Surface-Return Peaks and Channel Sources

Scientific and Engineering Release Notes for `wepp_261010`

Roger Lew
University of Idaho
rogerlew@uidaho.edu

2026-10-10 Document revision 1

PUBLICATION APPROVED – 2026-10-10

Release scope: This report describes changes made after `wepp_260803` through `wepp_261010` in the University of Idaho's `wepp-forest` fork of WEPP. It is not a release note for the upstream USDA WEPP model.

Abstract

WEPP 261010 combines an hourly surface-return peak estimator with explicit component transport, bounded channel-source construction, channel-profile continuation and a correction to source-sample normalization. The changes address two distinct problems: returned water concentrated into a short rainfall-excess duration at the hillslope, and an existing indexing defect that could amplify a small separated rainfall component into an unrealistic channel pulse. This report explains the legacy and revised calculations, their hydrological and sediment consequences, and their operating limits. The evidence supports a bounded engineering release, not a routing rewrite, complete conservation closure or demonstrated agreement with field observations.

AI authoring disclosure: This report was prepared with AI authoring assistance from OpenAI Codex, including source and evidence analysis, prose drafting, figure preparation, original explanatory code and document checks. AI assistance does not constitute independent scientific review or publication approval; the review status is recorded separately.

The binary release was authorized by Roger Lew and vendored on October 10, 2026. The preceding 261009 build was never deployed and was withdrawn from WEPPpy. Roger Lew approved this document for publication in WEPPpy on October 10, 2026. This approval is distinct from binary release or deployment. NSERL has not agreed to public release of the WEPP source; this report contains detailed technical descriptions and newly authored explanatory stubs, not legacy source listings. No NSERL endorsement is implied.

Contents

1 Release Scope and Main Conclusions **3**

2 Engineering Strategy and Scope **4**

3 Scientific Context and Terminology **5**

4 Hourly Surface-Return Peak Estimation **6**

5 Why Erosion and Sediment Can Change **8**

6 PASS v3: An Explicit Component Contract **8**

7 Bounded Return-Source Construction **9**

8 Channel State Across Daily Calls **11**

9 CHRQIN Rainfall-Source Normalization **11**

10 Native Composition and Workflow Compatibility **13**

11 Alternatives Not Shipped **13**

12 Evaluation Design and Interpretation **13**

13 Mutation Study: Figures 1–3 **15**

14 Topanga: Target Correction and Its Boundary **19**

15 Independent Watershed and Ordinary Events **19**

16 Cedar Creek: Complete Yield and Discharge Comparison **20**

17 Disturbed Rankings and Parameterization Context **29**

18 Remaining Limits and Practical Judgment **30**

19 Release Artifacts and Operational Use **31**

20 Quality Gates and Reproducibility **31**

A Symbols, Units and Interface Semantics **32**

B Descriptive Fit Metrics **33**

C Implementation Coverage **33**

D Release Identity and Evidence Index **34**

1 Release Scope and Main Conclusions

The aim is to correct an anomalous surface-return peak calculation without materially degrading established streamflow, timing, event behavior, stability or workflows relative to `wepp_260803`. Cedar Creek comparisons retain 10 cm initial random roughness. Other comparisons retain their own frozen project inputs; the 10 cm setting is not imposed on every watershed or burn class. Exact parity is neither required nor desirable for erroneous peaks.

Principal findings. Hillslope return peaks now use WEPP's calculated hourly return maximum, added to an independently recomputed rainfall-only peak. Erosion receives these physical inputs, not merely corrected reports. The channel interface separates the components and uses a deliberately simple return-source approximation. Channel state is retained across daily calls. Finally, source normalization counts the positive-time samples actually delivered, correcting the Topanga January 18, 1993 outlet peak from 707.63910 to 12.50683 $m^3\,s^{-1}$. Ordinary-event consequences, including large individual sediment changes, remain important findings rather than being erased by the release decision.[6, 11]

The complete Cedar record preserves hillslope-derived yield: the candidate and 261009 both give 1,267,735,665.682 m^3, only 0.393 m^3 above 260803. The final sampled outlet integral is 0.001008% above 260803 but 0.049858% below 261009. The remaining integral–ledger discrepancy is close to the 260803 discrepancy, not zero. This distinction is central: unchanged water production does not establish unchanged routed representation or accounting closure.[12]

1.1 Three Different Baselines

Baseline	Meaning in this report
Source-change baseline	260803 source `f24c957e3` to 261010 source `7471bb5e9`. The net implementation change covers 14 non-binary files under `src/`; two tracked executables are build products.
Immediate predecessor	261009 source `5a01758b7`. The sole subsequent model-code change is the CHRQIN normalization correction. Earlier scientific changes must not be attributed to that final edit.
Behavioral reference	260803 with matched physical inputs, with 261009 retained as a diagnostic comparator. Comparisons are model-to-model unless explicitly stated otherwise; no new field calibration is claimed.
Previous long-form report	This is the first report in this document series. Earlier Markdown release summaries and work-package studies provide evidence but are not silently relabeled as a prior public edition.

The old 261007 SURPEAK/channel-patch bundle is not layered into 261010. The shipped code descends from the accepted hourly MIXPEAK/component-v3 candidate. Daily-mean physical-peak substitutions, endpoint-based support rules and broader routing reconstructions were not retained. A concise disposition of these alternatives appears in Section 11.

1.2 Reader-Facing Change Map

ID	Shipped change or coordinated change	Account
RN-01	Hourly mixed-flow physical peak; isolate native rainfall calculation	Section 4
RN-02	Event-lifetime producer state and diagnostic hooks	Sections 4.6, 5
RN-03	In-band PASS v3 records, header/version handling and daily clearing	Section 6
RN-04	Caller-volume partition and bounded return-source support	Section 7
RN-05	Channel-profile retention and grid transfer between daily calls	Section 8
RN-06	Positive-time rainfall-source normalization	Section 9
RN-07	Native reader and Roads/AgFields component-preserving composition	Section 10
RN-08	Build linkage, sidecars, paired binaries and withdrawn-version handling	Section 19
RN-09	Committed regressions, discovery/resource policy, evidence packaging	Section 20
RN-10	Concurrent WEPPpy disturbed roughness revision and published benchmarks; not a binary parameter change	Section 17

The internal change register maps every changed implementation file to these IDs and accounts for documentation, archived experiments and build products separately. This report does not reproduce restricted source to establish coverage. Routine names, equations and state descriptions provide the technical account. New code listings are original explanatory Python written for this report, with their relationship to production explicitly limited.

1.3 What Does Not Change

The water-balance scheduling, rainfall-routing equations, soils, roughness inputs, channel routing recurrence and coefficient-selection algorithms are not rewritten. No new model parameter, rainfall-fraction threshold, runtime flag, companion transport file or cross-day hillslope-source reservoir is introduced. The numerical banner remains 2020.500; dated executable names and sidecars distinguish releases. Capacities and general project defaults are unchanged. These statements concern the binary change; the independent WEPPpy roughness update is described separately.

2 Engineering Strategy and Scope

The simplest scientifically defensible correction, with bounded regression risk. WEPP couples runoff generation, erosion, sediment transport and channel routing through legacy numerical methods and shared interfaces. A locally reasonable change can therefore alter processes well beyond its intended target. Attempting to resolve every related inconsistency within one release would enlarge the change surface and make its consequences harder to diagnose and validate.

This release favors targeted corrections and explicit approximations over comprehensive model reconstruction. The objective is to correct demonstrated defects while preserving useful baseline behavior,

not to achieve exact conservation, universal numerical convergence or a fully resolved hydrograph where the model does not calculate one. The baseline is a behavioral reference, not an authority that requires preservation of erroneous peaks.

Simplicity constrains implementation complexity; it does not justify discarding water, concealing accounting discrepancies or retaining a serious regression. Acceptance considers the magnitude, frequency, affected conditions and practical consequences of changes in runoff, peaks, timing, erosion and sediment delivery. Aggregate agreement alone is insufficient when individual events show consequential departures. Model-to-model agreement also does not establish agreement with field observations.

The implemented choices distinguish calculated quantities, approximations and validation evidence:

- **Calculated quantities.** MIXPEAK uses existing hourly return depths and an independently recomputed rainfall-only peak. It retains the water-balance runoff quantity; it does not invent a time-resolved mixed hydrograph.
- **Explicit approximations.** Adding component peaks assumes their coincidence. Channel return delivery uses bounded, constant-rate support, not the literal timing of the hourly return buckets. PASS v3 carries the component information needed for that separation, while the channel routing recurrence is retained. Channel-profile continuation and the normalization repair address identified defects without reconstructing the routing system.
- **Observed evidence.** The comparisons below test water yield, routed discharge, timing and sediment consequences separately. Close Cedar record totals coexist with event-level departures and a nonzero integral–ledger discrepancy; neither is hidden by aggregate fit statistics.

Remaining limitations are documented separately from the defects addressed. They are not automatically prerequisites for this release, nor are they declared harmless merely because a broader repair would be difficult. Broader changes require their own scientific justification and validation. This strategy supports a bounded engineering correction, not a claim of complete physical or numerical closure.

3 Scientific Context and Terminology

WEPP's published surface-hydrology documentation distinguishes rainfall excess, runoff volume, runoff duration and peak discharge; hydrological outputs supply erosion calculations. Its peak methods include a kinematic-wave solution and an approximation to that solution.[1] The channel-routing literature for WEPP describes numerical kinematic-wave and Muskingum–Cunge alternatives and the importance of discretization.[3] Those references establish context, not validation of the new return rule. The implementation-specific account below is based on the shipped-source audit and retained experiments.[6]

An *overland flow element* (OFE) is an element of a hillslope profile. Surface return here is water delivered back to the surface by the modeled soil-water calculation. It is not synonymous with rainfall excess and is not a new water source added by the estimator. The hourly return buckets are calculated model quantities, not measured within-day hydrographs. The word "event" may mean a model record/date or a selected multi-day analysis window; neither automatically identifies a complete isolated storm.

Use R for final runoff depth, S for surface-return depth, and A for the applicable area. Depths are in metres; p denotes a depth-rate peak in $\mathrm{m\,s^{-1}}$. Volumes are V in $\mathrm{m^3}$ and volume-rate peaks are Q in $\mathrm{m^3\,s^{-1}}$. Multiplication by A converts matching-area depth quantities to volume quantities. A rate in $\mathrm{m\,s^{-1}}$ is multiplied by 3.6×10^6 to report $\mathrm{mm\,h^{-1}}$. Definitions and interface units are collected in Appendix A.

Three distinct durations: $D_{\text{eff}} = R/p$ is an equivalent rectangular duration used by hillslope erosion bookkeeping; T_r is the existing rainfall/storm support supplied to channel reconstruction; and T_s is the local support chosen for returned water. None by itself establishes the actual timing of a mixed hydrograph. The time-of-concentration scalar `tcs` is a further, separate quantity.

4 Hourly Surface-Return Peak Estimation

4.1 Legacy Concentration Mechanism

The legacy IRS path computes surface surplus after water balance, maps it to the runoff-depth basis, and can augment rainfall-excess ordinates by an amount proportional to returned depth divided by rainfall-excess duration. It then uses the existing peak procedure. A sizeable returned depth can consequently be represented over a short rainfall-excess period. This is a mismatch of source volume and representative duration, not proof of an unusually intense rainstorm. The final peak affects physical calculations downstream.

In the two reproduced H106 Ksat-35 cases, approximately 33.690 and 40.075 mm of runoff coexist with legacy peaks of 240.274 and 274.596 $\text{mm}\,\text{h}^{-1}$. The revised peaks are approximately 1.605 and 2.038 $\text{mm}\,\text{h}^{-1}$ with runoff unchanged. These are targeted mutation reproductions, not measured Cedar flood peaks. Their small but positive rainfall contribution matters: a pure-return-only correction would miss the original triggers.[7]

4.2 Separation After Water Balance

The producer operates only on single-OFE, non-contoured, continuous-hourly cases with positive surface return. It takes final runoff depth R from `runtmp` and returned-water depth S from `surpls`, on the same area basis. In the eligible single-OFE setting the general equivalent-plane scaling does not introduce an unresolved multi-OFE area mapping. Define

$$R_r = R - S, \qquad V_r = AR_r, \qquad V_s = AS. \tag{1}$$

Finite R and S, positive S, and $R \geq S$ are checked. Hourly return depths s_h must be finite and nonnegative with a positive sum. Invalid separation is rejected; a small positive rainfall contribution is not rounded to zero or removed using a return-dominance threshold.

4.3 Independent Rainfall-Only Peak

MIXPEAK does not subtract a return rate from an already contaminated combined peak. For $R_r > 0$, it calls the existing applicable rainfall-only procedure with R_r and the native rainfall/routing operands. The approximation branch uses APPMTH; the alternate branch uses HDRIVE. Exactly zero R_r produces zero physical rainfall peak. The water-balance calculation is not repeated.

HDRIVE and its callees use shared integration and random-number scratch. The helper saves and restores the integration arrays, counts/time step, coefficient scratch and random-state values around that call. Otherwise an apparently local peak correction could alter later calculations by advancing shared state. Branch selection follows the existing approximation and rainfall-input conditions rather than a new fitted criterion. The legacy route can still execute, but its final combined peak is replaced when this eligible component calculation is active.[6]

4.4 Combination, Duration and Bounds

Let p_r be the independently calculated rainfall peak and s_h the existing hourly returned depth in metres. The pre-floor mixed-flow estimator is

$$p_s = \max_{h=1,\ldots,24} \frac{s_h}{3600}, \qquad p = p_r + p_s. \tag{2}$$

Rainfall and hourly return peaks are assumed to coincide. Return receives no additional hillslope attenuation in this estimator. The sum is an upper envelope of these two specified component peaks under that coincidence assumption, not a guarantee about unresolved subhourly return peaks or true field discharge. Low rainfall volume alone does not establish a low rainfall peak; its calculation is retained.

The existing physical floor sets a peak below 3.6×10^{-8} to 3.63×10^{-8} $\mathrm{m\,s^{-1}}$. IRS then obtains $D_{\mathrm{eff}} = R/p$ and retains the existing 86400 s cap. Later peak bookkeeping recomputes a rate from runoff and effective duration with the existing OFE-area mapping. Thus the final published peak need not equal the unbounded sum if a legacy floor/cap engages. A duration cap can raise the reassigned rate to the equivalent one-day mean; it must not be described as discarded volume or as an unconstrained new physical estimate.

The rainfall operand sent to the channel is separately protected by the existing floor/cap convention. For positive R_r, let p_r^* denote its floor-adjusted value, and define

$$D_r^* = \min(86400, R_r/p_r^*), \qquad Q_r^I = AR_r/D_r^*. \tag{3}$$

For exactly zero rainfall, $Q_r^I = 0$. This interface operand is not a manufactured physical rainfall component: p_r in Equation 2 remains zero for pure return. Nor is D_r^* the storm-support field T_r.

4.5 Eligibility and Continuity

No-return cases follow the baseline producer path. Multi-OFE, contoured and non-hourly cases are outside the targeted producer correction, not claimed fixed by inference. With positive return, every positive R_r is retained. As its independently calculated peak tends to zero, Equation 2 tends to p_s; this algebra does not guarantee that all native branches or legacy floors are continuous. Native tests include vanishing rainfall and return, pure return, no-return controls and support endpoints. A proof of continuity of the entire legacy model is not an acceptance prerequisite.

4.6 Event Lifetime and Observability

The new producer common block holds an active flag, rainfall/return depths, physical and interface peaks and 24 hourly values for the current event. Initialization and the continuous day loop reset this state. Consumer-side arrays are cleared before assembling or reading each day's sources. This prevents an active source packet from a previous event being reused when a later event has no eligible return.

Existing opt-in observation infrastructure receives mixed depths, component peaks, final duration/peak, hydraulic quantities and derived concentration time/alpha diagnostics. The release does not add a new user flag. Observer neutrality is tested against standard outputs; the act of logging is not assumed harmless without comparison.

5 Why Erosion and Sediment Can Change

WEPP's erosion calculation is coupled to hydrological and hydraulic inputs, not just to annual runoff volume.[2] In the inspected release, XINFLO uses the physical peak in runoff and shear-related hydraulic quantities; PARAM applies duration-dependent factors; SLOSS uses effective duration when integrating loads and delivery. Correcting the peak can therefore change hydraulic radius, width, effective friction and erosion even while runoff volume remains the same. These are model consequences, not a reporting patch.

Neither direction of sediment change proves accuracy by itself. Lower artificial peaks can reasonably reduce detachment/transport inputs, but changes to duration, thresholds, deposition and channel history need not produce a uniform decrease in exported sediment. The report distinguishes gross erosion or detachment, hillslope delivery and outlet export throughout.

For physical context, the published hillslope formulation makes rill detachment capacity depend on excess soil shear stress:

$$D_{\mathrm{cap}} = K_r(\tau_f - \tau_c), \qquad D_f = D_{\mathrm{cap}}(1 - G/T_c), \tag{4}$$

in the detachment domain $\tau_f > \tau_c$ and $G < T_c$. Here K_r is rill erodibility, τ_f and τ_c are applied and critical stress, and G and T_c are load and transport capacity on the same unit-width basis. Deposition is a different regime. Peak-driven hydraulics can change the threshold regime and transport capacity, while effective duration changes the time over which rates contribute to delivery. These published relationships explain why volume parity alone cannot establish sediment parity; they are background science, not new equations introduced by this release.[2]

The legacy `tcs` formula and overland alpha calculation are retained. Their values can nevertheless respond indirectly to changed hydraulic inputs, including effective Manning behavior. The PASS alpha uses concentration time in seconds, physical peak and area divided by runoff volume, with its existing minimum bound. Saying "no tcs formula change" is not a claim that every tcs or alpha value is identical. Importantly, the return-only support introduced below is local to source assembly: it is not written into erosion duration, storm duration, tcs or alpha. In CHRQIN's exponential construction the internal shape time is $T_r/2.67$, not an imposed new tcs smoothing time.

Soil-water scheduling, snow processes and nutrient equations are not directly rewritten by these changes. Their presence in the executable does not constitute independent validation of every coupled output; nutrient consequences were not a separate acceptance study for this release.

6 PASS v3: An Explicit Component Contract

One combined peak, duration and volume cannot in general distinguish a short rainfall event from soil water returned over a longer period. Replacing only that combined peak can make the existing reconstruction inconsistent with the water it is expected to deliver. The release therefore uses in-band component information in the existing PASS transport, not a companion file or a guessed decomposition in the watershed consumer.[7, 10]

The leading marker is `WEPP_PASS_COMPONENTS 3`. Extended hillslope EVENT records have an accompanying `SRC3` record. Its date and active flag identify either an inactive packet or an active 27-value payload:

Payload	Meaning	Units/interpretation
V_r	Non-return runoff volume	m^3
Q_r^I	Protected rainfall-only peak operand	$m^3\,s^{-1}$
V_s	Returned-water volume	m^3
$w_1, \ldots, w_{24}$	Nonnegative hourly return weights	Relative weights

The producer initially writes hourly depths as weights. Native composition can instead produce volume weights. The v3 consumer uses ratios of these weights, so their common scale is immaterial. Their hour positions are retained as provenance but *do not prescribe v3 channel forcing times.* This is an intentional versioned interpretation, not a hidden serialization change.

Master files include a dated `SOURCES3` count and ordered active source records keyed by hillslope index. Initialization and reading are wired through INFILE, WSHPAS and WSHRED. Header handling accepts legacy input contracts for compatibility, accepts v3, and rejects unsupported component versions rather than guessing offsets. Reading legacy syntax is not authorization to mix builds in a new watershed simulation.

Parsers check record arity, dates, counts, source order/bounds and numeric validity. Active values must be finite and nonnegative with a nonempty return component; zero rainfall volume cannot carry a nonzero rainfall peak. The record tokenizer does not accept list-directed null/repeat/slash conventions as a fallback. The existing EVENT-selection logic, including its runoff threshold, is not replaced with a new universal event definition.

7 Bounded Return-Source Construction

7.1 Volume Allocation Comes First

The watershed caller may allocate only a portion V_c of the hillslope's surface volume to a particular top or lateral source. For an active packet, CHRQIN partitions that caller volume using the serialized proportions:

$$V_{r,c} = V_c \frac{V_r}{V_r + V_s}, \qquad V_{s,c} = V_c \frac{V_s}{V_r + V_s}. \tag{5}$$

This preserves the caller's total; it does not add another copy of hillslope water. Existing caller allocation semantics are retained. The component peak operand is not silently rescaled into a new invented rainfall hydrograph.

7.2 An Hourly Ceiling, Not Literal Hourly Timing

Let $W = \sum_h w_h > 0$, $w_{\max} = \max_h w_h > 0$, and let T_r be the existing storm support in seconds. Define

$$T_{\mathrm{eq}} = 3600 \frac{W}{w_{\max}}, \qquad T_s = \max(T_r, T_{\mathrm{eq}}), \qquad q_s = \frac{V_{s,c}}{T_s}. \tag{6}$$

The return source is a rectangle starting at the existing computational event origin and ending at T_s. Since $T_s \geq T_{\mathrm{eq}}$,

$$q_s \leq \frac{V_{s,c} w_{\max}}{3600 W}. \tag{7}$$

This ceiling refers to the allocated *return component*, not total rainfall-plus-return inflow or outlet discharge. The rate need not attain the hourly maximum. Return timing is redistributed onto a

scalar support anchored to an existing computational clock; it is not recovered from an observed mixed hydrograph or from the positions of the original hourly buckets.

For routing interval $[(i-1)\Delta t, i\Delta t]$, assembly adds its overlap mean:

$$\bar{q}_{s,i} = q_s \frac{\max\{0, \min(i\Delta t, T_s) - (i-1)\Delta t\}}{\Delta t}. \tag{8}$$

On a grid covering the support, these interval means represent exactly the allocated rectangular-source volume before downstream processing, apart from floating-point arithmetic. Partial final cells are not silently truncated. At initial setup the zero-time boundary receives the initial rate; subsequent days retain the caller's existing aggregate boundary convention.

The 24 nonnegative weights imply $3600 \leq T_{\text{eq}} \leq 86400$ seconds. The implementation rejects invalid operands or support beyond the available day rather than clipping water or introducing a new cross-day source store. Existing storm support and the full source grid still need to satisfy their contracts. Volume consistency of this assembly alone does not guarantee volume consistency through every downstream operator.

7.3 Why the Hourly Bins Are Not Used Literally

The investigation found that daily scheduling can concentrate modeled return in the last hourly bucket. The hour receives consequences of daily bookkeeping, including late ET/residue and redistribution operations; it is not simply an arbitrary "whatever remains" bucket with independently known physical timing. In a read-only Cedar population of 61,645 eligible return records, 4,450 had at least 90% of their return in hour 24, but those records represented only 0.67% of eligible return volume. Thus a small aggregate volume share could still create conspicuous event pulses. These are descriptive population bins, not eligibility thresholds.[9]

Keeping literal hourly timing produced problematic next-day pulses in ordinary controls. A global routing-mode switch, moving ET, repeating percolation, or attributing a residue release directly to returned-water parcels would change other assumptions without a demonstrated bounded benefit. The selected rule preserves the hourly physical estimator while permitting a coarser channel representation. It is an explicit simplicity-versus-timing-fidelity decision.

7.4 Original Explanatory Stub

The following is original Python authored by Codex for this report, not a legacy listing or a production substitute. It illustrates Equations 6–8. It omits legacy boundary handling, rainfall reconstruction, caller allocation, serialization precision and the routing solver. Inputs are already valid, positive-return operands. The companion examples are checked independently of any watershed claim.

```
def return_rectangle(volume_m3, weights, storm_s, step_s, steps):
    """Illustrate the v3 support rule and interval-overlap means."""
    if len(weights) != 24 or min(weights) < 0 or max(weights) <= 0:
        raise ValueError("Expected 24 nonnegative, nonempty weights")
    if volume_m3 <= 0 or not 0 <= storm_s <= 86400 or step_s <= 0:
        raise ValueError("Expected positive volume and a valid day/grid")
    support = max(storm_s, 3600 * sum(weights) / max(weights))
    if support > 86400 or steps * step_s < support:
        raise ValueError("Support must fit the represented day/grid")
    rate = volume_m3 / support
    means = []
    for i in range(1, steps + 1):
        overlap = max(0, min(i * step_s, support) - (i - 1) * step_s)
        means.append(rate * overlap / step_s)
    return support, means
```

Listing 1: Original return-rectangle illustration; not WEPP production code.

8 Channel State Across Daily Calls

The old daily initialization retained boundary information but reconstructed internal channel discharge by drawing a straight line between inlet and outlet. Different internal profiles can have the same endpoints. Replacing the profile can therefore alter subsequent routed discharge even when the next day's forcing is unchanged. This is distinct from source timing and source-volume normalization.[8]

The new CHRSTATE helper stores each channel's final nodal discharge profile and segment count after routing. On the next call, an equal grid restores nodes exactly. If the uniform grid changes, interior nodes are linearly interpolated in normalized reach distance. For old segment count n_o and new node j on n_n segments, let $x = jn_o/n_n$, $k = \lfloor x \rfloor$ and $\lambda = x - k$; the transfer uses

$$q_j^{\text{new}} = (1 - \lambda)q_k^{\text{old}} + \lambda q_{k+1}^{\text{old}}. \tag{9}$$

The downstream endpoint is preserved. The caller's current upstream boundary remains authoritative, and an initialization call resets saved state. This bounded interpolation does not create an extremum between the bracketing old values, but it is not a storage-conservative remap theorem.

In native 48-hour diagnostics, daily restart errors for a 3600 m^3 last-hour source reached roughly -1121 m^3 on a lateral case and +2414 m^3 on a top case, depending on routing method. Same-grid retention reproduced uninterrupted native trajectories in the tested controls. These are synthetic diagnostic results, not observed watershed errors or proof of universal convergence. The repair preserves channel routing memory, not a new hillslope reservoir.

9 CHRQIN Rainfall-Source Normalization

9.1 The Existing Reconstruction Contract

CHRQIN supplies top or lateral source rates from volume, peak and support. For rainfall volume exceeding its existing small absolute threshold, it uses the provided support T_r and peak operand. In the active component case, these are the rainfall volume allocation and Q_r^I, not the combined physical MIXPEAK peak. One legacy branch is flat when volume is at least approximately peak times support. The other constructs a double exponential.

For $a = V_{r,c}/(Q_r^I T_r)$, the routine solves

$$1 - \exp(-u) = au, \qquad t_c = T_r/2.67, \tag{10}$$

and defines rates proportional to

$$r(t) = \begin{cases} Q_r^I \exp[(u/t_c)(t - t_c)], & t \le t_c, \\ Q_r^I \exp[-u(t - t_c)/(T_r - t_c)], & t_c < t \le T_r. \end{cases} \tag{11}$$

The configured pointwise path samples this shape and scales it to its supplied volume. A separate existing cell-average path is not replaced by this release. The delivered sampled maximum need not equal Q_r^I after normalization; continuous shape identities do not automatically survive finite sampling.

9.2 The Indexing Error and Its Correction

Write $S_q = \sum_{i=1}^{n} r_i$ for the positive-time samples that contribute to the chosen integral, at interval Δt. On the initialized path the old normalizer removed the first positive-time sample even though the addition loop still delivered it. When the remaining sum was positive,

$$B_{\text{old}} = \Delta t(S_q - r_1), \quad V_{\text{represented,old}} = V_{r,c}\frac{S_q}{S_q - r_1}. \tag{12}$$

If nearly all sampled shape lies in r_1, the factor can become enormous. If the remaining denominator vanishes, the scale-application guard can instead suppress the source. Floating-point cancellation can aggravate the symptom, but increasing precision alone does not repair inconsistent sample accounting.

The correction retains every positive-time sample in the denominator. A time-zero sample is subtracted only on the path where it was accumulated:

$$B_{\text{new}} = \Delta t S_q, \qquad q_i = V_{r,c}\, r_i / B_{\text{new}}. \tag{13}$$

This is not a global change to the initialization flag and not a peak cap. It applies to ordinary rainfall-source reconstructions using this branch as well as active component sources. The flat branch, existing absolute volume threshold and short-support limitations remain. No universal closure assertion follows from the local identity.[11]

9.3 Captured Topanga Operands

For H101 on January 18, 1993, the allocated rainfall component was only 0.090845406 m^3 within 274.970 m^3 total surface volume. Its peak operand was 0.0014701501 $\text{m}^3\,\text{s}^{-1}$, support 1440 s and interval 600 s. The unscaled first sample was approximately 3.059043×10^{-4} $\text{m}^3\,\text{s}^{-1}$, while the next was only 5.54092×10^{-11} $\text{m}^3\,\text{s}^{-1}$. The faulty single-precision denominator was approximately 3.49246×10^{-8} m^3, producing a multiplier of about 2.60 million. Native routine probes reproduced source pulses near 796 and 185 $\text{m}^3\,\text{s}^{-1}$ from H101 and H234, before channel routing.

The earlier combined-volume operands took a different branch, so a legacy defect became conspicuous only after component separation. Read-only observer instrumentation reproduced 293 output files exactly over 45 years. The pulses were already present at headwater sources; channel-state carry was not their origin. Correctness therefore required repairing this arithmetic boundary, not zeroing small rainfall, discarding the hourly estimator or reconstructing the routing architecture.

10 Native Composition and Workflow Compatibility

The paired native reader/combiners support legacy syntax and component v3, reject experimental/unknown component versions, and preserve ordinary public Parquet columns. Roads metadata operations preserve the leading version marker and climate identity. AgFields reads modeled area and climate after the marker rather than using stale fixed offsets. No Python physical-composition fallback silently drops the new operands.[10]

Roads applies additive scaling; AgFields applies existing represented-area weights. Active V_r and V_s are scaled and summed. Rainfall peaks use the existing triangular peak-combination procedure with source concentration times. Inactive sources contribute legacy volume and peak as an unsplit component; the combiner does not infer unsupported return fractions. Each active return weight vector is normalized, multiplied by represented V_s, and accumulated by hour. Zero-area sources contribute nothing; zero aggregate return produces an inactive record.

Existing combined EVENT volume/peak/duration, tcs, alpha and sediment rules remain in place. The channel consumes separately combined source operands; composition does not rerun hillslope erosion. This aggregate construction is not exact superposition of separately routed source hydrographs. Its support and timing assumptions remain the bounded v3 approximation.

11 Alternatives Not Shipped

Alternative	Disposition and reason
Pure-return-only eligibility	Demonstrated a subset, but missed original return-dominated triggers with small positive rainfall. No return-fraction threshold was introduced to disguise that boundary.
Daily-mean physical return peak	Not retained. It discards the hourly estimator used for physical erosion; favorable routing alone does not justify it.
Support from the last return endpoint	Rejected: a tiny late contribution could impose an hours-long support/timing change. The 43.484 mm rainfall control with 0.01 mm late return remains a relevant boundary test.
Literal hourly channel forcing alone	Produced problematic late pulses and recession enlargement. State continuation fixed a separate restart error but did not by itself remove those source-timing effects.
New source carry, water-balance rescheduling or general routing rewrite	Not part of the bounded release. No unvalidated soil-water change or new transport architecture is used to obtain favorable totals.
Old 261007 repair bundle	Superseded, not silently accumulated. The net shipped-source audit identifies the implementation actually released.

12 Evaluation Design and Interpretation

The governing standard was a demonstrated correction of the original peak anomaly without material degradation relative to 260803 under matched inputs. No new calibration, roughness workaround, universal convergence requirement or strict zero-residual tolerance was introduced. Materiality was assessed through magnitude, frequency, affected conditions and practical consequences; ambiguous scientific decisions remained with Roger Lew.

The evidence combines native-routine probes, generated full-period watersheds, same-build hillslope/-PASS regeneration, observer-neutrality checks, small-input mutation experiments and a canonical disturbed matrix. These are different forms of evidence. An isolated arithmetic identity is not watershed validation; high model-to-model agreement is not observed-flow skill. The release binaries match the studied candidate after removing debug data and GNU build identity, with executable code and runtime data retained in the comparison. Fresh release checks supplement that equivalence.[17]

Study	Extent	Question answered
H106 native triggers	Two Ksat-35 cases; Ksat/no-return controls	Physical peak correction at the producer, not only a lower outlet peak.
Topanga	Burned/undisturbed; 284 hills each; 1980–2024	Localization and correction of the January 1993 source inflation, with full-record consequences.
Rattlesnake	Burned/undisturbed; 73 hills each; 100 synthetic years	Independent watershed compatibility and ordinary-event sensitivity to normalization.
Cedar Creek	864 hills, 371 channels; 1980–2003	Yield, routed-discharge sampling, ledger discrepancies, events, timing and sediment under matched inputs.
Mutation census	140 hills in each of two scenarios; 1980–2024	Sensitivity to small conductivity/cover perturbations within the corrected model.
Disturbed matrix	96 cases; four soils, four severities, six vegetation types; 100 synthetic years	Parameterization expectations on one common slope/climate, including young forest, not universal severity ranking.

Use fresh matching-build passes in simulations; reader backward compatibility does not justify substituting old PASS files. Diagnostic observer output was compared with uninstrumented standard output before interpretation. Retained failed branches and incomplete attempts are distinguished from passing results.

12.1 Volume, Peak and Timing Measurements

Daily hillslope streamflow from totalwatsed is converted to volume as depth times contributing area divided by 1000 when depth is in millimetres. The watershed EBE outlet volume is a separate daily ledger. Cedar's sampled integral uses the 144 positive-time discharge samples per day:

$$V_{\mathrm{sampled},d} = 600 \sum_{i=1}^{144} Q_{d,i}, \qquad G_d = V_{\mathrm{sampled},d} - V_{\mathrm{ledger},d}. \tag{14}$$

Samples, including the 86400-second endpoint, remain assigned to the printed simulation day. This reproduces the comparison convention; it is not exact continuous integration. Report both $\sum_d G_d$ and $\sum_d |G_d|$, because opposite daily discrepancies can cancel.

Extra digits in comparison tables preserve recorded output identities; they do not imply sub-cubic-metre physical accuracy over decades. Printed sampling, rounding, discretization and the model's assumptions limit interpretation.

Peak comparisons distinguish internally reported maxima and maxima of printed samples. First-argmax timing can jump between nearly equal crests or rounded plateaus. Flow centroids use all sampled discharge in the selected window; they are not automatically storm travel times. Event windows include background flow unless stated otherwise. NSE, KGE and R^2 are descriptive comparisons of model outputs, not automatic acceptance gates or field validation scores. Definitions are given in Appendix B.[4, 5]

13 Mutation Study: Figures 1–3

The companion investigation by Lew, *Small-Mutation Screening of WEPP Hillslope Peak Flow: Remediation Strategies and Validation Findings*, version 0.2 (October 6, 2026), explains the original Topanga census, perturbation design and earlier remediation assessment.[13] It is historical context: its earlier candidate and unresolved Cedar routing results are not the shipped 261010 implementation or its final validation. The results and Figures 1–3 below describe the retained 261010 candidate study.

All 280 unmutated controls and 1088 eligible mutations completed, plus ten uninstrumented H106 checks of observer neutrality. The design retains first-horizon conductivity factors 0.99/1.01 and paired initial interrill/rill-cover changes of -0.01/+0.01. Thirty-two originally excluded cover mutations remain excluded rather than clipped. All nonmutated inputs, including frozen roughness, stay unchanged.[14]

Both axes in the following figures use the validated candidate: one axis is unmutated input and the other mutated input. They are *not* corrected versus 260803 plots. The report renderer changes layout only; input-ledger hashes, filters, marker meanings and counts are checked against the retained study. The candidate's event and sediment ledgers match 261009 exactly, and the release's executable equivalence establishes applicability to 261010.

The outer ledger has 225,658 hillslope/date/mutation rows: 225,042 paired events, 320 unmutated-only and 296 mutated-only rows. Missing events are not zero-filled. An event can appear for multiple perturbations, so these are not independent storm counts. Sediment is the original EBE hillslope-delivery metric in kg/m, not gross detachment or watershed export.

Metric	Eligible pairs	Beyond 2× guides	Beyond 5× guides	Outside axes
Runoff	218,974	453	34	86
Sediment delivery	129,496	329	72	0
Peak flow	200,984	966	209	179

The twofold guides mean ratios below 0.5 or above 2; fivefold guides mean below 0.2 or above 5. They are inherited descriptive guides, not thresholds invented to accept the release. Existing nonlinear sensitivity remains: 966 eligible peak pairs lie beyond twofold guides, including 543 with surface return and 423 without it. Small input perturbations can change event behavior even in a successfully repaired release.

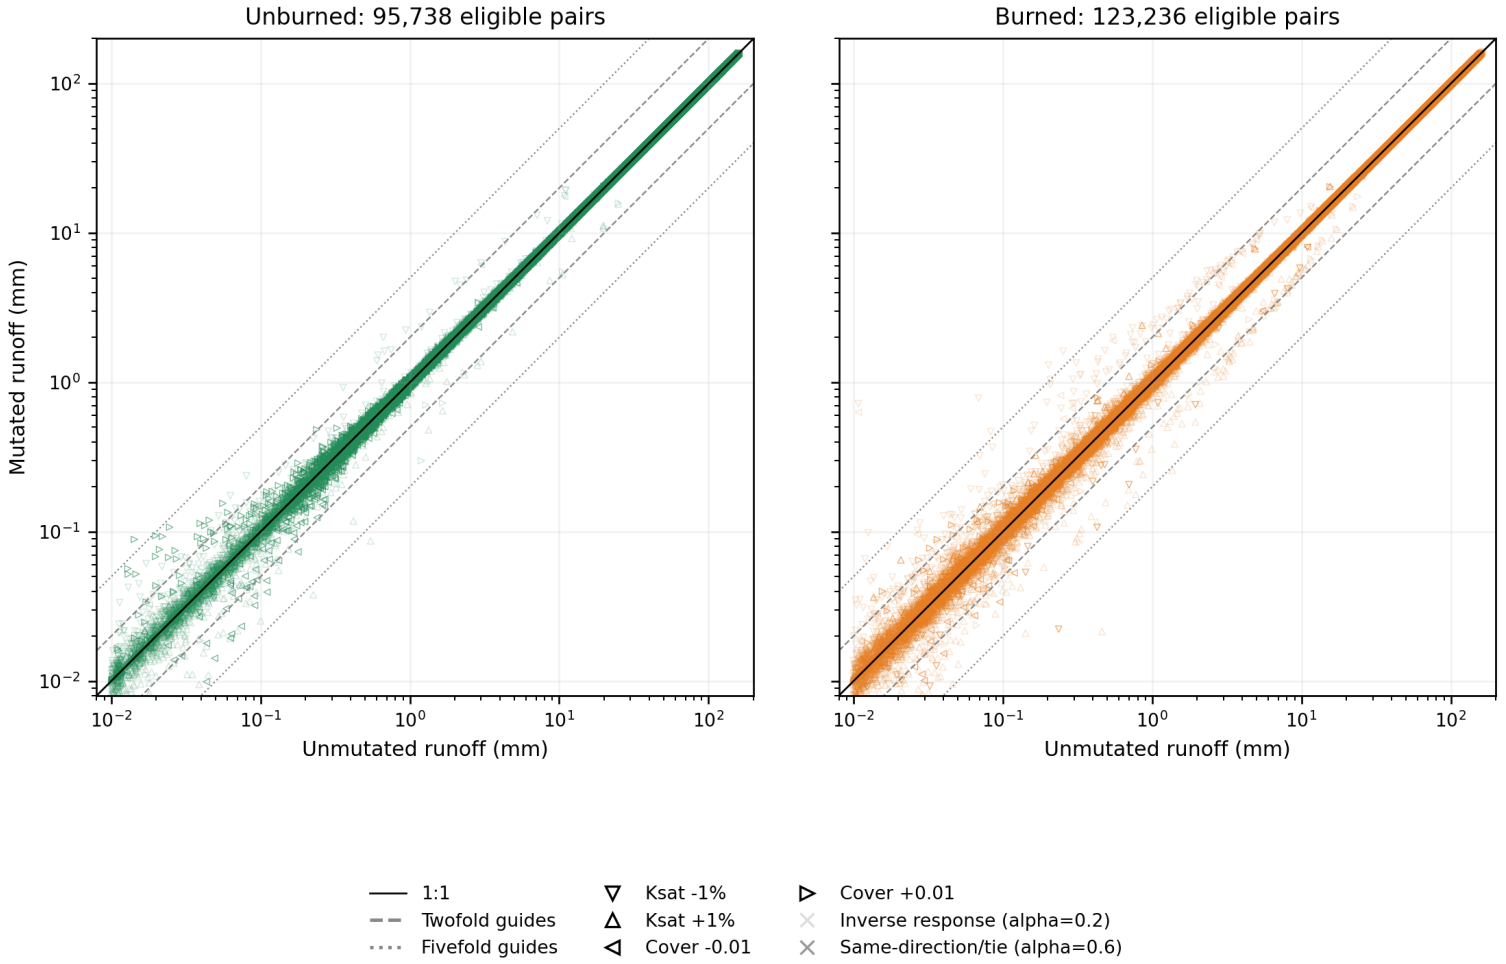


Figure 1: Runoff sensitivity within the corrected candidate, with full 1980–2024 histories. Both events must be present; unmutated runoff is at least 0.01 mm and mutated runoff is positive. Original axes span 0.008–200 mm. The 86 eligible pairs outside the displayed limits remain in the statistics. Triangles identify conductivity/cover directions; opacity distinguishes inverse responses from same-direction responses and ties. Layout-only rerender of retained observations, not new simulations or a cross-version comparison.

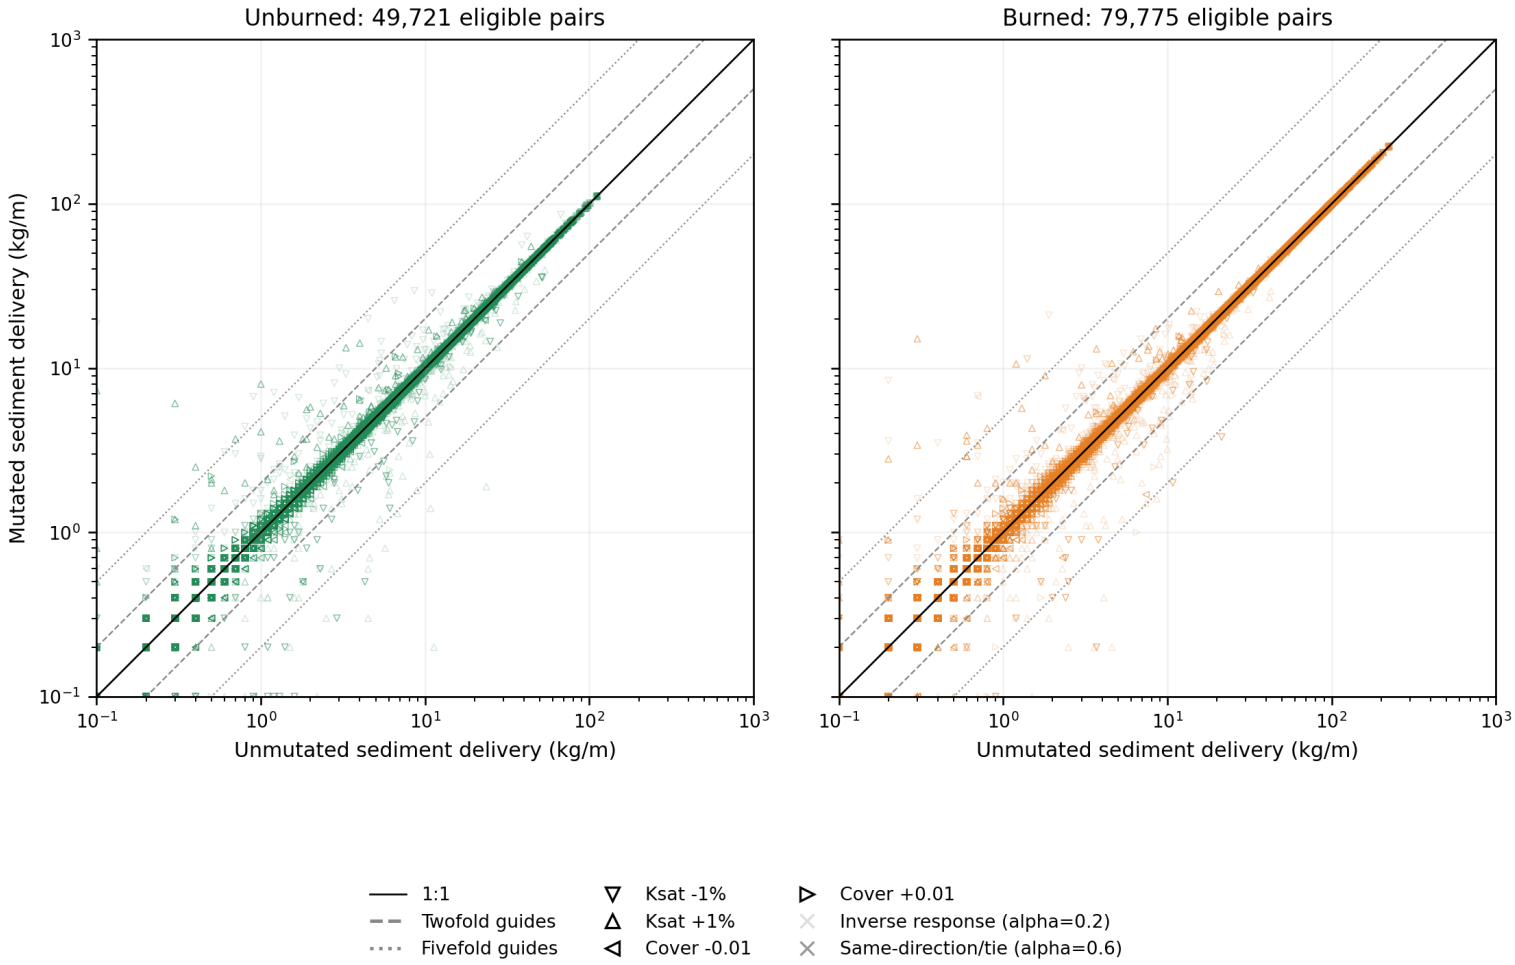


Figure 2: Hillslope sediment-delivery sensitivity in kg/m, the report's erosion metric. Both printed deliveries must be positive. There are also 540 unmutated-positive-only and 583 mutated-positive-only records, retained outside this paired-positive plot. The EBE 0.1 kg/m precision produces visible steps and many ties. This is not a plot of gross detachment or channel sediment export.

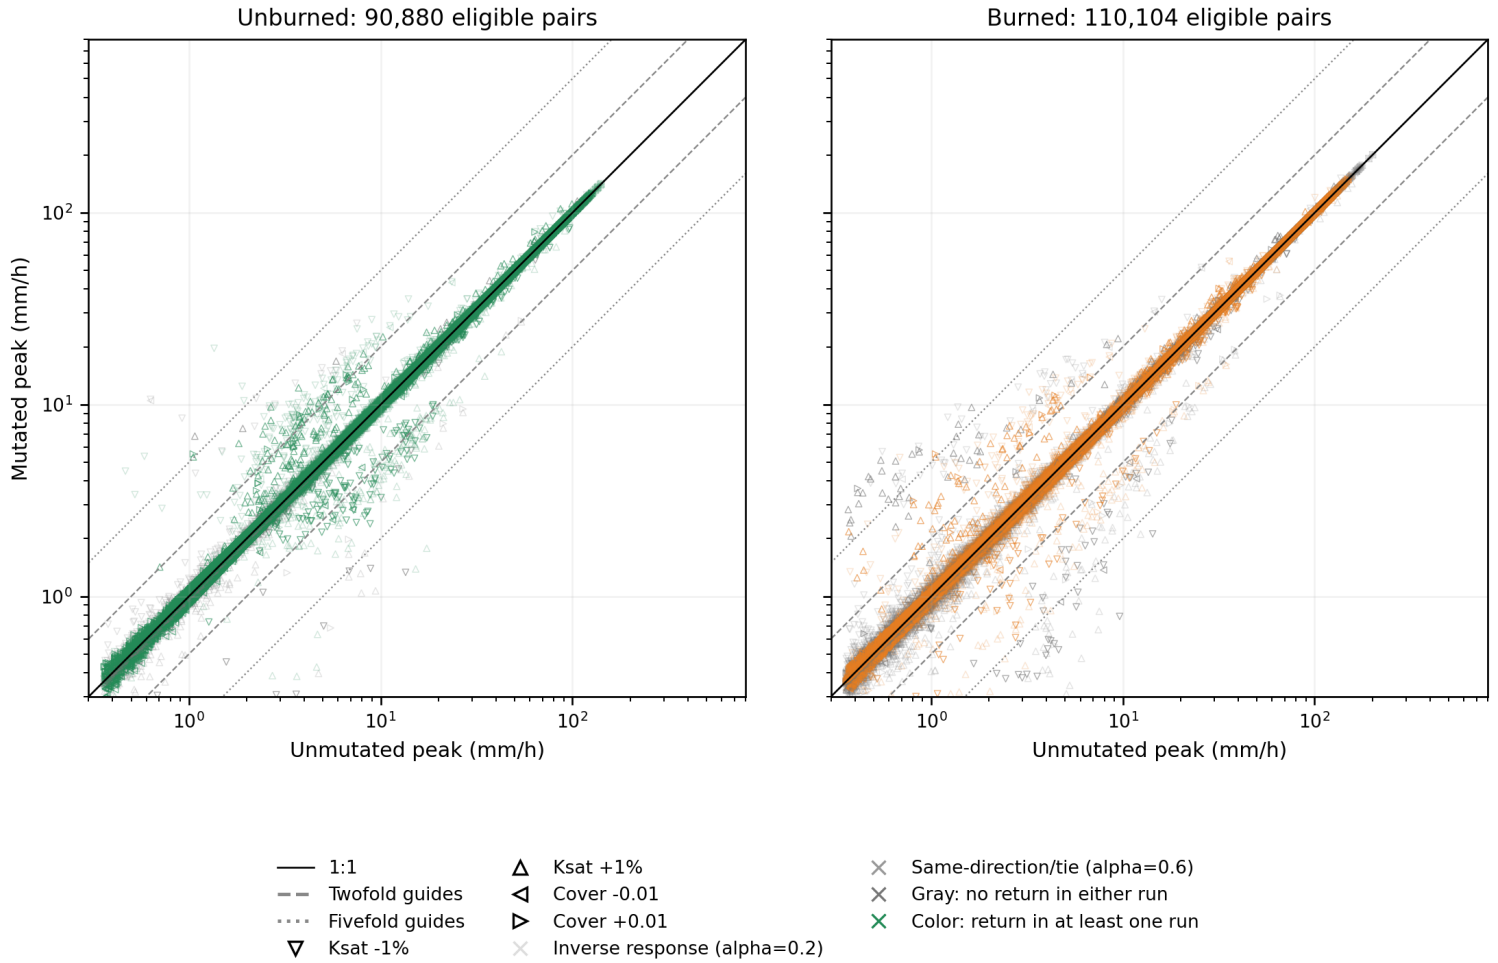


Figure 3: Peak-flow sensitivity within the corrected candidate. Both events must be present; unmutated peak is at least 0.36 mm h^{-1} and mutated peak is positive. Axes span 0.3–800 mm h^{-1}; 179 eligible pairs outside the axes remain counted. Gray denotes no return in either run; color denotes return in at least one. The hourly physical estimator is retained, not replaced by a daily mean. Both axes use the same corrected implementation.

14 Topanga: Target Correction and Its Boundary

The January 18, 1993 GridMET case was initially investigated as an outlier, not presumed defective solely from a burned/unburned reversal. Matching inputs, the 260803 comparison, native source probes and observer-neutral traces then established the normalization mechanism. Roger judged the roughly 708 $m^3\,s^{-1}$ result physically unacceptable at this location. This is a diagnosis supported by implementation evidence, not an assumption that every unusual peak is wrong.

Scenario	Build	Outlet volume (m^3)	Peak ($m^3\,s^{-1}$)	Sediment (kg)
Burned	260803	213215.02	106.42034	212686.55
Burned	261009	213282.16	76.03957	222187.20
Burned	Candidate	213281.84	35.72905	223117.84
Undisturbed	260803	219667.84	27.85154	39192.88
Undisturbed	261009	218608.19	707.63910	26789.54
Undisturbed	Candidate	219692.70	12.50683	27201.02

The corrected undisturbed volume is +0.0113% versus 260803. Its peak occurs at 1800 s instead of 1200 s in 261009. Removing the manufactured source pulse is the evidence of repair; a lower outlet peak by itself would not have been sufficient. Candidate watershed output for this study was peak-only, so it does not provide a new full subdaily recession or integral–ledger audit. The figure below connects daily maxima, not hydrograph samples.[11]

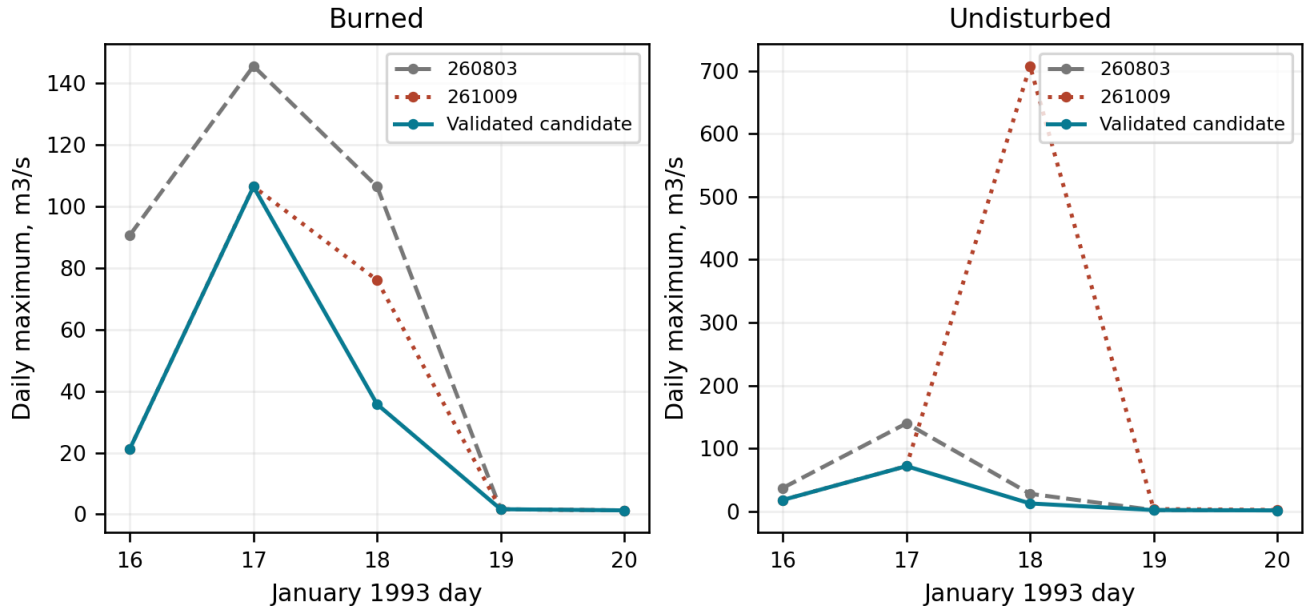


Figure 4: Retained January 1993 daily peak comparison. The candidate label identifies the studied, pre-packaging executable; runtime equivalence to the 261010 release is documented. Lines between daily maxima must not be read as a continuous subdaily hydrograph.

15 Independent Watershed and Ordinary Events

Both Topanga scenarios and both Rattlesnake scenarios completed with 714 fresh hillslope runs in the candidate comparison. Topanga covers 45 years; Rattlesnake covers 100 synthetic years. Signed full-record changes versus 260803 are:

Case	Reported outlet volume change (%)	Outlet sediment change (%)
Topanga burned	-0.0000510	+1.30014
Topanga undisturbed	-0.0000501	-0.61171
Rattlesnake burned	-0.0000002	+0.49250
Rattlesnake undisturbed	+0.0000001	+1.98547

These aggregate values do not establish that individual events are unchanged. For burned Rattlesnake synthetic year 40, May 15, both released comparators have a 44.54780 $m^3 s^{-1}$ peak; the candidate gives 20.21181 $m^3 s^{-1}$ (-54.6%). Volume changes only from 28250.79 to 28250.22 m^3. The twelve EVENT source packets that date are inactive for mixed-return transport. This establishes an ordinary legacy-source-path consequence; an inactive packet is not by itself proof of physically zero return in every unsupported producer mode.

The following sediment examples also have essentially unchanged event volume:

Burned Rattlesnake event	Both older builds (kg)	Candidate (kg)	Difference (kg)
Year 41, May 28	140954.78	621811.88	+480857.10
Year 89, June 6	438143.62	99119.68	-339023.94

All checked hillslope sediment fields match 261009 across the 146 Rattlesnake hillslope runs. The differences arise downstream, but a unique channel transport/storage explanation was not isolated. An adjacent-storm redistribution explanation is not established merely by its plausibility. Year-41 sediment export increases 19.746%; year 89 decreases 13.837% versus 261009, despite the full-record incremental sediment change of only +0.346%. These findings remain disclosed through promotion.[11]

Correcting a known source-accounting defect is justified independently of whether a legacy number is reproduced. It also affects practical peak and sediment interpretation. Roger explicitly authorized the release after the follow-up studies; that decision is not evidence that these individual sediment magnitudes have been confirmed against observations.

16 Cedar Creek: Complete Yield and Discharge Comparison

The complete 1980–2003 record contains 8766 days, 864 hillslopes and 371 channels, with frozen 10 cm initial random roughness. Fresh candidate and 261009 runs regenerated 1728 hillslopes; their 6048 standard hillslope files match exactly. Native conversion accepted 7,573,824 PASS and 16,690,464 WAT records per build with no rejected records. Both watersheds completed.

16.1 Water Volumes and Accounting

Full-record measure (m^3)	260803	261009	Candidate
Hillslope streamflow	1,267,735,665.289	1,267,735,665.682	1,267,735,665.682
Outlet daily ledger	1,272,337,716.46	1,272,337,716.29	1,272,337,711.84
Sampled outlet integral	1,266,448,890.30	1,267,093,408.92	1,266,461,659.92
Integral minus ledger	-5,888,826.16	-5,244,307.37	-5,876,051.92
Sum of absolute daily gaps	6,441,305.44	5,849,097.91	6,380,525.06

Candidate totalwatsed matches 261009 on every day. The largest daily yield difference versus 260803 is 0.217 m^3. The reported outlet ledger changes -4.62 m^3 over the full record versus 260803 and -4.45 m^3 versus 261009. Sampled outlet volume changes +12,769.62 m^3 (+0.001008%) versus 260803,

and -631,749.00 m^3 (-0.049858%) versus 261009. Annual sampled-volume changes range -0.02413% to +0.04228% against 260803 and -0.08420% to -0.03010% against 261009.[12]

The normalization repair therefore causes a small systematic reduction in represented discharge relative to 261009, not only a change of day assignment. It does not reduce hillslope water production. Its signed integral–ledger gap is 12,774.24 m^3 smaller in magnitude than 260803's, and absolute daily gaps sum to 0.944% less. Relative to 261009, however, signed-gap magnitude grows 631,744.55 m^3 and absolute daily gaps grow 9.086%. Much of that release's apparent accounting improvement is lost when the incorrect source weighting is repaired. This is not a claim of restored conservation.

The remaining approximately 4.602 million m^3 difference between channel ledger and hillslope yield is also distinct from the sampled-integral gap. Neither is explained simply by saying that flow moved across midnight. Existing discrepancies are retained as limitations, not automatically made prerequisites for a routing reconstruction campaign.

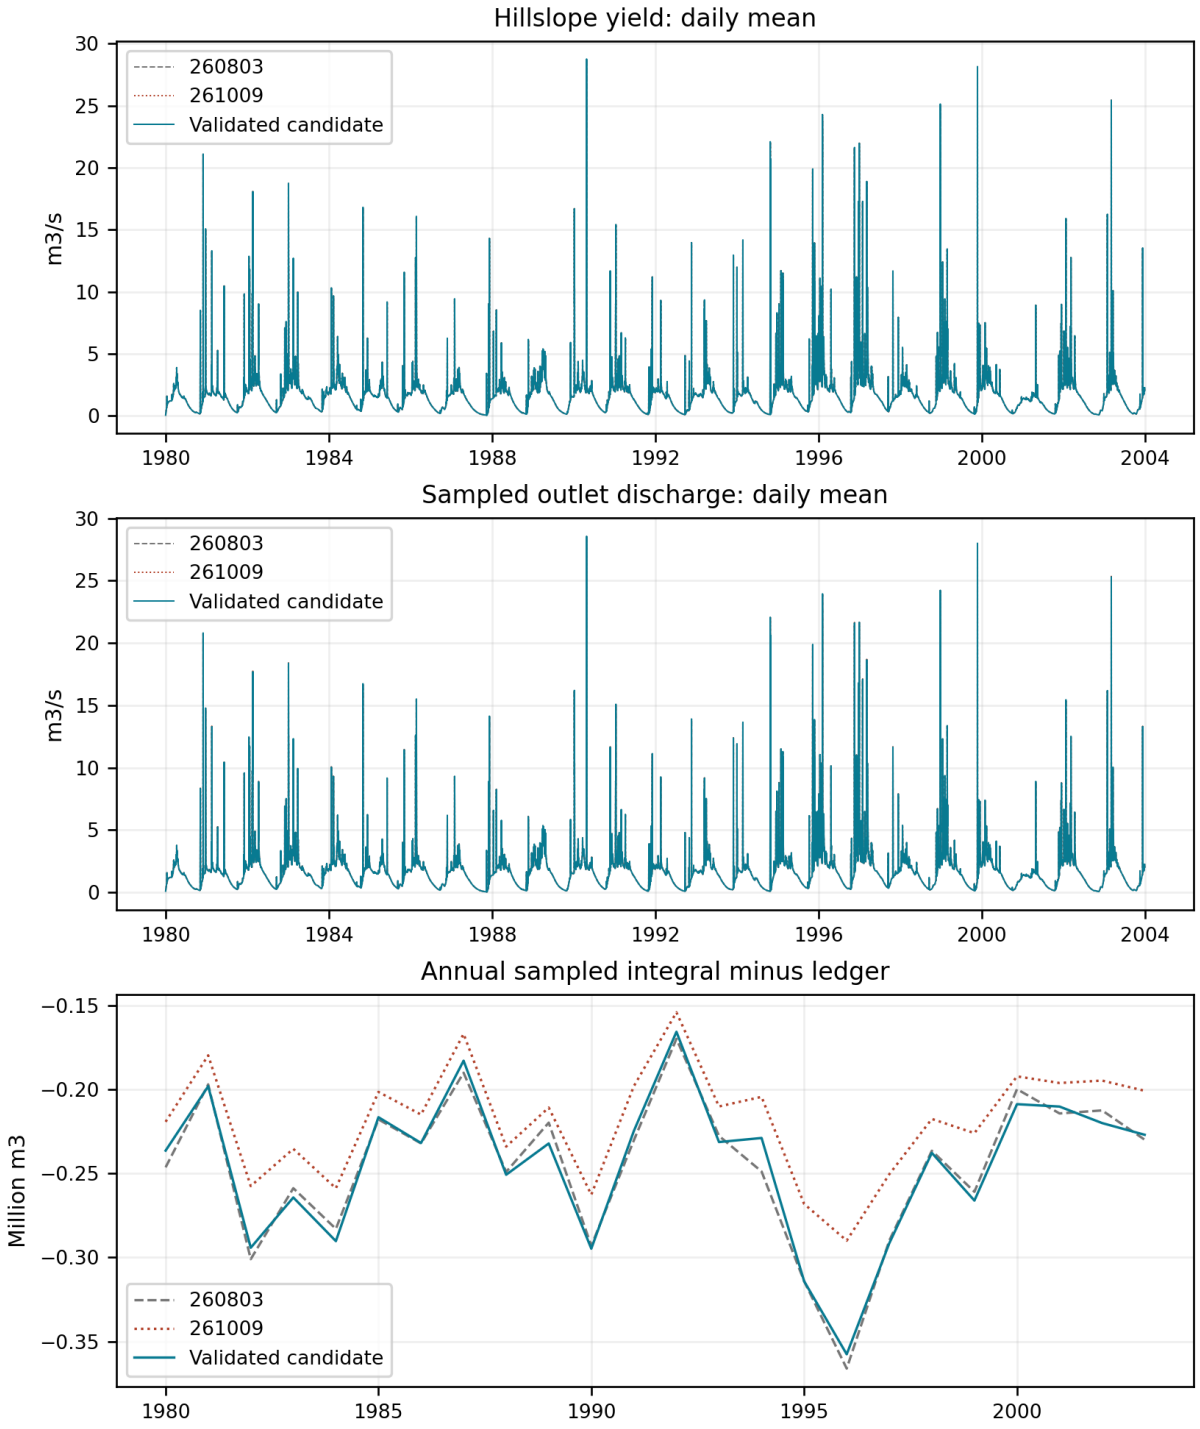


Figure 5: Cedar's complete record. The first two panels are daily means derived from hillslope yield and sampled outlet discharge, respectively; the third shows annual sampled-integral minus ledger volume. Closely overlapping series do not erase the accounting differences tabulated in the text. Original study build labels are retained.

16.2 Events, Recessions and Small Changes

For October 25–30, 1994, sampled volume is +4524.18 m^3 (+0.179%) versus 260803. The sampled peak changes 50.5 to 46.2 $m^3\,s^{-1}$ with unchanged first-peak time; the centroid is 20.82 minutes later. October 28 volume increases only 220.20 m^3 (+0.573%), and its gap improves from -7260.18 to -7172.74 m^3. The earlier large recession enlargement is not reintroduced in this control. The full-window gap is nevertheless worse than 261009's: -17918.33 rather than -10929.95 m^3. Relative to 261009, event volume changes -0.275%, peak -0.2 $m^3\,s^{-1}$ and centroid +0.47 minutes.

The pre-existing screen selects 322 seven-day windows from 260803 daily-mean peaks, using seven-day spacing and 0.1 $m^3\,s^{-1}$ prominence, then retaining three days before and after each peak. These are selection conventions, not new acceptance tolerances or universal storm boundaries. Window-volume changes range -0.154% to +0.619% versus 260803, with centroid shifts -4.38 to +20.78 minutes. Versus 261009, volumes range -0.374% to unchanged and centroids -1.26 to +2.44 minutes.

The retained 1991 small-event window is unchanged. The 1990 and 1998 windows change +0.355% and +0.619% in volume versus 260803, with centroid changes -2.23 and -0.76 minutes. Across all days, the largest baseline-relative absolute daily-volume change is +6720 m^3 on February 7, 1996, from 2,061,960 to 2,068,680 m^3. The largest incremental reduction versus 261009 is 6900 m^3 on February 8, 1996, from 1,959,780 to 1,952,880 m^3.

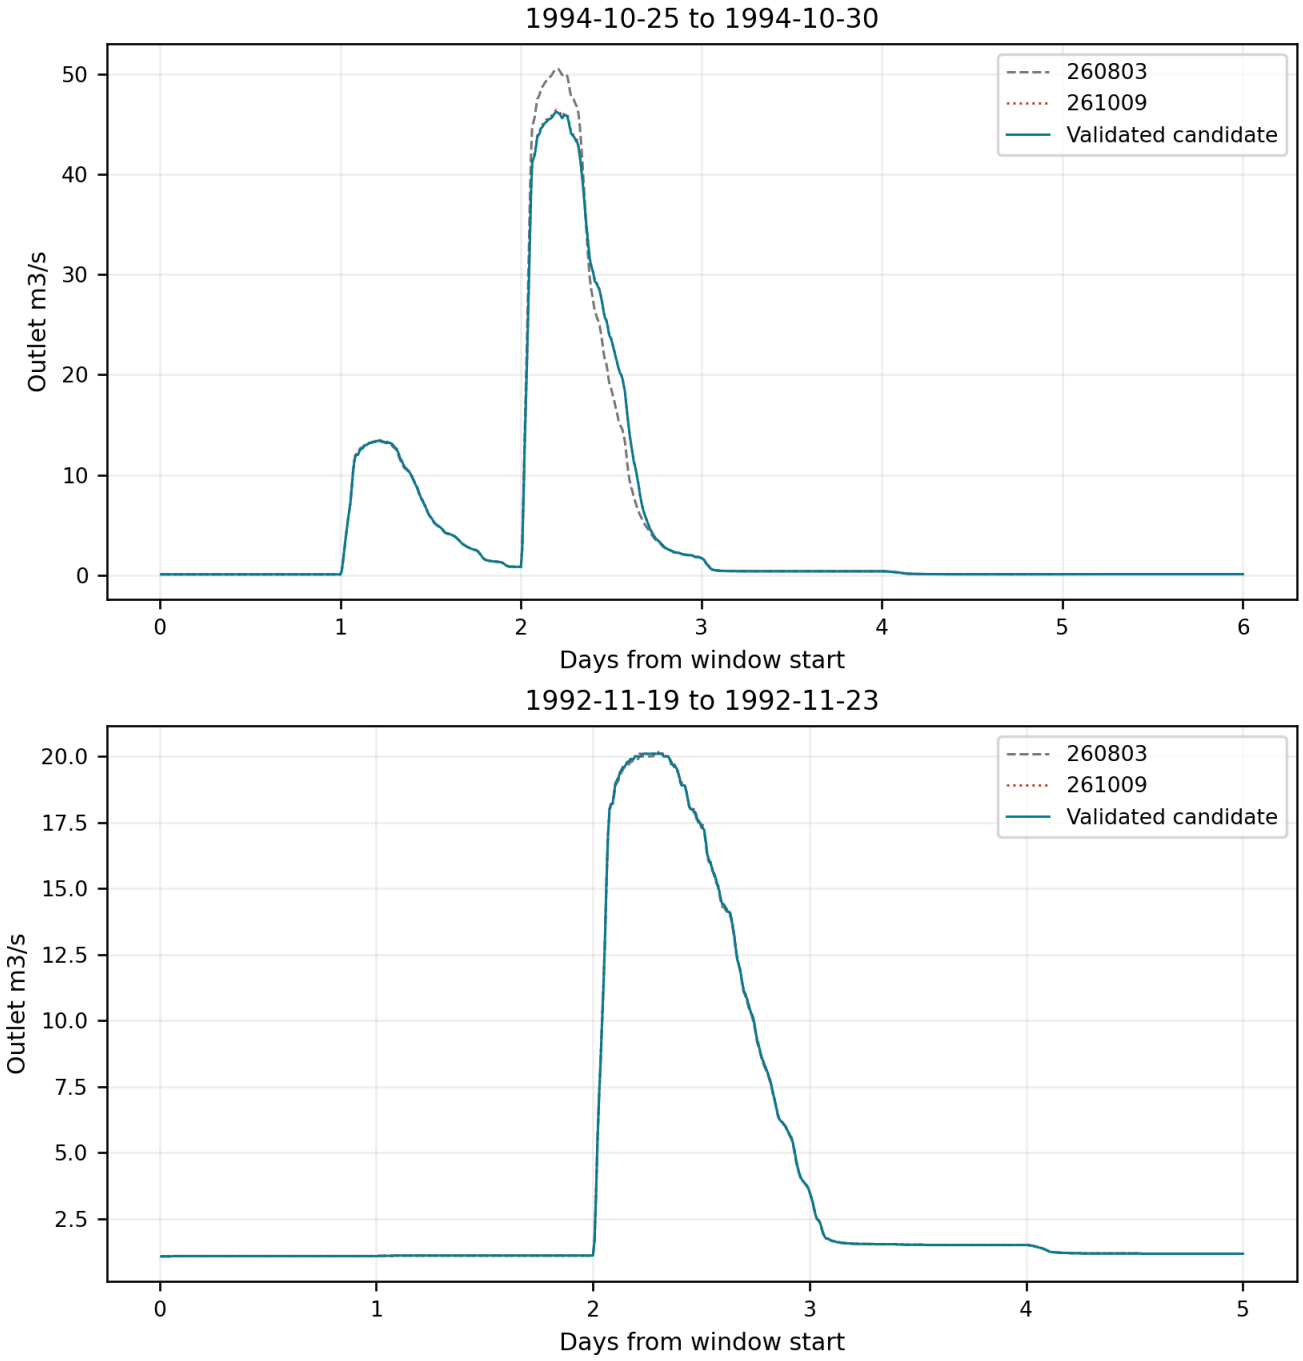


Figure 6: Retained Cedar October 1994 and November 1992 windows. Curves are 600-second output samples, not daily maxima. The comparison includes the complete frozen antecedent record. The selected windows are a consistent diagnostic, not a claim that every panel isolates one complete physical storm.

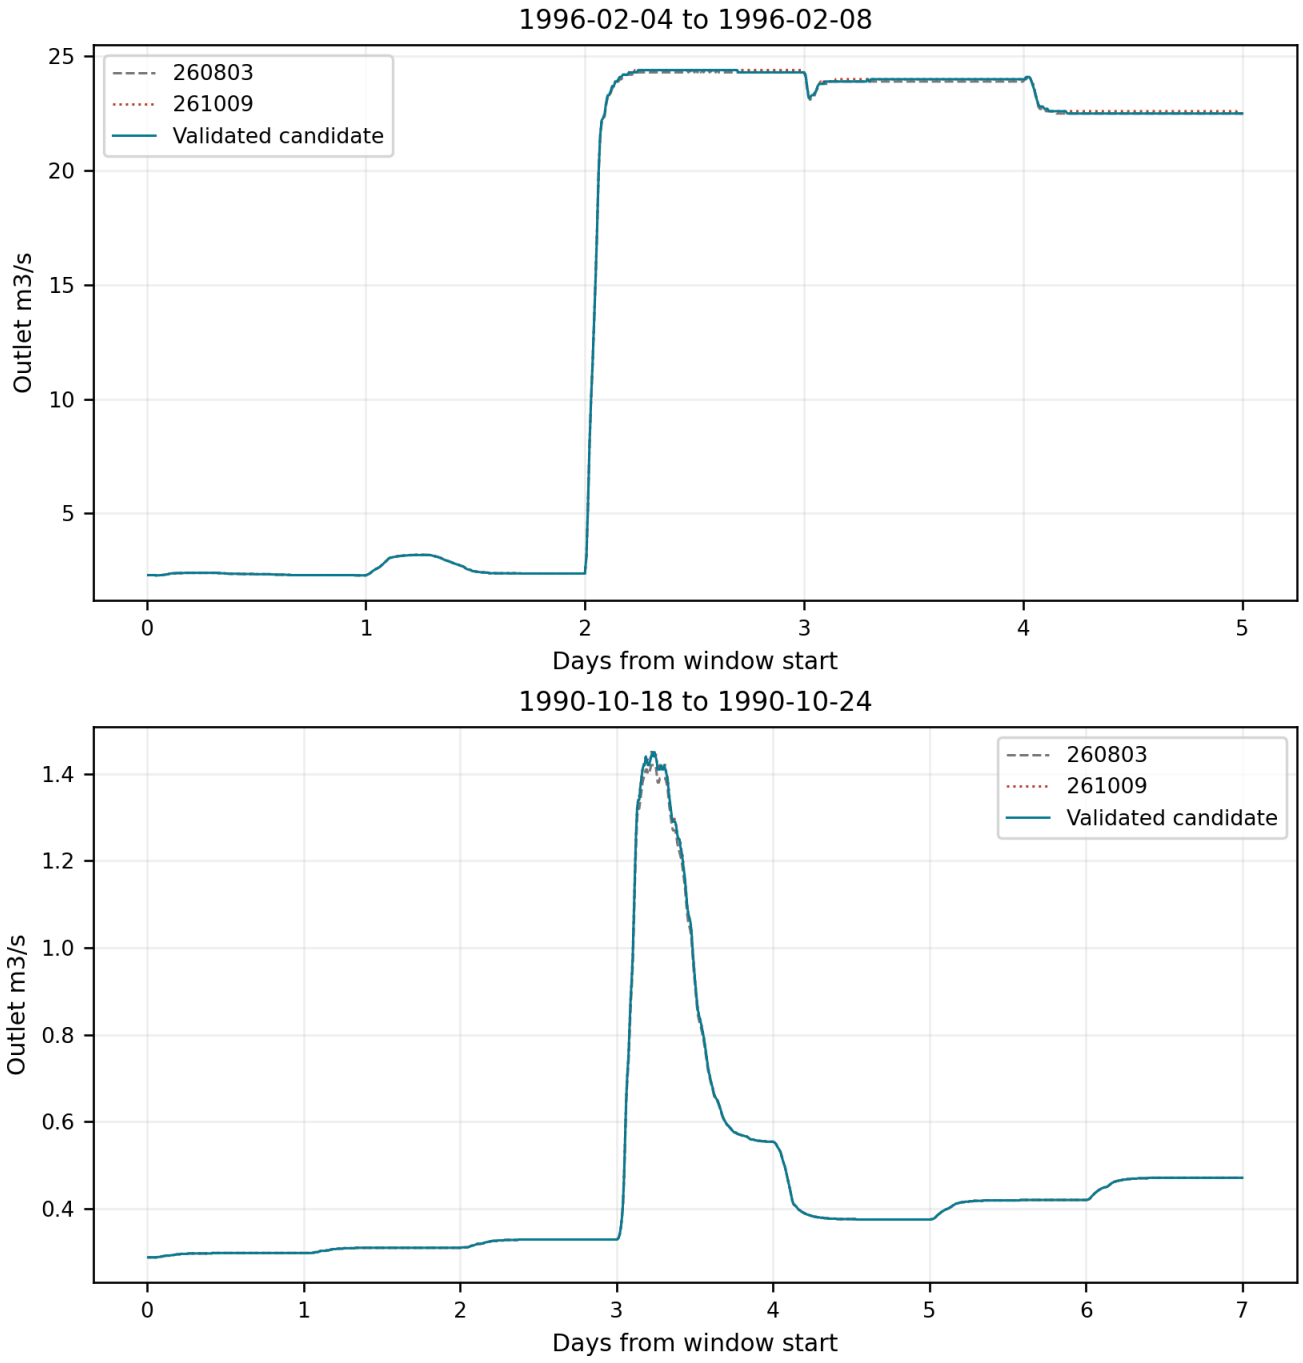


Figure 7: Retained Cedar February 1996 and October 1990 windows. The large event has a broad, rounded crest; the smaller control helps expose changes that full-record water totals alone would not identify.

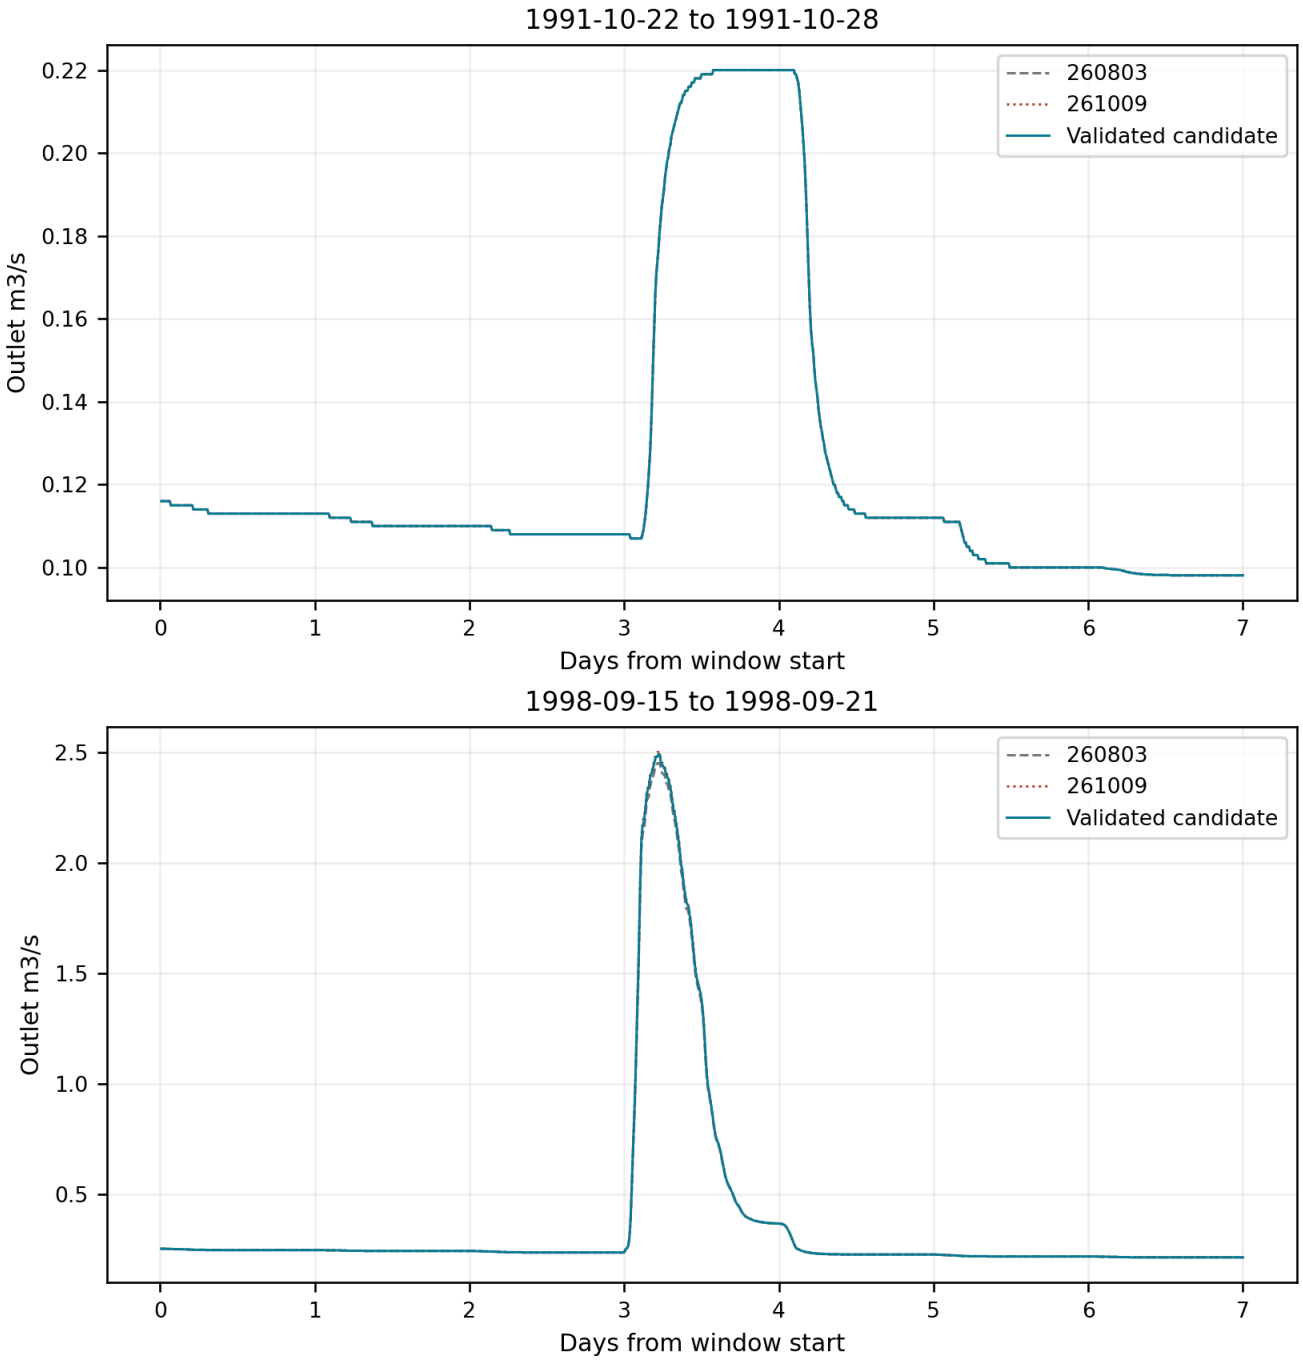


Figure 8: Retained Cedar smaller-event controls. The October 1991 window is unchanged; the September 1998 window retains a small volume and timing difference. Both comparisons use unchanged forcing and parameters.

16.3 Why First-Peak Times Can Mislead

The largest first-maximum shift in the window screen is -1220 minutes versus 260803 in late November 1984, while peak magnitude changes 3.94 to 3.93 $m^3\,s^{-1}$ and the centroid shifts -0.68 minutes. The candidate has equal printed maxima on both November 27 and 28; the baseline's slightly higher maximum occurs on the latter day. This is selection between near-equal crests, not a 20-hour translation of the whole hydrograph.

In the February 1996 retained window, the first maximum moves -320 minutes versus 260803, but the centroid shifts -0.29 minutes. Both printed maxima are 24.4 $m^3\,s^{-1}$; the baseline has 11 tied samples and the candidate 66. The timing audit preserves these facts rather than silently replacing an unfavorable peak-time statistic with a favorable centroid.

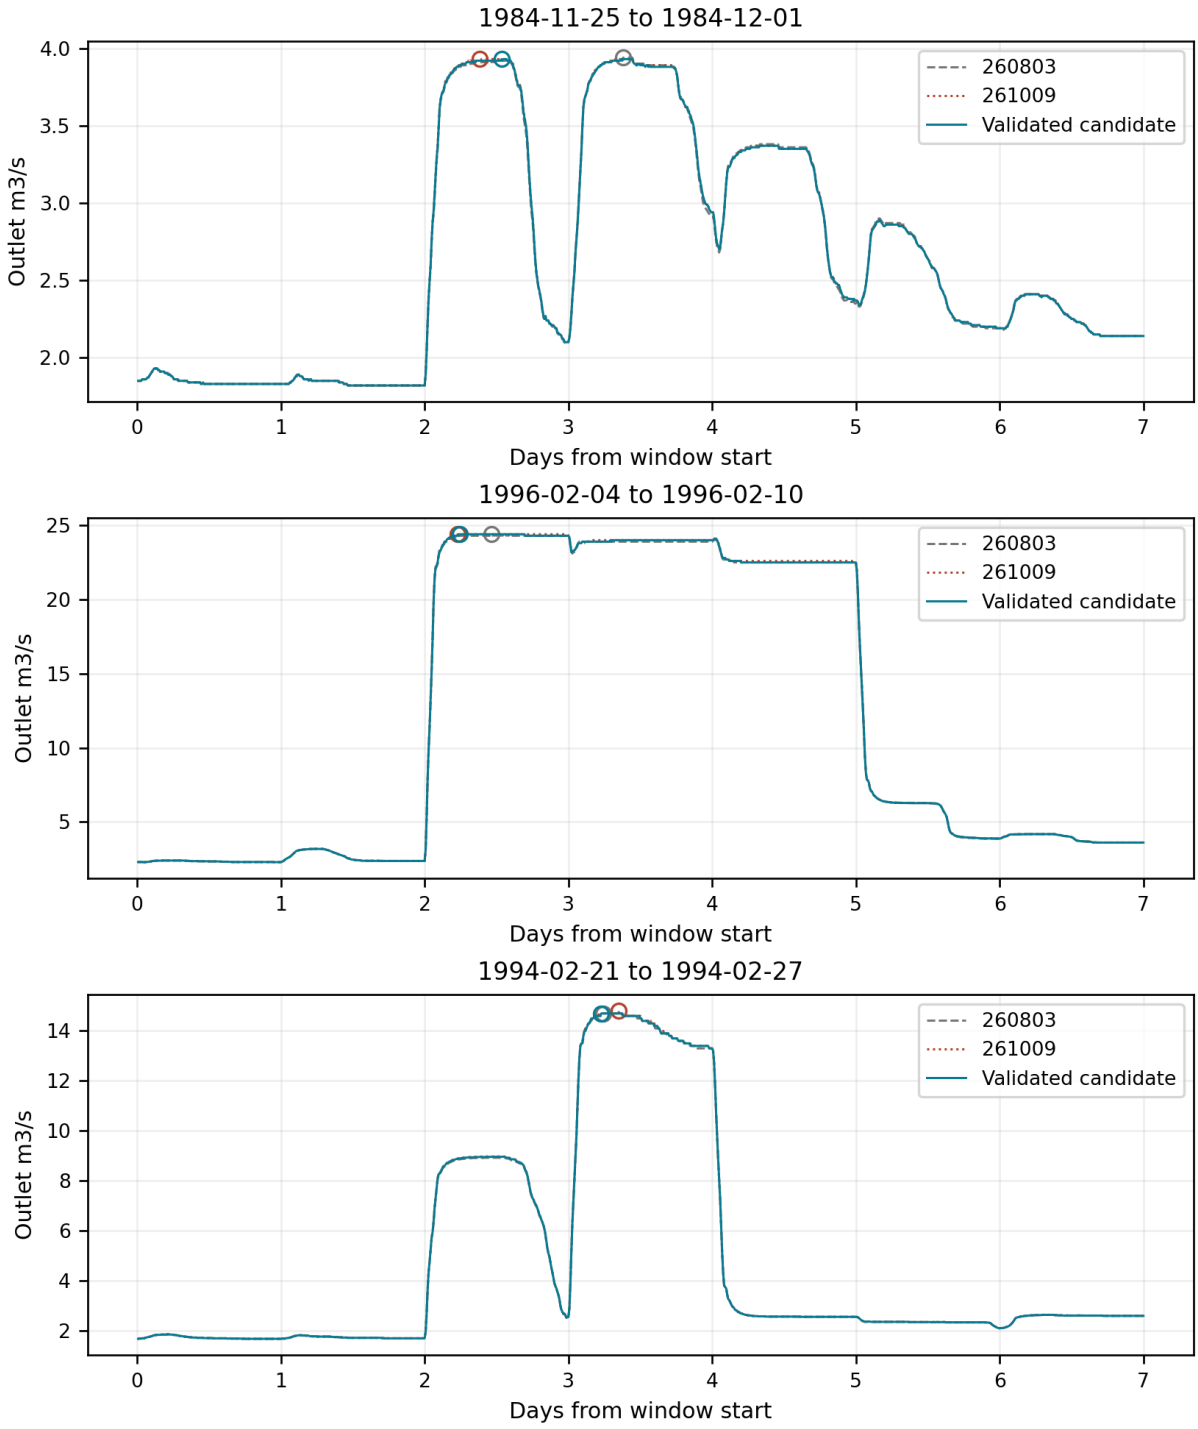


Figure 9: Timing audit. Markers identify the first printed maximum. Small changes to near-equal crests or rounded plateaus can move this selected time without a comparable movement of the whole hydrograph. This ambiguity is reported, not used as a blanket dismissal of timing effects.

At printed precision, 3348 days contain changed samples versus 261009; daily peaks decrease on 1038 days and increase on none. Of 1,262,304 samples, 67,080 decrease and 182 increase, with extremes -0.20 and +0.01 $m^3 s^{-1}$. No new positive sample appears on a previously zero sample. Across 2909 all-OFE dry-weather days, the net volume difference is -1679.04 m^3 versus 260803 and -1701.60 m^3 versus 261009. These counts do not certify the absence of every possible recession-shape change.

Daily-volume NSE/KGE/R^2 versus 260803 are approximately 0.999996/0.999013/0.999997; subdaily NSE is 0.999318. The practical conclusion rests on magnitudes, conditions, the event checks and examined timing ambiguities, not on these fit values alone.

16.4 Final-Release Sediment Readback

The following values were read for this document from the already completed Cedar runs, not from a new simulation. Hillslope delivery and channel erosion are sums of printed annual report fields; outlet export is summed from EBE. The measures have different roles and precision and are not an independently closed sediment budget.[12, 18]

Reported full-record quantity (kg)	260803	261009	Candidate
Hillslope delivery	62336.8	57898.0	57898.0
Channel erosion	442190.7	437916.7	437027.7
Outlet sediment export	202652.51	204130.91	203706.71

The final candidate changes hillslope delivery by -7.12%, reported channel erosion by about -1.17% and outlet export by about +0.52% versus 260803. The earlier +0.73% outlet figure described the 261009/v3 parent and must not be silently assigned to 261010. Identical hillslope delivery versus 261009 does not require identical channel export; source shape and channel sediment history can change. These totals do not settle the independent Rattlesnake event-level consequences in Section 15.

17 Disturbed Rankings and Parameterization Context

The canonical matrix uses four soil textures, four severity states and six vegetation types, including young forest, under one common 100-year synthetic McKenzie Bridge climate and an 87.9 m steep profile. Young, deciduous and mixed forest share generic forest burn templates. The release-build rerun completed 96 simulations plus three completion checks; all 768 output files match the validated candidate.[15]

Across the 24 soil/vegetation combinations, full-record weak ordering unburned $\leq$ low $\leq$ moderate $\leq$ high holds for runoff in 23/24, delivered sediment in 24/24 and largest event peak in 18/24. Exceptions remain: sandy-loam tall grass for runoff, and the four clay-loam forest types plus clay-loam and sandy-loam tall grass for maximum peak. Ranking exceptions are findings, not automatic implementation defects.

The current WEPPpy low-severity forest roughness of 6 cm is a separate parameterization revision, not a hidden binary change or a calibration aid used in Cedar. It replaces 4 cm in a shared template and four extended-lookup rows, aligning low burn with the other burned-forest defaults. The targeted assessment removed 17 sandy-loam low-over-moderate sediment inversions and was adopted as correction of a parameter inconsistency, not a claim that less sediment is inherently worse or better. Existing project-local overrides are not automatically migrated.[16]

The sandy-loam mature-forest sequence retains mean annual runoff of 92.7, 119.8, 121.3 and 376.4 mm and maximum peaks of 0.020084, 0.170380, 0.173370 and 0.313510 $m^3\,s^{-1}$ across the four severity states. This helps set expectations for these inputs; it does not establish universal burn ordering, accuracy under other climates/slopes, or observational sediment skill. Mutation inputs retain their historical frozen roughness, while this matrix uses adopted defaults. The two designs are not interchangeable.

18 Remaining Limits and Practical Judgment

The release is a bounded repair with a deliberately limited claim. It corrects two demonstrated implementation mechanisms and provides a workable interface for the hourly physical estimator. It does not establish exact mixed-flow timing, universal numerical conservation or observational validity for all hydrological and sediment quantities.

1. **Hourly means do not resolve subhourly return peaks.** Coincidence of rainfall and return maxima is assumed in the physical estimator. Its scope does not extend by assertion to multi-OFE or contoured producers.
2. **Return forcing is intentionally coarser than its input weights.** V3 uses their equivalent support, not their within-day positions. The support starts on a computational clock. It is neither a recovered mixed hydrograph nor an additional hillslope routing calculation.
3. **Local assembly and downstream integration differ.** The legacy lateral point-mode operator can average adjacent source ordinates even when return assembly supplies interval means. Earlier native probes showed 3300 of a 3600 m^3 last-hour source before midnight and 300 m^3 in the next first interval under that interpretation; a cold-start first-hour example gave 3900 m^3. These were controlled source-operator diagnostics, not Cedar water-loss estimates. No global mode switch was bundled to remove them.
4. **Short-support and flat-branch defects are not comprehensively repaired.** Fixing the normalization denominator does not make every supplied peak/volume/duration combination simultaneously attainable on a finite grid. The regression set explicitly retains an existing substep flat-source limit.
5. **State transfer is not storage normalization.** Equal-grid profile retention is exact for saved nodal discharge; changed-grid interpolation is bounded but does not prove exact water storage conservation.
6. **Sediment response remains nonlinear and incompletely explained.** The large Rattlesnake event changes require visibility in application decisions. The release decision does not establish that their absolute magnitudes are observationally correct or merely shifts between adjacent storms.
7. **Validation coverage is finite.** Three watersheds, targeted controls, one canonical soil/cover matrix and a mutation census do not cover all channel geometries, forcing resolutions, managements, coupled chemistry or impoundment behavior. The report does not claim such coverage.

The practical standard is not served by retaining known source inflation to improve an unrelated aggregate discrepancy, nor by using a favorable aggregate fit to dismiss a consequential event. Roger's promotion decision followed the recorded studies. Future contradictory evidence should trigger a focused reassessment with the original results preserved, not undocumented calibration or cumulative patches to obtain an attractive metric.[17]

19 Release Artifacts and Operational Use

The dated pair is `wepp_261010` and `wepp_261010_hill`. Both were built with system gfortran 13.3 from source `7471bb5e9` on the verified default branch, `wepp_260430_negmeltfix_comparator`. The source-state decision is distinct from the subsequent artifact/documentation commit. Full identities appear in Appendix D. The original runtime banner is retained; it is not the sole way to identify a dated build.

The build links MIXPEAK, component transport/assembly and CHRSTATE into both executables. Watershed and hillslope builds are performed sequentially because their build rules select shared include profiles. An earlier concurrent mixed-profile attempt produced an invalid watershed capacity profile and was rejected; it is not hydrological evidence. Capacity includes were restored and verified without a capacity-contract change.

The existing paired native reader supports component v3 and retains public Parquet schemas. The reader/composition implementation is associated with source `c3d8481d1c7d`; the binary sidecars pin its packaged artifact identity. No additional native-reader algorithm change accompanies the final CHRQIN normalization edit.

19.1 Selecting the Build and Regenerating Inputs

The never-deployed 261009 binary pair and sidecars were removed from WEPPpy at Roger's request; historical results and Forest release artifacts remain identified as 261009. A local project that explicitly selects that removed name must select 261010 and regenerate every contributing hillslope PASS with the matching build. There is no hidden alias, cross-release PASS swap or automatic project migration. Roads and AgFields combinations require matching versions and component-preserving native composition.

General project defaults were not switched by this vendoring action, and no production deployment was performed. Changing a binary selection does not automatically rewrite project-local management/soil overrides. Users adopting the independent low-burn roughness revision must review regenerated inputs separately from the binary change.

19.2 Source and Documentation Publication

This document explains legacy implementation details through prose, equations and interface descriptions. Such explanations and original code we author are explicitly allowed under the release-note authoring policy. It does not embed legacy source, mixed-origin diffs or raw source attachments. The Python listing is independently authored for this report and is clearly illustrative.

The current report remains a review draft. Scientific/implementation checks performed during authoring are not a substitute for Roger's publication decision or confirmation of third-party/project-data disclosure scope. The PDF has not been vendored to WEPPpy. Once approved, the frozen PDF and public manifest can be copied without rebuilding, with hash verification and a tested download path. Neither PDF publication nor binary vendoring is authorization for production deployment.

20 Quality Gates and Reproducibility

Required gates use committed inputs or reproducibly generated outputs from committed inputs. Private project directories are supplementary scientific evidence, not prerequisites that a new checkout cannot

satisfy. Historical replay gates requiring retired private resources remain retired.

Check	Recorded release result
Maintained Forest suite	163 passed
Committed-resource hillslope watchlist	12/12 passed
Forest host smokes and runtime provenance	Both executables passed
Post-vendor host/container smokes and provenance	Both executables passed
WEPPpy runner/output/interchange	114 passed; one existing skip
Fresh 261010 disturbed matrix	99 passed, including 96 simulations
Report freshness, route coefficients, Usersum rendering	63 passed
Artifact policy and scoped documentation checks	Passed

The native sample-normalization regression failed in six targeted cases before the fix; the combined focused source/return/state checks passed afterward. The complete maintained suite includes those checks. Default pytest discovery now targets `tests/`; otherwise archived research modules introduced duplicate names and missing research dependencies during collection. No maintained numerical test was excluded or weakened to obtain a pass.

Supplementary exact-release validation regenerated 140 hills for the full 1980–2024 period. Plain, Roads additive and AgFields weighted watershed cases completed all 45 years. Native hillslope conversion accepted 2,301,180 records with zero rejection. Doubling a selected source through additive composition doubled its supplied volume from 195011.111 to 390022.222 m^3; the 0.3/0.7 weighted composition retained 195011.111 m^3. These are interface controls, not calibrated disturbed or Roads scenario predictions.

Incomplete checks remain distinguishable from successful ones. An optional additional committed-fixture watershed could not start because three required topology/channel/impoundment files were absent; 301 hillslope runs had completed. That attempt is unavailable, not passed, and no private substitute or invented topology was introduced. One WEPPpy interchange test skips because its static PASS dataset is absent. The full unrelated WEPPpy suite was not run for this handoff. Early command/environment failures and an analysis-only column-name error were corrected and retained separately from valid results.[17]

The figures and tables in this report identify their actual study cohort and baseline. They are not newly simulated outputs. The report's retained input inventory records data and figure identities; local raw project data may be needed to reproduce the studies, but not to compile this self-contained document or understand its reported conclusions.

A Symbols, Units and Interface Semantics

Symbol/field	Units	Meaning and caution
R, S, R_r	m	Final runoff, return and non-return depths on one area basis.
p_r, p_s, p	$m\,s^{-1}$	Rainfall, maximum hourly return and combined physical depth-rate peaks before legacy final bookkeeping.
V_r, V_s	m^3	Serialized component volumes.
V_c	m^3	Surface volume allocated by the watershed caller to this source.

Symbol/field	Units	Meaning and caution
Q_r^I	$\mathrm{m^3\,s^{-1}}$	Rainfall-only interface operand after existing floor/cap conventions; not necessarily the reconstructed sampled maximum.
w_h	relative	Twenty-four finite nonnegative return weights; common scale cancels in the v3 support ratio. Hour positions are not v3 forcing times.
D_{eff}	s	Equivalent runoff/peak duration used in erosion bookkeeping.
T_r, T_s	s	Existing rainfall support and local return-only support, respectively. Neither alone establishes actual event timing.
`tcs`	h	Existing concentration-time scalar; may change indirectly with hydraulics although its formula is unchanged.
`oalpha`	dimensionless	Existing ratio involving concentration time, peak, area and volume, with its legacy minimum bound.
`nt0`	index	Caller initialization convention. Zero includes a time-zero boundary; one starts at the first positive-time sample.
`qs` profile	$\mathrm{m^3\,s^{-1}}$	Channel nodal discharge state, not a new stored hillslope-return volume.
G_d	$\mathrm{m^3}$	Sampled outlet integral minus the reported daily outlet ledger. A signed sum can hide cancellation; also inspect absolute daily gaps.

B Descriptive Fit Metrics

For reference series x_i and candidate series y_i, the retained analysis uses the following familiar definitions, with r the Pearson correlation:

$$\mathrm{NSE} = 1 - \frac{\sum_i (y_i - x_i)^2}{\sum_i (x_i - \bar{x})^2}, \tag{15}$$

$$\alpha = \sigma_y/\sigma_x, \qquad \beta = \bar{y}/\bar{x}, \tag{16}$$

$$\mathrm{KGE}_{2009} = 1 - \sqrt{(r-1)^2 + (\alpha-1)^2 + (\beta-1)^2}, \qquad R^2 = r^2. \tag{17}$$

Undefined constant/zero-reference cases require explicit handling, not a fabricated perfect score. The reference here is normally another model run, not observed discharge. Correlation, variability and mean agreement are different properties, and their aggregate combination does not replace event-level examination.[4, 5]

C Implementation Coverage

The complete net non-binary `src/` difference between the identified 260803 and 261010 source states is mapped below. The internal audit also accounts for tests, tools, documentation and retained release artifacts. Source paths are identifiers, not links exposing restricted code.

Component	Change IDs	Responsibility
`mixpeak.for`	RN-01, RN-02	Physical component peaks, rainfall-only call/state isolation, interface operand and reset helper.
`irs.for`	RN-01, RN-02	Eligibility, producer invocation and physical peak replacement before retained duration/reassignment behavior.
`cmixprod.inc`	RN-02	Current-event producer state.
`contin.for`	RN-02	Daily reset and hydraulic observations.
`inidat.for`	RN-02, RN-03	Producer/consumer initialization.
`mixpass.for`	RN-03, RN-04	Headers, validation, serialization, master records, daily clear and bounded return assembly.
`cmixsrc.inc`	RN-03	Current-day source operands and flags.
`infile.for`	RN-03	Master header interpretation.
`wshpas.for`	RN-03	Extended hill/master writers and ingestion, plus derived-scalar observations.
`wshred.for`	RN-03	Daily clear and master source consumption.
`chrqin.for`	RN-04, RN-06	Component allocation, rainfall/source selection, return call and corrected normalization.
`chrstate.for`	RN-05	Saved channel profile and bounded grid transfer.
`wshchr.for`	RN-05	Restore/save calls around existing routing.
`makefile`	RN-08	Matching linkage in watershed and hill executables.

The immediate 261009-to-261010 model delta is confined to RN-06. Companion RN-07/RN-10 changes belong to the reader/WEPPpy integration and parameterization records, not additional undisclosed Fortran files. RN-09 accounts for captured regressions, observer tests, resource/discovery policy and documentation. Historical artifact retention is not a claim that every experimental binary or patch is part of the selected release.

D Release Identity and Evidence Index

Selected model source

`7471bb5e981d14d0b8c1cdb88a16af305aed1b67`

Watershed executable SHA256

`1dd1ca75cf53f9a0606cf5a598312d4e680e161df631156360d4a20dbcb6f17e`

Hillslope executable SHA256

`d8ea3a07a29ef1e754c5362bd7931cc3a93486d75faa32df5c5f2e821c22d698`

Paired native reader artifact SHA256

`06319303a2337cc75a3c57eb594161565d8a2af1d58cee18c17abcfe6eea8cf8`

Paired native reader source

`c3d8481d1c7dcaf59dceb37d9ea2e860b89f4307`

Studied candidate watershed SHA256

`0eb062812786064b230091ac0144f7502bc225964ea982dd76ebfc68abc6b76d`

Studied candidate hillslope SHA256

`219be50ac94a7ffa01589eb22251586caea53d5e856f7af0feaea2c12fa6433a`

Common normalized watershed SHA256

`d3b5fb0fb75fa1fbe0777e234388404239dcfd0b88dea09f2c7172a139a59df5`

Common normalized hillslope SHA256

`5a65377c5f923de3d0fbb9db4e608408fe06f159f6a874f57afafa4c75ceab89`

The studies distinguish original candidate hashes from packaged release hashes. Binary equivalence excludes debug data and GNU build identity only, not runtime code or data. Document revision 0.1 is an authored review draft; its PDF hash and frozen authoring inputs are recorded separately from these model identities. Publication approval is pending.

The project evidence references below identify retained records rather than claiming public availability of private source or raw project data. The report contains the principal equations, tables, figures, interpretation and limits needed to understand the release without opening those repositories. No source-code attachment is embedded in the PDF.

References

[1] Stone, J.J., Lane, L.J., Shirley, E.D., and Hernandez, M. (1995). Hillslope Surface Hydrology. Chapter 4 in *USDA–Water Erosion Prediction Project: Hillslope Profile and Watershed Model Documentation*, NSERL Report 10. `https://www.ars.usda.gov/ARSUserFiles/50201000/WEPP/chap4.pdf`.

[2] Foster, G.R., Flanagan, D.C., Nearing, M.A., Lane, L.J., Risse, L.M., and Finkner, S.C. (1995). Hillslope Erosion Component. Chapter 11 in *USDA–Water Erosion Prediction Project: Hillslope Profile and Watershed Model Documentation*, NSERL Report 10. `https://www.ars.usda.gov/ARSUserFiles/50201000/WEPP/chap11.pdf`.

[3] Wang, L., Wu, J.Q., Elliot, W.J., Dun, S., Lapin, S., Fiedler, F.R., and Flanagan, D.C. (2010). Implementation of channel-routing routines in the Water Erosion Prediction Project (WEPP) model. *SIAM Mathematics for Industry: Challenges and Frontiers (MI09)*, pp. 120–127. `https://research.fs.usda.gov/treesearch/40457`.

[4] Nash, J.E., and Sutcliffe, J.V. (1970). River flow forecasting through conceptual models, Part I: A discussion of principles. *Journal of Hydrology*, 10(3), 282–290. `https://doi.org/10.1016/0022-1694(70)90255-6`.

[5] Gupta, H.V., Kling, H., Yilmaz, K.K., and Martinez, G.F. (2009). Decomposition of the mean squared error and NSE performance criteria: Implications for improving hydrological modelling. *Journal of Hydrology*, 377(1–2), 80–91. `https://doi.org/10.1016/j.jhydrol.2009.08.003`.

[6] WEPP Forest project record (2026-10-10). Shipped-source audit, 260803 source `f24c957e3` to 261010 source `7471bb5e9`; source/component tests and release-sidecar identities. Internal implementation evidence; legacy source not reproduced in this document.

[7] WEPP Forest project record (2026-10-08, with subsequent promotion decisions). *Return support feasibility contract* and *Return-support v3 feasibility result*. Native source controls, H106 reproductions, matching-build watershed and interface evidence. Historical experiment status is superseded by the documented release decision, not by alteration of the retained measurements.

[8] WEPP Forest project record (2026-10-08). *Hourly routing: midnight diagnosis and implementation results*. Native 48-hour restart, grid-transfer and source-operator diagnostics; no claim of universal storage conservation.

[9] WEPP Forest project record (2026-10-08). *Return alternatives assessment*. Eligible Cedar return population, hourly timing/scheduling evidence and bounded source-rule alternatives.

[10] WEPPpy/wepppyo3 project records (2026-10-09–10). *PASS components v3 release compatibility* and *PASS components v3 integration*. Native reader/composition contracts and generated plain/Roads/AgFields evidence.

[11] WEPPpy project records (2026-10-10). *Mechanism Located: CHRQIN Sample Normalization* and *CHRQIN Normalization Candidate Results*. Retained Topanga/Rattlesnake source probes, observer-neutrality checks, paired outputs, event/annual statistics and limitation records.

[12] WEPPpy project record (2026-10-10). *Cedar Hillslope Yield and Channel Discharge*. Fresh paired 24-year simulations, totalwatsed conversion, sampled outlet series, ledgers, event/timing audits and provenance.

[13] Lew, R. (2026). *Small-Mutation Screening of WEPP Hillslope Peak Flow: Remediation Strategies and Validation Findings*. Version 0.2, October 6, 2026. University of Idaho; prepared with agentic coauthor support from OpenAI Codex. Companion investigation covering the original census and earlier remediation evidence, not final 261010 validation. Download the companion investigation report (PDF).

[14] WEPPpy project record (2026-10-10). *Candidate Mutation and Disturbed Ranking Studies*. Frozen 1980–2024 mutation census, event/sediment ledgers, observer checks, figure filters and exact 261009 comparisons.

[15] WEPPpy project record (2026-10-10). *Disturbed Parameterization: Current Canonical Results*. Fresh 261010 96-case report and source/output provenance.

[16] WEPPpy ADR-0083 (2026-10-09). *Low-Severity Forest Initial Random Roughness*. Decision provenance, 112-run targeted assessment and linked broader sensitivity study; parameterization change separate from model binaries.

[17] WEPP Forest and WEPPpy release records (2026-10-10). *WEPP 261010 Release Validation*, release handoff and ADR-0084, *CHRQIN Source Normalization and WEPP 261010*. Exact artifact identities, required and supplementary checks, skips/unavailable attempts, executable equivalence and release authorization.

[18] This report, revision 1 (2026-10-10). Additional read-only Cedar sediment summary from the retained 260803, 261009 and candidate annual/EBE output files. Input hashes, parsing method and numeric results are retained with the authoring evidence; no new model simulation was performed.