ. Analyzing animal movement patterns using potential functions. Ecosphere. 2013;4(3):32. https://doi.org/10.1890/ES12-00286.1.

[7] Gunner RM, Quintana F, Tonini MH, Holton MD, Yoda K, Crofoot MC, et al. Penguins exploit tidal currents for efficient navigation and opportunistic foraging. PLOS Biology. 2025;23(7):e3002981. https://doi.org/10.1371/journal.pbio.3002981.

[8] Gunner RM, Quintana F, Tonini MH, Holton MD, Yoda K, Crofoot MC, et al. Penguins exploit tidal currents for efficient navigation and opportunistic foraging: ocean currents and penguin biologging data. Figshare. 2025;p. Dataset. https://doi.org/10.6084/m9.figshare.28517873.

[9] Goto Y, Weimerskirch H, Fukaya K, Yoda K, Naruoka M, Sato K. Albatrosses employ orientation and routing strategies similar to yacht racers. Proceedings of the National Academy of Sciences of the United States of America. 2024;121(23):e2312851121. https://doi.org/10.1073/pnas.2312851121.

[10] Gurarie E, Andrews RD, Laidre KL. A novel method for identifying behavioural changes in animal movement data. Ecology Letters. 2009;12(5):395–408. https://doi.org/10.1111/j.1461-0248.2009.01293.x.