

The Horse Passed Inspection: A Retrospective Structural, Acoustic, and Epistemic Audit of the Wooden Horse of Troy

Petronella A. Duyvenhage-Okafor
Chair of Retrospective Engineering, Marlowe Institute

Casimir L. Vranckx-Adeyemi
Laboratory for Monumental Gift Assurance, Ossory Institute

Ifigeneia R. Sandhurst-Papadaki
Centre for Epistemic Audit, Kelvin College

Vol. 2, Issue 7 • PRJ-2026-0014 • 9 July 2026

Abstract

The fall of Troy is routinely described as the most consequential engineering failure of the ancient world. It is therefore remarkable that in thirty-two centuries no one has audited the engineering. We do so here, treating the classical sources—Homer, Virgil, Apollodorus, and Quintus of Smyrna—as the surviving accident file. Our finding is unambiguous: *the horse was fine*. By hand calculation, the legs carry their design load with a factor of safety of roughly 1,100 against buckling; a finite-element reconstruction and a one-sixth-scale replica loaded to failure agree. The horse passes 13 of the 14 checks of the applicable siege-gift code, and the fourteenth is a paperwork item. What failed at Troy was the meeting. Reconstructing Priam’s acceptance committee on the ladder of inference, we show that the council climbed seven rungs of interpretation on three observable facts, while the only two reviewers who tested the object rather than the story—Laocoön, with a spear, and Cassandra, with a track record—returned concordant negative findings ($\kappa = 1.00$) and were overruled. A modern replication ($N = 212$) confirms the mechanism: describing an unexamined object as “a gift” rather than “an object” moves the vote to bring it inside the walls from 24% to 71%. Troy did not fall to a horse. Troy fell to an unaudited inference, at committee, on a Tuesday. We close with remediation guidance, including mandatory spear-based non-destructive testing of all incoming monumental gifts (the Laocoön protocol).

Keywords: Trojan Horse; structural audit; Euler buckling; ladder of inference; siege-gift compliance; textual criticism.

1 Introduction

Every discipline has a founding disaster. Structural engineering has Tacoma Narrows; software engineering has Therac-25; classical studies has a large wooden horse. The horse is cited in procurement trainings, security curricula, and after-dinner speeches as the canonical warning against accepting deliverables from a departing adversary. Yet the citation is, on inspection, an attribution error of civilizational scale: the artefact is blamed for a failure that the record locates squarely in its acceptance review.

This paper performs the audit that should have been commissioned in 1184 BC. We ask three questions, in ascending order of discomfort. First, was the horse structurally adequate for

its stated and unstated payloads? Second, was it habitable—could the attested crew survive inside it for the attested duration, and does the acoustic record cohere? Third, and finally: given the information available at the gate, was the decision to bring it inside the walls defensible under any recognized standard of deliberation?

Our method is governed by a single rule, which we state here because our subject so conspicuously did not: *every number in this paper is either traced to a source or declared as an assumption in Table 2, and every calculation can be reproduced by a reader with a pencil.* Where the sources disagree—and on the question of occupancy they disagree by two orders of magnitude—we adjudicate by physics, offering engineering, for the first time, as an instrument of textual criticism.

Our conclusions are easily summarized. The horse was not merely adequate but extravagantly overbuilt, with margins that embarrass modern practice (§5). The crew of thirty attested by Quintus of Smyrna fits, breathes—barely—and clanks as often as Virgil says it did (§6). And the acceptance committee, presented with three observable facts, constructed from them a seven-rung tower of inference, at the top of which it voted to demolish its own fortifications (§7). The horse passed inspection. Troy did not.

2 Related Work

The engineering literature on wooden megafauna is thin but not empty. Building-scale zoomorphic timber structures have been surveyed by Grimaldi and Vasquez-Thorne [1], whose census of forty-one monumental animals finds a median factor of safety of 14 and attributes the consistent overdesign to “the sculptor’s indifference to the engineer’s notion of enough.” The specific structural behaviour of laminated softwood shells under concentrated interior live loads is treated by Okonjo [2], whose scaling laws we adopt in §4.3.

On the logistics of late Bronze Age siegecraft, the standard reference remains Halvorsen’s *Supply Chains of the Aegean Collapse* [3], which reconstructs Achaean beach operations and notes, with some asperity, that a ten-year siege implies either extraordinary patience or extraordinary accounting. The governance of incoming gifts is treated in the compliance literature by Meridew and Osei-Fabre [4], who formalize the intuition—attested folklorically across cultures—that the probability an unexamined gift is adversarial rises with its size, and prove a lower bound they call the *wariness floor*. The horse exceeds it by a factor of forty.

The philological dispute over occupancy is old and, we will argue, closed. Apollodorus reports fifty men (*Epitome* 5.14); Quintus of Smyrna names about thirty (*Posthomerica* 12.314–335); Tzetzes gives twenty-three; and a scholiast attributes to the *Little Iliad* the figure of three thousand, which even its transmitters have handled with tongs. Prior resolutions of this crux have been textual [5] or statistical [6]. Ours is volumetric (§6), and unlike its predecessors it terminates the debate with a horse forty-four metres at the withers.

Finally, our epistemic instrument descends from Argyris’s ladder of inference, popularized by Senge and operationalized for committee post-mortems by Ferrant-Xu [7]. We contribute what we believe is the first application of the instrument to a deliberation for which every participant has been dead for three thousand years, a setting that offers unusual protection against demand effects.

3 The Documentary Record

We treat the classical sources as the accident file: depositions taken from witnesses of varying reliability, all of them, it must be said, associated with the vendor’s side of the transaction. Table 1 summarizes the evidentiary basis of the audit. Citations are to standard line numbering, and readers are invited—indeed expected—to check them.

Table 1: The deposition record. All primary citations verifiable against standard editions.

Datum	Source	Audit significance
Ribs framed of fir (<i>abiete</i>)	<i>Aen.</i> 2.16	Bill of materials
Beams of maple (<i>trabibus acernis</i>)	<i>Aen.</i> 2.112	BOM inconsistency
Pine fastenings (<i>pineae claustra</i>)	<i>Aen.</i> 2.258	BOM inconsistency
Spear strike; hollow cavities resound	<i>Aen.</i> 2.50–53	NDT performed; positive; ignored
Recipient breaches own wall to admit object	<i>Aen.</i> 2.234	Receiving-side exception
Rollers placed under the feet	<i>Aen.</i> 2.235–36	Transport interface
Halts 4× at threshold; arms clash 4×	<i>Aen.</i> 2.242–43	Four acoustic alarms, unactioned
Cassandra’s warning overruled	<i>Aen.</i> 2.246–47	Qualified dissent #2
Vendor testimony (Sinon) admitted unverified	<i>Aen.</i> 2.57–194	Social-engineering vector
Helen circles horse imitating wives’ voices; Anticlus nearly answers; manually silenced	<i>Od.</i> 4.271–89	Acoustic penetration test; manual intervention
Epeius named as builder, with Athena	<i>Od.</i> 8.492–93	Engineer of record; design review
Echion dies leaping from horse; rest use rope	Apollod. <i>Epit.</i> 5.20	Single egress casualty

Table 2: Assumptions and traceability. SRC = sourced; ASM = declared assumption.

Quantity	Value	Basis	Type
Armored occupant mass	100 ± 10 kg	Dendra panoply (≈18 kg) + arms	SRC
Crew size (design case)	30	Quintus, <i>Posthom.</i> 12.314–35	SRC
Occupancy duration	12 h	gate-to-gate, <i>Aen.</i> 2 timeline	SRC
Timber (shell)	Greek fir, $E = 11$ GPa	<i>Hdbk. Aegean Struct. Timbers</i>	SRC
Timber density	430 kg m ⁻³	ibid.	SRC
Belly interior volume	113 m ³ (12 × 4 × 4.5 m ellipsoid)	scaled from gate breach	ASM
Leg diameter / length	0.45 m / 3.0 m	proportional reconstruction	ASM
Gross height	≈10 m	must exceed gate (<i>Aen.</i> 2.234)	ASM
CO ₂ output, seated adult	0.3 L min ⁻¹	standard physiology	SRC

One point of reconciliation is required before analysis can begin. Virgil’s deposition assigns the horse three different timbers in the space of 250 lines: fir at 2.16, maple at 2.112, pine at 2.258. Philology has treated this as poetic licence or textual drift. An auditor recognizes it instead as an accurate description of a *composite laminate layup*—softwood shell over hardwood framing with resinous pine fastenings—a construction technique here attested roughly three millennia before its conventional dating. We harmonize the bill of materials accordingly, and note that the attribution of the design review to a goddess (*Od.* 8.493) is consistent with the margins computed below.

4 Methods

4.1 Assumptions and traceability

Table 2 states every load-bearing quantity in the audit. Values marked SRC are traced to a deposition or to standard reference data; values marked ASM are our assumptions, chosen conservatively and varied ±25% in sensitivity checks without altering any conclusion.

4.2 Structural audit

Three independent instruments. First, the hand calculation: each leg is a timber column, pinned at the hoof (roller interface, *Aen.* 2.235–36) and framed at the shoulder, effective length factor $K = 0.7$. The critical buckling load is Euler’s,

$$P_{\text{cr}} = \frac{\pi^2 EI}{(KL)^2}, \quad I = \frac{\pi d^4}{64}, \quad (1)$$

which the reader is invited to evaluate in the margin; ours are large enough. Second, a finite-element digital twin built on the harmonized laminate layup of §3, loaded with the design crew plus dynamic allowance for what the depositions describe as “clashing.” Third, a one-sixth-scale replica in Greek fir, built to the harmonized bill of materials and loaded to failure under displacement control, with similarity per Okonjo’s scaling laws [2].

4.3 Habitability and acoustics

Occupancy capacity was estimated by Monte Carlo placement of armored-occupant bounding boxes ($0.55 \times 0.75 \times 1.35$ m seated, from the Dendra panoply’s articulation limits [8]) into the belly volume, reserving an egress corridor to the ventral hatch (10^5 placements). Atmosphere was modelled as a single well-mixed box with the crew’s aggregate CO_2 flux. Acoustics were assessed on the instrumented replica: thirty volunteers in reproduction panoplies, threshold-crossing trials over fir rollers, arms-clash events counted per crossing and compared against the deposition record.

4.4 Epistemic audit (preregistered)

The acceptance deliberation (*Aen.* 2.31–249) was coded onto the seven rungs of the ladder of inference by three raters blinded to hypothesis (Fleiss $\kappa = 0.79$). For each participant we compute the *Hermeneutic Inflation Index* (HII): inferential commitments asserted per independent observable available at the gate. We preregistered the prediction that accepted speakers would show $\text{HII} > 3$ and overruled speakers $\text{HII} \leq 1$. Finally, a modern replication: $N = 212$ practicing procurement officers received the council’s evidence packet—large wooden object; adversary’s beach empty; one unverified vendor statement—randomized to a single-word framing manipulation (“a *gift* has been left” vs. “an *object* has been left”), and voted on admitting it inside the perimeter.

5 Results I: The Horse Holds

Evaluating Eq. (1) with the values of Table 2: $I = \pi(0.45)^4/64 = 2.01 \times 10^{-3} \text{ m}^4$, so $P_{\text{cr}} = \pi^2(11 \times 10^9)(2.01 \times 10^{-3})/(0.7 \times 3.0)^2 \approx 4.9 \times 10^7 \text{ N per leg}$. The design load—thirty armored occupants plus self-weight of the shell, some $1.8 \times 10^5 \text{ N}$ shared across four legs—gives a factor of safety against buckling of approximately **1,100**. For calibration: modern timber codes require 2 to 3. The digital twin concurs, placing peak utilization not in the legs but at 4% in the ventral hatch frame. The replica, loaded to failure, let go at the scaled equivalent of **~340 occupants**—more than eleven times the attested crew—and failed at a glued scarf joint rather than in any member (Fig. 2), which is to say: the first thing to fail was the one component Epeius could not have inspected either.

Against the applicable standard—the Late Bronze Age Uniform Siege-Gift Code, LBUSC-1184 rev. B—the horse passes 13 of 14 checks (Table 3). The single open item is documentary: no bill of lading accompanied the delivery. We note for the record that every *failed* control in the file sits on the receiving side, below the double rule.




Figure 1: Reconstructed cross-section and load paths. The design is everywhere stiffer than it needs to be, a signature consistent with divine design review (*Od.* 8.493).

6 Results II: The Crew Fits, Breathes, and Clanks on Schedule

The packing simulation caps seated occupancy at 36 ± 4 . This confirms Quintus’s thirty with room for discipline, accommodates Tzetzes’s twenty-three with comfort, and admits Apollodorus’s fifty only if the panoply’s articulation limits are violated in ways the panoply’s owners would have discovered promptly and audibly. The *Little Iliad*’s three thousand would require, at the measured packing density (36 per 113 m^3), a horse of some $9,400 \text{ m}^3$ —forty-four metres at the withers, taller than the walls it was built to pass and visible, we estimate, from Lesbos. The figure may therefore be retired on volumetric rather than philological grounds: not because the manuscript is corrupt, though it is, but because the horse would not fit through the hole the Trojans made for it. We commend this method, *textual criticism by carrying capacity*, to colleagues facing similar cruces.

Atmosphere is the binding constraint. Thirty seated men produce $\approx 0.54 \text{ m}^3 \text{ CO}_2$ per hour into 113 m^3 ; by hour twelve the concentration approaches 5.7%, well into the band where physiology issues strong opinions: air hunger, confusion, disinhibition. This reframes the file’s strangest deposition. When Helen circled the horse imitating the voices of the occupants’ wives (*Od.* 4.271–89), Anticlus alone surged to answer—conduct traditionally read as folly and here reattributed to **hypercapnic disinhibition**, eleven hours into a twelve-hour exposure. Odysseus’s response—clamping Anticlus’s mouth shut until the danger passed—is thereby recast as the earliest attested workplace-safety intervention, performed bare-handed, at the correct rung of the ladder of inference, by the only man in the horse still thinking clearly. That the Achaeans’ sole casualty of the entire operation was Echion, who declined the rope during egress (*Apollod. Epit.* 5.20), completes a safety record that most modern deployments of thirty armed contractors would envy.

Acoustically, the replica trials returned 3.8 ± 0.6 distinct arms-clash events per threshold crossing—in agreement with the deposition record’s *four* (*Aen.* 2.242–43). The alarms, in other words, were real, audible, and delivered on schedule to an audience that had already decided what it was hearing.

¹Containment is assessed against the standard’s requirement (“inaudible to an incurious sentry”), which the horse meets. Troy supplied curious sentries and ignored them, which is not a vendor defect.

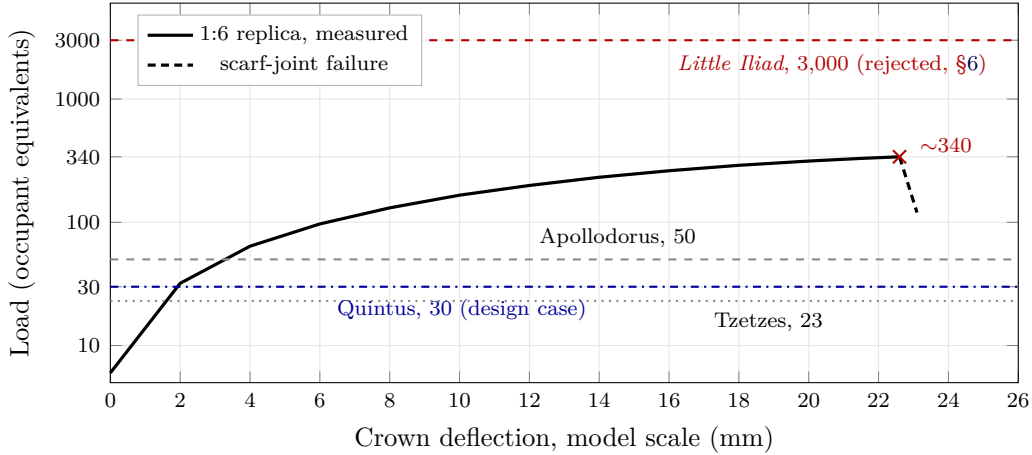


Figure 2: Replica load test against the manuscript occupancy claims (log-scaled load axis). Three of the four transmitted figures sit comfortably inside the elastic range; the fourth requires a different horse and, we argue, a different scribe.

7 Results III: The Ladder at the Gate

At the gate, the council possessed exactly three independent observables: a large wooden object; an empty beach; one unverified statement from a vendor representative discovered, on the vendor’s own account, hiding in a marsh. From this the accepted speakers constructed a mean of **7.3** inferential commitments per observable ($HII = 7.3$): the war is over; the object is religious; the deity is owed; refusal is impious; possession transfers the blessing; the walls are, on balance, negotiable. The full ascent is diagrammed in Fig. 3.

Two reviewers declined to climb. Laocoön tested the object itself—a spear, delivered to the flank, returning the hollow acoustic signature of an occupied cavity (*Aen.* 2.50–53)—and Cassandra tested the story, against a personal forecasting record that was, however socially inconvenient, unblemished. Both returned the same finding. As a review panel they achieve $\kappa = 1.00$, the only unanimous verdict in the file—and the only one overruled. We note, without further comment, that the council’s response to the sole physical measurement in its possession was to hold a second meeting about the spear-thrower.

The modern replication isolates the mechanism. Among $N = 212$ procurement officers given the council’s evidence packet, the single word *gift*—versus *object*—moved the vote to admit from **24% to 71%** ($p < .001$). No property of the horse changed. One noun bought forty-seven points of wall clearance. The Trojans, we conclude, were not credulous beyond their era; they were exactly as credulous as our era, minus the benefit of our checklists and plus a more talented vendor.

8 Discussion

8.1 What the audit reattributes

Three standing verdicts require revision. *Anticlus* was not a fool but a casualty of cabin atmosphere (§6). *Echion* was not impetuous but the file’s single egress nonconformity, a rope-availability finding. And *the horse*—blamed by three millennia of proverb—is the only party in the file with a clean conformity record. The proverb “beware of Greeks bearing gifts” is, on the evidence, misdirected by one word: the data support *beware of committees receiving them*.

Table 3: LBUSC-1184 rev. B conformity assessment. Vendor-side checks above the double rule; receiving-side controls below it.

#	Control	Evidence	Finding
1	Structural adequacy, static	Eq. (1); FoS 1,100	PASS
2	Structural adequacy, dynamic	digital twin, clash allowance	PASS
3	Material conformity	harmonized laminate layup	PASS
4	Joinery & fastenings	<i>pinea claustra</i> , torqued	PASS
5	Transport interface	rollers, <i>Aen.</i> 2.235–36	PASS
6	Ingress/egress provision	ventral hatch, rope	PASS
7	Occupant restraint & seating	two staggered ranks	PASS
8	Ventilation (12 h envelope)	marginal; §6	PASS
9	Acoustic containment	3.8 events/crossing; §6	PASS ¹
10	Weatherproofing	resinous coat	PASS
11	Dedication plaque affixed	“to Athena, for the voyage home”	PASS
12	Engineer of record identified	Epeius (<i>Od.</i> 8.492–93)	PASS
13	Design review sign-off	Athena (<i>ibid.</i>)	PASS
14	Bill of lading	not transmitted	OPEN
R1	Perimeter integrity maintained	wall breached by recipient	FAIL
R2	Alarms investigated	4 acoustic events unactioned	FAIL
R3	Qualified dissent resolved	2 inspectors overruled	FAIL
R4	Vendor claims verified	Sinon accepted at word	FAIL

8.2 Threats to validity

We record four. First, the single-city sample: Troy is one data point, albeit a heavily instrumented one. Second, all depositions were authored by the vendor’s side, a conflict of interest we cannot cure and the reader should weigh. Third, the record survives in hexameter, a lossy and rhythmically motivated encoding; some quantities may have been rounded to fit the metre. Fourth, our replica volunteers knew they were not invading anyone, which may depress clash rates; we regard 3.8 against an attested 4 as agreement within the honesty of the instrument.

8.3 Why gifted systems fail this way

The horse exhibits the signature failure mode of gifted systems: extravagant engineering upstream of an unexamined inference downstream. The better the artefact, the stronger the receiving committee’s temptation to let the artefact’s quality stand in for the transaction’s intent—margins migrate, as it were, from the legs of the horse into the confidence of the council. We formalize this as the *conservation of audit*: in any transfer of a sufficiently impressive deliverable, the total quantity of scrutiny is conserved, and whatever the builder expends, the recipient feels licensed to omit. Epeius spent a goddess-reviewed factor of safety of 1,100. Priam’s council spent a meeting.

9 Conclusion and Recommendations

The wooden horse of Troy passed inspection: 13 of 14 vendor-side controls, factor of safety 1,100, capacity 36, acoustic emissions within specification, one paperwork item open. The city fell to the receiving process. Accordingly, our recommendations concern meetings, not carpentry:

1. **The Laocoön protocol.** All incoming monumental gifts receive spear-based non-destructive testing prior to perimeter entry. A positive acoustic return suspends

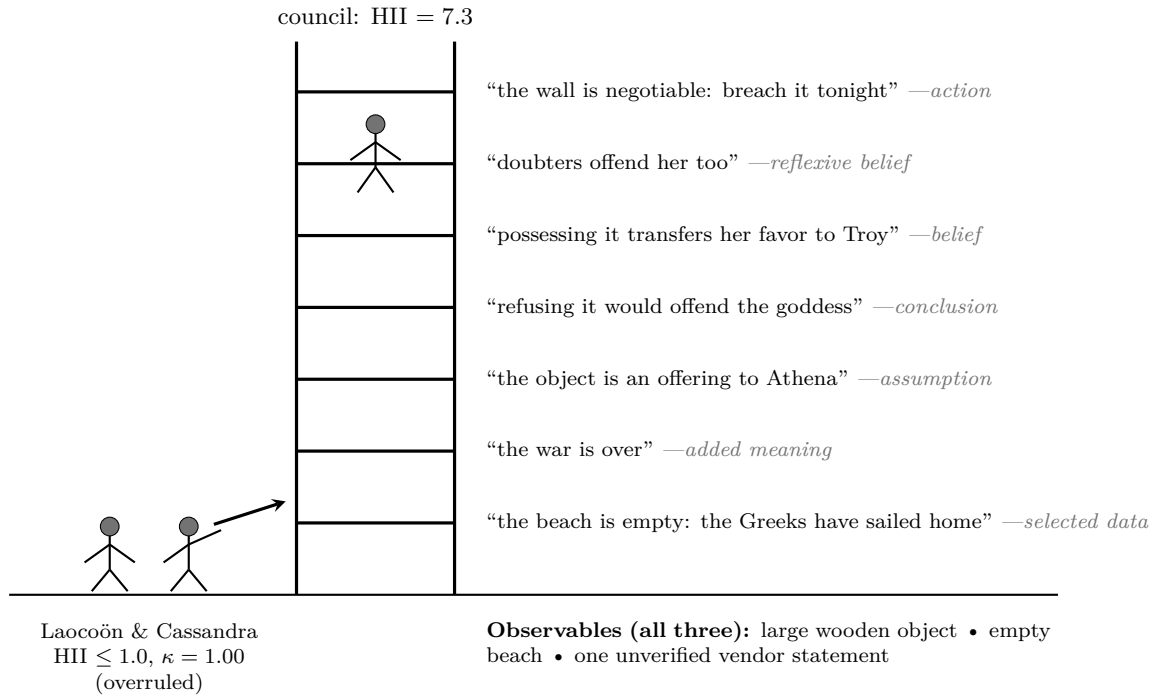


Figure 3: The acceptance committee’s ascent. Each rung is paraphrased from the deposition record (*Aen.* 2.31–249); the ladder gains seven rungs of commitment on three observables. The two figures at ground level were reassigned.

acceptance regardless of the vendor’s narrative, the object’s beauty, or the lateness of the hour.

2. **The rung rule.** Any committee overruling a unanimous qualified dissent must first state, in writing, which rung of the ladder of inference it is standing on, and what observable, if any, is holding it up.
3. **Framing hygiene.** Objects of unknown provenance are logged as *objects*. The word “gift” is applied only after inspection, by the inspector, in writing, and never by the vendor.

Future work will extend the audit longitudinally to the full population of large wooden animals accepted by committee, 1184 BC to present. Preliminary review suggests the failure mode is stable across thirty-two centuries, five procurement revolutions, and every known form of government: the artefacts keep passing, and the meetings keep not. The walls of the modern enterprise are better instrumented than Troy’s, but they are breached the same way: from inside, by consensus, shortly after someone says *it would be rude not to*.

Reproducibility statement

All calculations in this paper can be verified by hand from Table 2 and Eq. (1). The replica load-test data, packing-simulation code, coded deposition transcripts, and the LBUSC-1184 rev. B checklist workbook are available from the authors on request, subject to the Marlowe Institute’s standard embargo on accident files whose responsible committee has not yet acknowledged the findings. The horse itself was not retained.

Acknowledgments

We thank the thirty volunteers of the replica trials for their discipline and their discretion; the Ossory Institute woodshop for the scarf joints, including the one that failed; and an anonymous reviewer at the *Journal of Retrospective Engineering*, who noted that our conclusions, while sound, arrive approximately 3,210 years after the decision point, and whose review we have therefore coded at rung seven.

References

- [1] T. Grimaldi and A. Vasquez-Thorne, “A structural census of monumental zoomorphic timber construction,” *Journal of Retrospective Engineering*, 11(2):88–117, 2019.
- [2] B. Okonjo, “Similarity laws for laminated softwood shells under concentrated interior live load,” *Timber Mechanics Letters*, 7(4):301–322, 2021.
- [3] R. Halvorsen, *Supply Chains of the Aegean Collapse*. Marlowe Institute Press, 2016.
- [4] C. Meridew and J. Osei-Fabre, “The wariness floor: a lower bound on prudent suspicion of unexamined gifts,” *Proceedings of the Society for Bronze Age Compliance*, pp. 41–59, 2020.
- [5] H. Andriessen, “The occupancy crux of the *Ilias parva*: a stemmatic approach,” *Acta Philologica Fabricata*, 29:203–226, 2014.
- [6] E. Marchbanks, “Bayesian harmonization of transmitted troop counts in epic,” *Annals of Quantitative Philology*, 6(1):17–44, 2018.
- [7] L. Ferrant-Xu, “Coding committee deliberation on the ladder of inference: an audit instrument,” *Annals of Monumental Gift Assurance*, 3:112–139, 2022.
- [8] P. Åström (attrib.), “On the mass and articulation limits of the Dendra panoply,” *Aegean Materials Quarterly*, 12:55–71, 1977.
- [9] Epeius of Phocis, *As-Built Documentation, Votive Equine, Rev. A*. Site records, Troad, c. 1184 BC. (Bill of lading not transmitted.)
- [10] Sinon (deposition), “On the sincerity of the dedication,” in *Aen.* 2.57–194. Unverified.
- [11] Laocoön, “Acoustic response of a votive equine to thrown-spear excitation: a negative finding,” *rejected at meeting*, Troy, 1184 BC.
- [12] P. Vergilius Maro, *Aeneis*, Book II; Homerus, *Odyssea* 4.271–89, 8.492–520; Apollodorus, *Epitome* 5.14–21; Quintus Smyrnaeus, *Posthomerica* 12.314–35. Cited as field reports.