Response to the Third-Party Audit of Package Revision 4

Site: followtheepicenter.com **Analyst:** Jon Bray **Date:** 10 Nov 2026 **Supersedes:** package rev. 4 → rev. 5 (`forensic_acoustic_package_r5.zip`)

The audit confirmed the six input hashes, all six onset picks to the sample, the geometry, and the 200-trial null tests (1.88 ms; 0/200 and 1/200), and raised five issues. All five are accepted in substance. What each one changes is set out below; the corrections are logged as C25–C29 in `docs/results_v6.md`, Appendix A. The auditor's full report ("Independent technical verification … R4", 10 Nov 2026) was received after the first draft of this response; the additions it prompted are in the final section.

1. PA synchronization uncertainty — accepted; my stated allowance was wrong

The auditor's per-camera deviations reproduce exactly from the KML speaker pins: 6.6 ms (Cam 6), 7.0 (Cam 11), 9.2 (Cam 7), 9.3 (Cam 8); also 8.9 (2.MOV), 8.7 (IMG_6368), 6.8 (Canon). Revision 5 had said ±3 ms for the far cameras and "negligible" for the Canon, and `predictions.py` carried a comment that the speakers were "within ~1 m" for path purposes. The speakers are up to 5.7 m apart; both statements were wrong (C26).

The error is per camera and common to every pair that camera takes part in, so the loop closures I relied on cannot detect it. Rev. 5 treats it as a systematic and enumerates the speaker assignments (16 per pair, 4^7 for seven-camera fits) rather than folding it into a Gaussian σ.

What it changes. The auditor's 17.4 ms figure is for the Cam 8–Cam 7 pair, and it is right: speaker choice can reduce that pair's excess from +17.6 to +1.7 ms. It cannot close Cam 6–Cam 7 (minimum +15.7 ms over all 16 assignments) or Cam 11–Cam 7 (+13.3 ms), because Cam 6 and Cam 11 are far enough from the PA that their own speaker spread is only 4–6 ms. The Canon's violation stays ≥ +10.3 ms. So Cam 7's exclusion from the geometry fits rests on Cam 6 and Cam 11, not on Cam 8, and survives; the model-free ordering result (tent excluded as the source) also survives — the near cameras remain 37–65 ms after Cam 6 under every assignment. What does not survive unchanged is the pinned-firing-point test (item 3) and the 247° significance (item 4).

Speaker identities. Not verified. They could be, from the pre-event audio itself: with two speakers contributing comparably at a camera, the cross-correlation should show side-peaks at the inter-speaker delays (3–17 ms). That analysis is listed as open item 1a in Part V and has not been done.

2. Single-wavefront analysis — accepted

Agreed on both halves: passing the bound is necessary, not sufficient, for a common wavefront; failing it establishes that at least one of {same wavefront, positions, sync, PA path, onset pick, unobstructed path} is false, not which. Rev. 5 says this in §8a and no longer describes Cam 7's pick as "a different arrival" as if that were established; it is excluded because it cannot be fitted as the group's wavefront, and the cause is open.

The 20,000-draw Monte Carlo was run from code that was never shipped (C27). It is now `scripts/bound_mc.py`, seeded (20261110), reading `output/results_arrivals.json`, with speaker assignment as a fourth perturbation under four priors (tent-centre, nearest, $\propto 1/r^2$, uniform). The shipped values differ from the ones I quoted — at 1 m position σ, P(excess ≤ 0) is 0.008 / 0.017 / 0.060 / 0.002 for Cam 6–7, Cam 11–7, Cam 8–7, Cam 6–Canon under the rev.-4 assumptions, versus the 0.016 / 0.028 / 0.076 / 0.016 I published

— and the published numbers should be treated as unreproducible. With a uniform speaker prior: 0.009 / 0.018 / 0.10 / 0.005. The joint probability that Cam 7 is bound-consistent with Cam 6, Cam 11 and Cam 8 at once is 0.004–0.005 at 1 m, 0.01 at 2 m, 0.02 at 3 m, 0.06–0.07 at 5 m. The σ values remain assumptions; the auditor's own sensitivity analysis should be compared against the shipped script when the report arrives.

3. Firing-location exclusions — accepted; rev. 4's rejections are withdrawn

The audit is right that rev. 4 excluded Cam 7 from the model fits but kept it in the fixed-pin χ^2 (4 dof), which is inconsistent (C25). Recomputed against Cam 11, 2.MOV, Cam 8 (3 dof, no free parameters):

pin	v = 830 m/s, tent-centre PA correction	minimum over 4^4 speaker assignments
"shooters position" (61.8°, 150 m)	χ^2 7.96, p 0.047 (auditor's value reproduced)	χ^2 2.6, p 0.46
"n wave source" (57.3°, 161 m)	17.2, p 0.0006	6.2, p 0.11
prosecution pin (55.3°, 135 m)	27.7, p < 0.001	12.6, p 0.006

(780 and 880 m/s in `results_output.md`, Table R6.) The "shooters position" pin is not excluded. The "n wave source" pin is weakly disfavoured. The prosecution pin remains disfavoured but not at the level rev. 4 implied, and all three pins share a systematic — they predict the near cameras 10–18 ms earlier than observed — that a ~3 m error in Cam 6 or Cam 11 along the line of fire, or the speaker assignment, would absorb. The honest statement is that the timing data no longer reject the "shooters position" pin and that the test is dominated by systematics nobody has measured. It does not select a firing point.

4. The 247° result — agreed, with one addition

Bearing, emission time and null results reproduce. Rev. 5 adds that the fit is sensitive to the speaker systematic: over the 4^7 assignments the bearing stays within 236–261° (nearest-speaker assignment 247.5°, rms 2.2 ms), but the rms residual ranges 0.7–8.1 ms (median 4.6). The 1-in-200 figure was computed under the tent-centre correction and is now stated as conditional on it (C28). Re-deriving it fairly means running both observed and null under the same speaker freedom; at ~1 h per null run I have not done that in this revision, and the checklist (item G4) says so. Physical identity of the source and whether every selected peak is one wavefront remain open, as before.

5. Reproducibility — fixed

`extract.sh` now uses `-af "pan=mono|c0=c0"` (and `c0=c1` for the right channels), which runs on FFmpeg 6 and 7; the output is bit-identical to the `-map_channel` extraction (all ten derived WAV hashes unchanged against the manifest, re-verified end-to-end). Input names now follow the manifest (`img6368.mov`, `canon_ch3.wav`, `canon_full.wav`) (C29). The rev. 5 manifest is regenerated and verified against the zip.

Additions after reading the full report

- **Monte Carlo agreement.** The auditor's independent 200,000-draw simulation gives 0.7% (Cam 6–Cam 7) and 0.2% (Cam 6–Canon) at 1 m; `bound_mc.py` gives 0.8% and 0.2% under the same assumptions. The two implementations agree with each other and not with the numbers I published in rev. 4/5, which confirms C27.
- **Cam 6 k=3 detection at 7.004 s (report §2).** Checked, and it is more than the auditor said: with the package's own k=5 criterion there are crossings at 6.976–7.010 s (ratio up to 6.7×), 80–110 ms before the Cam 6 pick, and at Cam 11 at −149…−116 ms; my Cam 6 search window started at 7.0 s and hid them

(C30). Examined: both are 40–47 ms wide at half-maximum, 20–25× weaker in the 2–15 kHz band than the picks, and LF-dominated — a sustained rise in ambient level, not an impulsive arrival; nothing appears at the other five cameras in the corresponding true-time window. The picks stand, but the criterion as published would fire on a sustained rise, so rev. 5 adds an impulsiveness requirement (half-maximum width < 5 ms) that the picks pass and these do not.

- **Cam 7 onset spread** (k=3 to k=8: 25.3472–25.3605 s) is logged as an asymmetric uncertainty, −3.1/+10.2 ms (C31). Taking the earliest value lowers the Cam 6–Cam 7 minimum excess over speaker assignments from +15.7 to +12.6 ms; still positive.
- **Clopper–Pearson**: the one-sided 95% upper limit of 2.35% for 1/200 is the correct figure to quote alongside the two-sided 2.8%; rev. 5 quotes both, and both are now conditional on the tent-centre correction (item 4).
- **Waveform morphology (report §7, item 3).** Not done. Bound-consistency is treated in rev. 5 as necessary only; a same-wavefront argument from waveform shape and acoustic-path controls is listed as open work.
- **Canon files.** The two full Canon WAVs (65 MB each) were not available to the auditor. Their hashes are in the manifest; the files can be supplied directly. Until then every Canon statement in the package is unverified by anyone but me, and the package says so.

Requested items

- **Exact 20,000-draw simulation:** `scripts/bound_mc.py`, output `output/bound_mc_output.md`.
- **PA speaker/path assignments used for synchronization:** none were assigned — the correction used the tent centre for every camera, which is the point of item 1. The per-speaker path table is in `results_v6.md` §6; the identities remain unverified.
- **Independently calibrated camera positions:** none exist. Positions are the KML pins with an assumed 1 m σ. This is the single largest open item and is listed first in Part V.

What now stands, and what does not

Stands: seven recordings synchronized; the first arrival reached Cam 6 first and the tent-adjacent phones ~50 ms later under every speaker assignment, so the first sound did not originate at the tent; Cam 7's and the Canon's first-arrival picks are not that wavefront at the 1–3% level if positions are good to 1–2 m; no distinct tent-origin impulse apart from the coincident one; the 247° association reproduces and its bearing is robust.

No longer stands as stated in rev. 4: the rejection of all three pinned firing points (withdrawn; "shooters position" compatible); the ±3 ms PA allowance; the unconditional 1-in-200 significance for 247°; the Monte Carlo numbers quoted without code.

If the auditor's sensitivity analysis or report shows an error in any of the above, I would like to see it and will log it the same way.