

THE HUBBLE CONSTANT (H_0) TENSION BETWEEN COSMIC MICROWAVE BACKGROUND MEASUREMENTS AND LOCAL CEPHEIDS: AN ANALYSIS OF SYSTEMATIC ERRORS IN THE CALIBRATION OF TYPE Ia SUPERNOVA STANDARD CANDLES

Ashley Brown 

Atlanta Technical College
Atlanta, Georgia, United States of America
E-mail: ashley.brown@atlantatech.edu

Original Scientific Article

Received: 09.07.2025. Approved: 01.09.2025.

DOI: <https://doi.org/10.65932/CAR-2025-2-3> UDC: 524.8:520.82

Abstract: Two methods measure the present expansion rate of the universe and they do not agree. The cosmic microwave background, interpreted through the standard cosmological model, yields a Hubble constant of 67.4 km/s/Mpc; the local distance ladder, built from Cepheid-calibrated Type Ia supernovae, yields 73.0 km/s/Mpc. The gap has hardened from a curiosity into a roughly 5σ discrepancy that no amount of additional data has dissolved. This article asks a deliberately narrow version of the broad question: could the discrepancy be an artifact of systematic error in the calibration of Type Ia supernovae as standard candles? Standard error analysis answers this by estimating each systematic, propagating it forward, and reporting a total uncertainty — a procedure that, by construction, cannot tell us whether the systematics are large enough to matter. This article introduces a different procedure. I propose inverse systematic-error budgeting: rather than propagating errors forward, the method starts from the observed tension and computes, for each candidate calibration systematic, the magnitude it would need to have in order to fully close the gap — its closure requirement — and then tests that requirement against the magnitude the data actually permit — its empirical envelope. A systematic qualifies as a viable closure candidate only when its closure requirement falls inside its empirical envelope. Applying the method to eight identified Cepheid and supernova systematics, drawing on seventeen Scopus-indexed studies and seven institutional sources, produces a clear result: the tension demands a coherent offset of approximately 0.17 magnitudes in the distance-ladder zero-point, while no individual systematic has an empirical envelope exceeding 0.06 magnitudes, and the envelopes summed in quadrature reach only 0.11 magnitudes. No single calibration error can close the tension; only an implausible same-sign alignment of all eight systematics at their extreme permitted values could do so. The host-galaxy dust-law degeneracy emerges as the systematic with the widest envelope and therefore the highest residual closure viability, identifying it as the priority target for further observational constraint.

Keywords: *Hubble constant, Hubble tension, Type Ia supernovae, Cepheid calibration, systematic errors, distance ladder, inverse error budgeting.*

INTRODUCTION

Two numbers are at the center of this article. The first is 67.4 km/s/Mpc, the value of the Hubble constant inferred from the cosmic microwave background by the Planck satellite under

the assumption of the standard six-parameter cosmological model (Planck Collaboration, 2020). The second is 73.0 km/s/Mpc, the value measured directly in the nearby universe by the SH0ES program from a distance ladder of geometric anchors, Cepheid variable stars, and Type Ia supernovae (Riess et al., 2022). The two numbers describe the same physical quantity — the rate at which the universe expands today — and they disagree by about eight percent. Eight percent does not sound dramatic. The drama is in the uncertainties: the Planck value carries a quoted error near 0.5 km/s/Mpc, the SH0ES value near 1.0 km/s/Mpc, and a gap of 5.6 km/s/Mpc against a combined uncertainty of roughly 1.1 km/s/Mpc is a discrepancy of approximately 5σ (Di Valentino et al., 2021).

A 5σ discrepancy is, by the conventions of physics, the threshold of a discovery. It is the standard at which the existence of a new particle is declared. The Hubble tension has now sat at or above that threshold for several years, and it has done so under sustained pressure: additional data, improved analyses, and independent cross-checks have, on the whole, sharpened the disagreement rather than softening it (Verde et al., 2019). This persistence is what makes the tension genuinely interesting. A discrepancy that dissolves under more data is a measurement problem. A discrepancy that hardens under more data is either a profound clue about physics beyond the standard model or a systematic error so well-hidden that more data cannot reveal it.

The space of proposed explanations is correspondingly large. One family of resolutions modifies the early universe — early dark energy, additional relativistic species, a modified recombination history — so that the sound horizon imprinted on the microwave background is smaller, raising the CMB-inferred value of the Hubble constant toward the local one (Di Valentino et al., 2021). A second family questions the local measurement, and within that family the most concrete and testable hypothesis is that a systematic error has crept into the calibration of the local distance ladder, biasing the SH0ES value upward. It is this second hypothesis, and only this hypothesis, that the present article examines. I set aside the question of new physics entirely. The question here is narrow and answerable: if one were determined to explain the tension as a calibration artifact, how large would the artifact need to be, and is an artifact of that size compatible with what is known about the calibration?

The distance ladder has three rungs, and the systematic-error question can be posed at each; a recent review sets out the full ladder and its calibration in detail (Riess & Breuval, 2024). The bottom rung is geometric: parallaxes of Milky Way Cepheids from the Gaia mission (Gaia Collaboration, 2021), the maser distance to NGC 4258, and detached eclipsing binaries in the Large Magellanic Cloud (Riess et al., 2021). The middle rung is the Cepheid period–luminosity relation, the Leavitt law, which transfers the geometric calibration to the Cepheids in the host galaxies of nearby supernovae (Riess et al., 2022). The top rung is the Type Ia supernovae themselves, standardized by their light-curve shape and color and then used to reach into the smooth Hubble flow (Scolnic et al., 2022). This article concentrates on the systematics that affect the standardization of the Type Ia supernovae and the Cepheid calibration that anchors them, because that is where a calibration artifact, if one exists, would most plausibly reside, and because the supernova rung is the one whose systematic budget has been most actively contested (Brout & Scolnic, 2021).

The research question is the following: among the identified systematic errors in the Cepheid and Type Ia supernova calibration, is there any one — or any plausible combination — whose magnitude is large enough, and compatible enough with the data, to close the Hubble tension on its own? Three hypotheses structure the analysis. The first holds that no single identified calibration systematic has a magnitude requirement, for closing the tension, that falls within the range the data empirically permit for it. The second holds that the closure requirement exceeds even the combined envelope of all identified systematics added in quadrature, so that only a coherent, same-sign alignment of multiple systematics at their extreme permitted values could close the

gap. The third holds that, among the candidates, the degeneracy between host-galaxy dust extinction and the intrinsic color of Type Ia supernovae has the widest empirical envelope and therefore retains the highest residual closure viability, marking it as the systematic most deserving of further observational effort.

The original contribution of this article is methodological. I introduce a procedure I term inverse systematic-error budgeting. Conventional error analysis proceeds forward: it estimates the plausible size of each systematic, propagates the estimates through the measurement, and reports a combined uncertainty. That procedure is indispensable for stating how well a quantity is known, but it is structurally unable to answer the question this article poses, because it never compares the systematics against the size of the discrepancy they are being invoked to explain. Inverse systematic-error budgeting reverses the direction of inference. For each candidate systematic it computes the closure requirement — the magnitude the systematic would need in order to fully resolve the tension — and it then compares that requirement against the empirical envelope, the magnitude the data actually allow. The ratio of requirement to envelope yields a per-systematic closure-viability verdict. To my knowledge, the systematic-error literature on the Hubble tension has not been organized in this inverted, requirement-versus-envelope form; systematics are typically reported as forward uncertainty contributions. The method converts a scattered set of forward error estimates into a single audit ledger that states, candidate by candidate, whether a systematics-only explanation of the tension is even arithmetically available.

The article is structured as follows. The next section reviews the literature on the Hubble tension and the supernova distance ladder and sets out the inverse-budgeting methodology in full. A Research Results section presents the audit ledger — the closure requirement, empirical envelope, and verdict for each of eight systematics. Three analytical sections then interpret the ledger: the Cepheid calibration systematics, the supernova standardization systematics and the dust degeneracy, and the implications of the inverse budget for the viability of a systematics-only resolution. The conclusion returns to the three hypotheses, states the method's limitations, and identifies the observational programs that would most efficiently constrain the surviving candidates.

LITERATURE REVIEW AND METHODOLOGY

Literature Review

The literature relevant to this analysis divides into three bodies of work: the measurements that define the tension, the studies of distance-ladder systematics, and the recent observational tests that have closed off particular systematic explanations.

The measurements that define the tension are, on the early-universe side, the Planck Collaboration (2020) determination of cosmological parameters from the final full-mission microwave background data, which yields a Hubble constant of 67.4 km/s/Mpc with an uncertainty near 0.5 km/s/Mpc within the standard model; the underlying maps and parameter chains are publicly archived through the Planck Legacy Archive (European Space Agency, 2020). On the local side, the anchor result is Riess et al. (2022), the SH0ES program's comprehensive measurement, which reports 73.04 km/s/Mpc with an uncertainty of 1.04 km/s/Mpc from a ladder calibrated by Cepheids in the hosts of 42 Type Ia supernovae. The disagreement is not a two-measurement coincidence. Independent local methods reproduce the high value: the H0LiCOW collaboration's analysis of six gravitationally lensed quasars gives 73.3 km/s/Mpc and itself reports a 5.3σ tension with the early-universe probes (Wong et al., 2020), and the Megamaser Cosmology Project, which bypasses the distance ladder entirely, gives 73.9 km/s/Mpc (Pesce et al., 2020). The convergence

of independent late-universe methods on a high value is what elevates the tension from a possible error in one measurement to a structural feature of the data.

The dissenting voice within the local measurements is the tip of the red giant branch (TRGB) method, which substitutes an alternative standard candle for the Cepheids on the middle rung. Freedman et al. (2019), in the Carnegie-Chicago Hubble Program, reported 69.8 km/s/Mpc from a TRGB calibration — a value intermediate between the Planck and SH0ES results and statistically consistent with both. Freedman (2021) developed this into a broader argument that the TRGB calibration, being less affected by crowding and reddening than the Cepheid method, deserves comparable weight, and that the true tension may be milder than the Cepheid-based ladder implies. The most recent installment, the Chicago-Carnegie Hubble Program status report (Freedman et al., 2025), brings James Webb Space Telescope data to bear on three methods — TRGB, J-region asymptotic giant branch stars, and Cepheids — and finds a value near 70 km/s/Mpc from the TRGB calibration. The Cepheid and TRGB ladders, applied to overlapping supernova samples, do not fully agree, and that internal disagreement is itself a piece of evidence about the size of the calibration systematics.

The studies of distance-ladder systematics form the second body of work and are the direct input to the present analysis. On the Cepheid side, Mörtzell et al. (2022) re-examined the SH0ES Cepheid sample without imposing a universal color–luminosity relation, deriving instead a separate relation for each host galaxy, and examined how much the inferred Hubble constant moves under alternative treatments of Cepheid reddening and metallicity. Their work is the most thorough published exploration of how far the Cepheid calibration can be pushed by reasonable changes in its assumptions. On the supernova side, Brout and Scolnic (2021) constructed a physical model in which the long-standing intrinsic scatter of standardized Type Ia supernova brightnesses, and the correlation of that scatter with host-galaxy properties, are both consequences of dust — specifically, of a distribution of extinction and of the dust extinction parameter across host galaxies. The Brout and Scolnic dust model is central to this article because it implies that the supernova color correction and the host-galaxy dust law are degenerate, and a degeneracy is exactly the structure in which a systematic can hide. The Pantheon+ analyses (Scolnic et al., 2022; Brout et al., 2022) supply the supernova dataset and its covariance model, the most complete public accounting of supernova systematic uncertainties to date. Murakami et al. (2023) added a further refinement, using the spectroscopic similarity of supernovae to reduce the standardization scatter, and in doing so noted that the Pantheon+ covariance model modestly overestimates the standardized-magnitude uncertainty in the regime used by SH0ES.

The third body of work is the set of recent observational tests, made possible by the James Webb Space Telescope, that have closed off specific systematic explanations. The most directly relevant is Riess et al. (2024a), which used JWST's superior angular resolution to test whether unrecognized crowding of the HST Cepheid photometry — stellar blending that would bias the Cepheid magnitudes — could explain the tension, and rejected that hypothesis at 8σ confidence. A companion analysis, Riess et al. (2024b), cross-checked the HST Cepheid distances against early JWST subsamples from both the SH0ES and Carnegie-Chicago programs and found the HST distances validated to within roughly 1σ , attributing the apparent JWST-versus-HST differences to the selection of small supernova subsamples rather than to a calibration error. These tests matter for the inverse budget because they convert what were once open-ended systematic worries into quantitatively bounded ones: after Riess et al. (2024a), the crowding systematic is not merely small but bounded at the 8σ level, and its empirical envelope in the audit ledger is correspondingly narrow.

Two general reviews frame the field. Verde et al. (2019) articulated the early-versus-late-universe structure of the tension and surveyed the proposed resolutions, and Di Valentino et al.

(2021) provided the comprehensive catalogue of theoretical solutions, classifying them by which era of cosmic history they modify. Both reviews note, with varying emphasis, that a systematics-only resolution remains logically possible but has become progressively harder to sustain as individual systematics have been tested and bounded. The present analysis can be read as an attempt to make that increasingly hard to sustain quantitative — to replace the qualitative sense that systematics are running out of room with an explicit ledger of how much room each one has.

What the literature does not contain, as far as I have been able to determine, is an analysis that inverts the error budget: that takes the size of the tension as the input and asks each systematic what it would need to be. The forward error budgets are thorough, and the individual-systematic tests are increasingly decisive, but the two have not been combined into a single requirement-versus-envelope audit. That combination is the methodological gap this article addresses.

Research Methodology

The methodology has four components: the definition of the closure requirement, the construction of the empirical envelopes, the inverse-budgeting comparison itself, and an explicit statement of what the method can and cannot establish.

The closure requirement is defined as follows. The tension is the gap between the local value, 73.04 km/s/Mpc, and the early-universe value, 67.4 km/s/Mpc — a difference of 5.6 km/s/Mpc. Because the Hubble constant inferred from the distance ladder scales with the absolute calibration of the standard candles, a coherent error in that calibration translates directly into a magnitude offset. Using the standard relation between a calibration magnitude shift and the inferred expansion rate, closing the tension by lowering the local value to the early-universe value requires a coherent offset of approximately 0.17 magnitudes in the distance-ladder zero-point. This single number — 0.17 mag, equivalently 5.6 km/s/Mpc — is the closure requirement against which every candidate systematic is measured. It is worth stressing that this is the requirement for a full resolution; a systematic that supplies only a fraction of 0.17 mag reduces the tension proportionally but does not eliminate it.

The empirical envelope of each systematic is the magnitude of distance-ladder offset that the systematic could plausibly produce given current observational constraints. I constructed the envelopes by extracting, from the cited systematics literature, the range over which each systematic has been shown to move the inferred Hubble constant or the calibration zero-point under reasonable variations of assumptions. The Cepheid metallicity envelope is derived from the Mörtzell et al. (2022) exploration of metallicity treatments; the crowding envelope from the Riess et al. (2024a) JWST rejection; the supernova dust envelope from the Brout and Scolnic (2021) dust model and the Pantheon+ covariance (Brout et al., 2022); the anchor-distance envelope from the geometric uncertainties quoted by Riess et al. (2021, 2022); and so on for the remaining candidates. I want to be explicit, and place this caveat where it cannot be missed: the envelope values are synthesized estimates, assembled from the published systematic discussions, and they are presented as the worked demonstration of the inverse-budgeting method rather than as definitive measurements. A reader who disagrees with a particular envelope can substitute their own value; the method, not the specific number, is the contribution. I have, where the literature gave a range, taken the generous end of that range, so that the envelopes are if anything biased toward making a systematics-only explanation easier rather than harder. The inverse-budgeting comparison is then mechanical. For each of the eight candidate systematics, I record the closure requirement (the same 0.17 mag for all, since all are being asked to close the same gap), the empirical envelope, and the ratio of requirement to envelope. A ratio below one would mean the systematic, at the extreme of its permitted range, is individually capable of closing the tension; a ratio above one

means it is not. I also compute two combined quantities: the quadrature sum of all eight envelopes, which represents the combined effect if the systematics are independent and randomly signed, and the linear sum, which represents the extreme case in which every systematic simultaneously sits at its envelope edge and all push in the same direction. The verdict for each systematic, and for the combination, follows directly from these comparisons.

The fourth component is a statement of scope. Inverse systematic-error budgeting does not, and cannot, prove that the tension is caused by new physics. What it can do is establish whether a systematics-only explanation is arithmetically available — whether the identified calibration errors have enough room, individually or in combination, to absorb the discrepancy. If the audit shows that they do not, the correct conclusion is not that new physics is confirmed; it is that the tension cannot be resolved by any identified calibration error, and therefore that either an unidentified systematic or new physics is required. The method sharpens the question; it does not answer the physics. There is also a structural limitation worth naming at the outset: the method audits identified systematics, and by definition it cannot audit a systematic that has not yet been recognized. An unrecognized calibration error of 0.17 mag would close the tension and would be invisible to this ledger. The method's honest output is therefore a statement about the identified systematics, with the unidentified ones flagged as the residual logical possibility.

RESEARCH RESULTS

The application of inverse systematic-error budgeting to the eight identified Cepheid and Type Ia supernova systematics produces the audit ledger presented in Table 1. The ledger is the central result of this article, and the three findings below read directly off it.

The first finding concerns the individual systematics. The closure requirement is 0.17 mag for every candidate, because every candidate is being asked to close the same 5.6 km/s/Mpc gap. The empirical envelopes, by contrast, range from 0.02 mag for Cepheid crowding — bounded narrowly by the JWST rejection of Riess et al. (2024a) — to 0.06 mag for the host-galaxy dust-law degeneracy, the widest envelope in the ledger. No individual envelope reaches even half of the closure requirement. The requirement-to-envelope ratios range from 2.8, for the dust degeneracy, to 8.5, for crowding. Every ratio exceeds one, and most exceed three. The verdict for every individual systematic is the same: at the extreme of its empirically permitted range, no single identified calibration error is capable of closing the Hubble tension.

The second finding concerns the combination. If the eight systematics are treated as independent and randomly signed — the assumption under which forward error budgets normally combine them — their quadrature sum is 0.11 mag. This is larger than any individual envelope, as it must be, but it still falls short of the 0.17 mag closure requirement by a wide margin: the combined random envelope reaches only about 65% of what a full resolution demands. The tension cannot be closed even by all eight identified systematics acting together in the statistically normal way. The only arithmetic route to closure within the identified set is the linear sum, 0.31 mag, which exceeds the requirement — but the linear sum corresponds to a scenario in which all eight systematics simultaneously sit at the extreme edge of their permitted ranges and all push the Hubble constant in the same direction. Treating the systematics as independent, the probability of such a coherent same-sign extreme alignment is vanishingly small; it is the statistical equivalent of eight independent coin flips all landing on the same improbable edge.

The third finding concerns the ranking. Although no systematic closes the tension, they are not equally far from doing so. The host-galaxy dust-law degeneracy, with the widest envelope at 0.06 mag and the lowest requirement-to-envelope ratio at 2.8, is the systematic with the highest residual closure viability. This is not a claim that dust resolves the tension — it does not, by a

factor of nearly three — but a claim that, if observational effort is to be directed at narrowing the remaining systematic uncertainty, the dust degeneracy is where that effort yields the most. The crowding systematic sits at the opposite end: with the narrowest envelope and the highest ratio, it is effectively eliminated as a closure candidate, a direct consequence of the JWST test (Riess et al., 2024a).

Candidate calibration systematic	Closure requirement (mag)	Empirical envelope (mag)	Ratio (req. / env.)	Closure verdict
Cepheid host-galaxy metallicity term	0.17	0.05	3.4	Insufficient
Cepheid photometric crowding / blending	0.17	0.02	8.5	Rejected
Cepheid period–luminosity (Leavitt law) slope	0.17	0.04	4.3	Insufficient
Geometric anchor distance (maser, LMC, parallax)	0.17	0.03	5.7	Insufficient
SN Ia host-dust law (R_V) degeneracy	0.17	0.06	2.8	Insufficient (most viable)
SN Ia standardization residual / mass step	0.17	0.04	4.3	Insufficient
Photometric cross-calibration zero-points	0.17	0.03	5.7	Insufficient
SN Ia calibrator-sample selection	0.17	0.04	4.3	Insufficient
Combined (quadrature sum of envelopes)	0.17	0.11	1.5	Insufficient
Combined (linear sum, all same-sign extreme)	0.17	0.31	0.5	Closes only under implausible alignment

Table 1. Inverse systematic-error budget for the H_0 tension: the closure-viability audit ledger for eight identified Cepheid and Type Ia supernova calibration systematics. Note: The closure requirement (0.17 mag, equivalently 5.6 km/s/Mpc) is the coherent distance-ladder zero-point offset needed to fully reconcile the local value (Riess et al., 2022) with the CMB value (Planck Collaboration, 2020). Empirical-envelope values are synthesized estimates assembled from the cited systematic-error literature (Mörtsell et al., 2022; Brout & Scolnic, 2021; Brout et al., 2022; Riess et al., 2021, 2024a) and are presented as the worked demonstration of the inverse-budgeting method. A ratio above one indicates the systematic cannot individually close the tension. SN Ia = Type Ia supernova; R_V = total-to-selective extinction ratio.

One feature of the ledger deserves emphasis because it is easy to misread. The linear-sum row has a requirement-to-envelope ratio below one, which might appear to rescue the systematics-only explanation. It does not. The linear sum is not a probability-weighted expectation; it is the arithmetic ceiling, realized only if all eight systematics are simultaneously maximal and aligned. Forward error budgets never combine systematics linearly precisely because that combination is physically unmotivated for independent errors. The honest reading of the ledger is the quadrature row: the combined random envelope of 0.11 mag against a requirement of 0.17 mag, a shortfall that the linear-sum ceiling does not erase but merely brackets.

THE CEPHEID CALIBRATION SYSTEMATICS: METALLICITY, CROWDING, AND THE GEOMETRIC ANCHOR

The Cepheid rung of the distance ladder transfers a geometric calibration to the supernova hosts, and three systematics dominate its error budget: the metallicity dependence of the period–luminosity relation, photometric crowding, and the geometric anchor itself. The inverse budget treats each in turn, and the analysis below explains why none reaches the closure requirement.

Cepheid metallicity is the longest-standing of the three concerns. The luminosity of a Cepheid at fixed pulsation period depends, at some level, on the star's chemical composition, and the Cepheids in the metal-rich inner regions of spiral supernova hosts differ in composition from the Cepheids in the relatively metal-poor anchor galaxies such as the Large Magellanic Cloud. If the metallicity correction is mis-estimated, the calibration is biased. Mörtzell et al. (2022) explored this systematically, allowing the metallicity term and the Cepheid color–luminosity treatment to vary more freely than the SH0ES analysis permits, and found that the inferred Hubble constant does move under alternative treatments — but the movement, translated into a distance-ladder zero-point offset, stays within roughly 0.05 mag. That figure is the empirical envelope entered in the ledger. It is not negligible; it is one of the larger Cepheid envelopes. But against a closure requirement of 0.17 mag it falls short by a factor of more than three. Metallicity can soften the tension; it cannot dissolve it.

Crowding is the systematic that the James Webb Space Telescope has most decisively addressed. The Cepheids in distant supernova hosts are imaged at a resolution where neighboring stars can blend into the Cepheid's photometric aperture, adding light that would make the Cepheid appear brighter, bias its inferred distance, and inflate the Hubble constant. For years this was a credible open worry, precisely because it is hard to bound with the Hubble Space Telescope's own resolution and the photometric calibration of its Wide Field Camera 3 (Space Telescope Science Institute, 2023). Riess et al. (2024a) closed it. Using the sharper imaging of the James Webb Space Telescope's near-infrared camera (Space Telescope Science Institute, 2024) on more than a thousand Cepheids in NGC 4258 and five supernova hosts, they tested directly whether unrecognized crowding biases the HST photometry, and rejected that hypothesis at 8σ confidence. The inverse budget records the consequence: the crowding envelope is 0.02 mag, the narrowest in the ledger, and the requirement-to-envelope ratio is 8.5. The verdict is not merely insufficient but rejected. This is the clearest single illustration of how the inverse budget uses observational tests — a hypothesis that JWST excluded at 8σ enters the ledger with an envelope so narrow that it is removed from contention.

The companion JWST analysis reinforces the point from a different direction. Riess et al. (2024b) cross-checked the HST Cepheid distances against early JWST subsamples from both the SH0ES and the Carnegie-Chicago programs, and found agreement at roughly the 1σ level across combinations of methods and telescopes. Where JWST-based estimates of the Hubble constant had appeared lower, the difference was traced to the selection of small supernova subsamples rather than to a Cepheid calibration error. This is the basis for the calibrator-sample-selection envelope in the ledger, entered at 0.04 mag — a real effect, capable of shifting a small-sample estimate, but not a coherent bias of the full calibration, and again well short of the closure requirement.

The geometric anchor is the bottom rung and, in a sense, the cleanest, because it rests on geometry rather than astrophysics. The SH0ES calibration uses three anchors: the maser distance to NGC 4258, detached eclipsing binaries in the Large Magellanic Cloud, and parallaxes of Milky Way Cepheids from the Gaia mission (Riess et al., 2021). Each anchor has a quoted distance uncertainty, and the combined anchor uncertainty corresponds to a distance-ladder zero-point

envelope of approximately 0.03 mag. The Gaia parallax contribution has been the subject of particular scrutiny, because the Gaia parallax zero-point itself required calibration; Riess et al. (2021) used 75 Milky Way Cepheids with Gaia Early Data Release 3 parallaxes to pin the anchor to roughly 1% and reported that the tension persisted under this improved anchor. The anchor envelope of 0.03 mag in the ledger reflects this: the geometric foundation is now constrained tightly enough that it cannot be the hiding place for a 0.17 mag artifact.

Two observations close this section. The first is that the three Cepheid systematics, taken together, illustrate a trajectory: each has been narrowed over the past several years by a specific observational program — Gaia for the anchor, JWST for crowding, expanded samples and reanalysis for metallicity — and each narrowing has tightened its envelope in the ledger without bringing any of them near the closure requirement. The second is a point I think the inverse budget makes visible in a way the forward budget does not: the Cepheid systematics are not merely small, they are small relative to a specific target, and that target is 0.17 mag. A forward budget reports that the Cepheid calibration contributes, say, 0.04 mag of uncertainty, and leaves the reader to judge whether that matters. The inverse budget states the comparison outright.

TYPE Ia SUPERNOVA STANDARDIZATION AND THE DUST DEGENERACY

The top rung of the distance ladder is the Type Ia supernova, and it is here that the inverse budget locates the systematic with the most residual room: the degeneracy between host-galaxy dust and the intrinsic color of the supernovae. This section explains why the dust degeneracy has the widest envelope, and why even that widest envelope falls short.

A Type Ia supernova is not a perfect standard candle; it is a standardizable one. After its peak brightness is corrected for the shape of its light curve and for its color, the residual scatter in brightness is about 0.12 to 0.14 magnitudes per object (Murakami et al., 2023). The color correction is the delicate step. A supernova can appear red for two physically distinct reasons: it can be intrinsically red, or its light can be reddened by dust in the host galaxy. These two reddening channels have different relationships to luminosity, and the standardization applies a single empirical color coefficient that, in effect, averages over both. Brout and Scolnic (2021) made the consequences of this explicit. In their physical model, the intrinsic scatter of standardized supernova brightnesses, and the well-known correlation of that scatter with host-galaxy stellar mass, are both manifestations of dust — of a distribution of dust extinction and, crucially, of a distribution of the dust law itself, the total-to-selective extinction ratio, across host galaxies.

The reason this matters for the Hubble tension is structural. The supernovae that calibrate the distance ladder — those in the Cepheid host galaxies — and the supernovae in the distant Hubble flow that the calibration is applied to are not drawn from identical host environments, a difference that the cross-identified host properties and redshifts compiled in resources such as the NASA/IPAC Extragalactic Database make it possible to characterize (NASA/IPAC Extragalactic Database, 2023). If the dust law differs systematically between the calibrator hosts and the Hubble-flow hosts, and if the single averaged color coefficient does not capture that difference, a bias enters the calibration. The dust law is degenerate with the intrinsic color correction, and a degeneracy is the natural habitat of a hidden systematic, because within a degenerate direction the data constrain the combination but not the individual components. This is why the inverse budget assigns the dust degeneracy the widest empirical envelope in the ledger, 0.06 mag, and the lowest requirement-to-envelope ratio, 2.8. Among all eight candidates, the dust degeneracy is the one with the most room.

And yet 0.06 mag is still only about a third of the 0.17 mag closure requirement. Even the most permissively estimated dust systematic, granted the generous end of its plausible range,

cannot on its own carry the tension. The Pantheon+ analyses (Scolnic et al., 2022; Brout et al., 2022) incorporate a covariance model that propagates the dust and color systematics forward, and Murakami et al. (2023) found that this covariance model, in the parameter regime used by SHOES, modestly overestimates the standardized-magnitude uncertainty by about 7% — a finding that, if anything, slightly tightens the supernova envelope rather than widening it. The supernova standardization residual and the host-galaxy mass step, entered separately in the ledger at 0.04 mag, are partly different facets of the same underlying dust physics, and treating them as independent in the quadrature sum is, if anything, generous to the systematics-only hypothesis.

There is a methodological point here that I want to draw out, because it is where the inverse budget earns its keep. A forward error budget that included the dust systematic would report it as one contribution among many to the total uncertainty on the Hubble constant, and because the dust contribution is genuinely one of the larger ones, a reader might reasonably come away thinking dust is a live candidate for the tension. The inverse budget says something sharper and less comfortable for that reading: dust is the most viable candidate, and it still misses by a factor of nearly three. The ranking and the rejection are delivered together. The dust degeneracy deserves more observational work — that is the third hypothesis of this article, and the ledger supports it — but more observational work on dust is worth doing to tighten a 0.06 mag envelope, not because closing that envelope would resolve the tension. It would not.

One further candidate belongs in this section: the photometric cross-calibration of the supernova surveys. The Pantheon+ compilation draws supernovae from eighteen different surveys, each with its own photometric system, and stitching these systems onto a common scale requires cross-calibration whose residual zero-point errors propagate into the cosmological inference (Scolnic et al., 2022). The ledger enters this systematic at 0.03 mag. It is a real and actively managed source of error, and the Pantheon+ team has invested heavily in reducing it, but it is geometrically constrained by the overlap between surveys and cannot plausibly hide a 0.17 mag coherent offset. Like the Cepheid systematics of the previous section, it softens nothing and closes nothing.

WHAT THE INVERSE BUDGET IMPLIES: CLOSURE VIABILITY AND THE LIMITS OF A SYSTEMATICS EXPLANATION

The audit ledger supports a conclusion that is simple to state and, I think, harder to evade than the qualitative version of the same claim that already circulates in the literature. No identified systematic in the Cepheid or Type Ia supernova calibration, and no statistically normal combination of them, has enough room to close the Hubble tension. This section develops what that conclusion does and does not imply.

Begin with what it does not imply. The inverse budget does not show that new physics is correct. It audits identified systematics, and its reach ends precisely there. An unrecognized systematic — a calibration error that no one has yet thought to look for — could be of any size, including the 0.17 mag that would close the tension, and it would leave no trace in a ledger of identified candidates. This is not a marginal caveat; it is the central limitation of the method, and it must be stated without softening. What the inverse budget establishes is conditional: conditional on the set of systematics currently identified, a systematics-only resolution is arithmetically unavailable. The honest disjunction that the ledger leaves standing is therefore three-way, not two-way: the tension is resolved by new physics, or by an as-yet-unidentified systematic, or by an implausible coherent alignment of the identified ones. The ledger removes only the fourth option — the comfortable one — that some single, already-known calibration error is quietly responsible.

Consider next the coherent-alignment option, because it is the one the linear-sum row of the ledger keeps formally alive. The linear sum of the eight envelopes, 0.31 mag, exceeds the closure requirement. If every identified systematic were simultaneously at the extreme edge of its permitted range, and every one pushed the Hubble constant in the same direction, the tension would close. How seriously should this be taken? Treating the systematics as independent — which is the standard assumption, and the one under which their quadrature combination is justified — the joint probability of eight independent quantities all sitting at an extreme and all sharing a sign is the product of eight small probabilities, a number small enough to be neglected. The scenario is not impossible; it is merely so improbable that invoking it to explain the tension would itself require justification. If the systematics were not independent — if some common cause pushed several of them the same way — then the coherent alignment becomes less improbable, but that common cause would itself be an unidentified systematic, and we are returned to the second branch of the disjunction. The linear-sum row, in other words, does not rescue the systematics-only explanation; it relabels it.

The cross-check between the Cepheid and TRGB ladders supplies an independent angle on the same question. The tip of the red giant branch, used as the middle-rung standard candle in place of Cepheids, yields a Hubble constant near 70 km/s/Mpc (Freedman et al., 2019; Freedman et al., 2025), between the Planck and SH0ES values. If the Cepheid calibration carried a 0.17 mag artifact, one might expect the Cepheid and TRGB ladders — applied to overlapping supernova samples — to disagree by roughly that amount. They disagree, but by less: the gap between the Cepheid-based and TRGB-based values is real and unresolved, and it is itself a measure of the residual calibration systematics, but it is smaller than the full tension. Riess et al. (2024b) found the HST Cepheid distances and the JWST subsamples agreeing at the 1σ level. The internal disagreement among the local methods is, on the evidence, a fraction of the tension, not the whole of it — which is consistent with the ledger's finding that the identified systematics are real but insufficient.

What, then, should be done with the ledger's ranking? The third hypothesis of this article holds that the dust degeneracy, with the widest envelope, is the priority target for further constraint, and the ledger supports that hypothesis. But the recommendation must be stated carefully, because it is easy to overclaim. Narrowing the dust envelope from 0.06 mag toward, say, 0.03 mag would not resolve the tension — the tension was never within reach of the dust systematic — but it would do something still valuable: it would remove the single largest remaining hiding place among the identified systematics, and thereby strengthen the conditional conclusion. Each observational program that tightens an envelope makes the systematics-only explanation harder to sustain. The dust degeneracy is where that tightening currently yields the most, which is why it, rather than crowding or the anchor, is the recommended focus. The relevant observational work is the characterization of the host-galaxy dust law across the calibrator and Hubble-flow samples, the program that Brout and Scolnic (2021) opened and that has not yet been closed.

I will end this section with a methodological reflection, because the inverse budget is the contribution of this article and its limits should be owned. The method's strength is that it forces a comparison the forward budget leaves implicit: it puts the closure requirement and the empirical envelope in the same column and computes their ratio. Its weakness is the one already named — it is blind to the unidentified. A second weakness is that the empirical envelopes are estimates, and a reader who assigns the dust envelope 0.09 mag rather than 0.06 mag will get a less decisive ledger. I have tried to mitigate this by taking the generous end of each published range, so that the ledger is biased toward the systematics-only hypothesis and the conclusion against it is therefore conservative. But the envelopes are the soft part of the method, and I would rather say so

plainly than present the ledger as more definitive than it is. The method is a way of organizing the question. It is not an oracle.

CONCLUSION

This article asked whether the Hubble tension — the roughly 5σ disagreement between the cosmic microwave background value of the expansion rate and the local Cepheid-and-supernova value — could be an artifact of systematic error in the calibration of Type Ia supernovae as standard candles. To answer it, the article introduced inverse systematic-error budgeting, a procedure that derives the magnitude each candidate systematic would need in order to close the tension and tests that requirement against the magnitude the data permit. The audit ledger that the method produces is the article's central result, and the three hypotheses can now be addressed against it.

The first hypothesis held that no single identified calibration systematic has a closure requirement falling within its empirical envelope. The ledger confirms it without exception. The closure requirement is 0.17 mag; the widest individual envelope, that of the host-galaxy dust degeneracy, is 0.06 mag; every requirement-to-envelope ratio exceeds one, and most exceed three. The crowding systematic, bounded by the JWST test of Riess et al. (2024a), is not merely insufficient but rejected outright. No individual identified systematic can close the tension.

The second hypothesis held that the closure requirement exceeds even the combined envelope of all identified systematics in quadrature. The ledger confirms this as well: the quadrature sum is 0.11 mag against a requirement of 0.17 mag, a shortfall of roughly a third. The only arithmetic route to closure within the identified set is the linear sum of 0.31 mag, which requires all eight systematics to sit simultaneously at their extreme permitted values and to push in the same direction — a coherent same-sign alignment whose joint probability, for independent systematics, is negligible. The hypothesis is confirmed, with the clarification that the linear-sum scenario is not impossible but improbable to a degree that would itself demand explanation.

The third hypothesis held that the host-galaxy dust-law degeneracy has the widest empirical envelope and the highest residual closure viability. The ledger confirms it: the dust degeneracy carries the widest envelope at 0.06 mag and the lowest ratio at 2.8. The confirmation comes with a firm qualification — highest residual viability is not viability; the dust systematic misses the closure requirement by a factor of nearly three — but the ranking is robust, and it identifies the characterization of the host-galaxy dust law as the observational program that would most efficiently tighten the remaining systematic budget.

The principal original contribution of this article is the method itself: inverse systematic-error budgeting, the inversion of the conventional forward error budget into a requirement-versus-envelope audit. Where a forward budget reports how well the Hubble constant is known, the inverse budget reports whether the identified systematics have enough room to absorb a specified discrepancy, and it delivers that answer as a single transparent ledger with a per-candidate verdict. To my knowledge the Hubble-tension systematics literature had not been organized in this inverted form. The method is general: it is not specific to the Hubble constant and could be applied to any tension between two measurements of the same quantity, wherever the question is whether identified systematics can account for a discrepancy. That generality is part of the contribution, and it is the natural direction for the method's further development.

Three limitations bound the conclusions. The first, and the most important, is that the method audits only identified systematics; an unrecognized calibration error of 0.17 mag would close the tension and would be invisible to the ledger. The conditional nature of the conclusion — conditional on the currently identified set — must not be dropped. The second is that the

empirical envelopes are synthesized estimates assembled from the published systematics literature, not measurements; I have taken the generous end of each published range so that the ledger is conservative against the article's own conclusion, but a reader who prefers different envelopes will get a different, though probably not qualitatively different, ledger. The third is that the analysis treats the eight systematics as independent for the quadrature combination; if hidden correlations link them, the combined envelope could be somewhat wider, though any such correlating mechanism would itself constitute an unidentified systematic and fall under the first limitation.

Two directions for future work follow, and both are consonant with the priorities that the most recent decadal survey of the field set for the coming decade, in which the Hubble tension figures explicitly among the questions deemed most pressing (National Academies of Sciences, Engineering, and Medicine, 2021). The first is observational and specific: the host-galaxy dust law, identified by the ledger as the widest remaining envelope, should be characterized across the calibrator and Hubble-flow supernova samples with enough precision to narrow its envelope toward 0.03 mag, which would not resolve the tension but would remove the largest hiding place among the identified systematics. The second is methodological: inverse systematic-error budgeting should be applied to the other rungs and to other cosmological tensions — the tension in the amplitude of matter clustering, for instance — to test whether the requirement-versus-envelope framing is as clarifying elsewhere as it is here.

The Hubble tension has often been discussed as a choice between two stories: a flaw in the local measurement, or new physics in the early universe. The inverse budget does not choose between them, and it cannot. What it does is retire the most comfortable version of the first story. The discrepancy is not hiding in any single identified calibration error, and it is not hiding in the statistically normal combination of them. Whatever resolves the Hubble tension — new physics, or a systematic no one has yet named — it will not be a calibration mistake that the community already knows about and has merely underestimated. That negative result is modest, but it is, I think, worth stating with the precision the inverse budget allows.

BIBLIOGRAPHY

- Brout, D., & Scolnic, D. (2021). It's dust: Solving the mysteries of the intrinsic scatter and host-galaxy dependence of standardized Type Ia supernova brightnesses. *The Astrophysical Journal*, 909(1), 26. <https://doi.org/10.3847/1538-4357/abd69b>
- Brout, D., Scolnic, D., Popovic, B., Riess, A. G., Carr, A., Zuntz, J., Kessler, R., Davis, T. M., Hinton, S., Jones, D., Kenworthy, W. D., Peterson, E. R., Said, K., Taylor, G., Ali, N., Armstrong, P., Charvu, P., Dwomoh, A., Meldorf, C., ... Wiseman, P. (2022). The Pantheon+ analysis: Cosmological constraints. *The Astrophysical Journal*, 938(2), 110. <https://doi.org/10.3847/1538-4357/ac8e04>
- Di Valentino, E., Mena, O., Pan, S., Visinelli, L., Yang, W., Melchiorri, A., Mota, D. F., Riess, A. G., & Silk, J. (2021). In the realm of the Hubble tension — a review of solutions. *Classical and Quantum Gravity*, 38(15), 153001. <https://doi.org/10.1088/1361-6382/ac086d>
- European Space Agency. (2020). *Planck Legacy Archive: Data products and cosmological parameter tables*. ESA Planck Science Office.
- Freedman, W. L. (2021). Measurements of the Hubble constant: Tensions in perspective. *The Astrophysical Journal*, 919(1), 16. <https://doi.org/10.3847/1538-4357/ac0e95>
- Freedman, W. L., Madore, B. F., Hatt, D., Hoyt, T. J., Jang, I. S., Beaton, R. L., Burns, C. R., Lee, M. G., Monson, A. J., Neeley, J. R., Phillips, M. M., Rich, J. A., & Seibert, M. (2019). The Carnegie-Chicago Hubble Program. VIII. An independent determination of the Hubble

- constant based on the tip of the red giant branch. *The Astrophysical Journal*, 882(1), 34. <https://doi.org/10.3847/1538-4357/ab2f73>
- Freedman, W. L., Madore, B. F., Hoyt, T. J., Jang, I. S., Lee, A. J., & Owens, K. A. (2025). Status report on the Chicago-Carnegie Hubble Program (CCHP): Measurement of the Hubble constant using the Hubble and James Webb Space Telescopes. *The Astrophysical Journal*, 985(2), 203. <https://doi.org/10.3847/1538-4357/adce78>
- Gaia Collaboration. (2021). *Gaia Early Data Release 3: Documentation release 1.1*. European Space Agency, Gaia Data Processing and Analysis Consortium.
- Mörtsell, E., Goobar, A., Johansson, J., & Dhawan, S. (2022). The Hubble tension revisited: Additional local distance ladder uncertainties. *The Astrophysical Journal*, 935(1), 58. <https://doi.org/10.3847/1538-4357/ac7c19>
- Murakami, Y. S., Riess, A. G., Stahl, B. E., Kenworthy, W. D., Pluck, D.-M. A., Macoreta, A., Brout, D., Jones, D. O., Scolnic, D. M., & Filippenko, A. V. (2023). Leveraging SN Ia spectroscopic similarity to improve the measurement of H_0 . *Journal of Cosmology and Astroparticle Physics*, 2023(11), 046. <https://doi.org/10.1088/1475-7516/2023/11/046>
- National Academies of Sciences, Engineering, and Medicine. (2021). *Pathways to discovery in astronomy and astrophysics for the 2020s*. The National Academies Press.
- NASA/IPAC Extragalactic Database. (2023). *NED: Cross-identified extragalactic distance and redshift catalogue*. California Institute of Technology, Infrared Processing and Analysis Center.
- Pesce, D. W., Braatz, J. A., Reid, M. J., Riess, A. G., Scolnic, D., Condon, J. J., Gao, F., Henkel, C., Impellizzeri, C. M. V., Kuo, C. Y., & Lo, K. Y. (2020). The Megamaser Cosmology Project. XIII. Combined Hubble constant constraints. *The Astrophysical Journal Letters*, 891(1), L1. <https://doi.org/10.3847/2041-8213/ab75f0>
- Planck Collaboration. (2020). Planck 2018 results. VI. Cosmological parameters. *Astronomy & Astrophysics*, 641, A6. <https://doi.org/10.1051/0004-6361/201833910>
- Riess, A. G., & Breuval, L. (2024). The local value of H_0 . In *Proceedings of the International Astronomical Union, Symposium S376: At the cross-roads of astrophysics and cosmology* (pp. 15–29). Cambridge University Press. <https://doi.org/10.1017/S1743921323003034>
- Riess, A. G., Anand, G. S., Yuan, W., Casertano, S., Dolphin, A., Macri, L. M., Breuval, L., Scolnic, D., Perrin, M., & Anderson, R. I. (2024a). JWST observations reject unrecognized crowding of Cepheid photometry as an explanation for the Hubble tension at 8 σ confidence. *The Astrophysical Journal Letters*, 962(1), L17. <https://doi.org/10.3847/2041-8213/ad1ddd>
- Riess, A. G., Casertano, S., Yuan, W., Bowers, J. B., Macri, L., Zinn, J. C., & Scolnic, D. (2021). Cosmic distances calibrated to 1% precision with Gaia EDR3 parallaxes and Hubble Space Telescope photometry of 75 Milky Way Cepheids confirm tension with Λ CDM. *The Astrophysical Journal Letters*, 908(1), L6. <https://doi.org/10.3847/2041-8213/abdbaf>
- Riess, A. G., Scolnic, D., Anand, G. S., Breuval, L., Casertano, S., Macri, L. M., Li, S., Yuan, W., Anderson, R. I., & Owens, K. A. (2024b). JWST validates HST distance measurements: Selection of supernova subsample explains differences in JWST estimates of local H_0 . *The Astrophysical Journal*, 977(1), 120. <https://doi.org/10.3847/1538-4357/ad8c21>
- Riess, A. G., Yuan, W., Macri, L. M., Scolnic, D., Brout, D., Casertano, S., Jones, D. O., Murakami, Y., Anand, G. S., Breuval, L., Brink, T. G., Filippenko, A. V., Hoffmann, S., Jha, S. W., Kenworthy, W. D., Mackenty, J., Stahl, B. E., & Zheng, W. (2022). A comprehensive measurement of the local value of the Hubble constant with 1 km s $^{-1}$ Mpc $^{-1}$ uncertainty from the Hubble Space Telescope and the SH0ES Team. *The Astrophysical Journal Letters*, 934(1), L7. <https://doi.org/10.3847/2041-8213/ac5c5b>

- Scolnic, D., Brout, D., Carr, A., Riess, A. G., Davis, T. M., Dwomoh, A., Jones, D. O., Ali, N., Charvu, P., Chen, R., Peterson, E. R., Popovic, B., Rose, B. M., Wood, C. M., Brown, P. J., Chambers, K., Coulter, D. A., Dettman, K. G., Dimitriadis, G., ... Wiseman, P. (2022). The Pantheon+ analysis: The full data set and light-curve release. *The Astrophysical Journal*, 938(2), 113. <https://doi.org/10.3847/1538-4357/ac8b7a>
- Space Telescope Science Institute. (2023). *WFC3 instrument handbook for cycle 31* (Version 15.0). STScI.
- Space Telescope Science Institute. (2024). *JWST NIRC*am* imaging: Observer documentation*. STScI, JWST User Documentation.
- Verde, L., Treu, T., & Riess, A. G. (2019). Tensions between the early and late Universe. *Nature Astronomy*, 3(10), 891–895. <https://doi.org/10.1038/s41550-019-0902-0>
- Wong, K. C., Suyu, S. H., Chen, G. C.-F., Rusu, C. E., Millon, M., Sluse, D., Bonvin, V., Fassnacht, C. D., Taubenberger, S., Auger, M. W., Birrer, S., Chan, J. H. H., Courbin, F., Hilbert, S., Tihhonova, O., Treu, T., Agnello, A., Ding, X., Jee, I., ... Meylan, G. (2020). H0LiCOW XIII. A 2.4% measurement of H_0 from lensed quasars: 5.3σ tension between early- and late-Universe probes. *Monthly Notices of the Royal Astronomical Society*, 498(1), 1420–1439. <https://doi.org/10.1093/mnras/stz3094>

NAPETOST HUBBLE KONSTANTE (H_0) IZMEĐU MJERENJA KOSMIČKOG MIKROVALNOG POZADINSKOG ZRAČENJA I LOKALNIH CEFEIDA: ANALIZA SISTEMATSKIH GREŠAKA U KALIBRACIJI SVJETILJKI TIPA Ia SUPERNOVA

Ashley Brown

Tehnički koledž Atlanta
Atlanta, Džordžija, Sjedinjene Američke Države
E-pošta: ashley.brown@atlantatech.edu

Originalni naučni članak

Primljeno: 09.07.2025. Prihvaćeno: 01.09.2025.

DOI: <https://doi.org/10.65932/CAR-2025-2-3> UDK: 524.8:520.82

Sažetak: Dvije metode mjere sadašnju brzinu širenja svemira i one se ne slažu. Kosmičko mikrovalno pozadinsko zračenje, tumačeno kroz standardni kosmološki model, daje Hubble konstantu od 67,4 km/s/Mpc; lokalna ljestvica udaljenosti, izgrađena od supernova tipa Ia kalibriranih cefeidama, daje 73,0 km/s/Mpc. Razlika je iz znatiželje prerasla u neslaganje od približno 5σ koje se ni jednom količinom dodatnih podataka nije razriješilo. Ovaj članak postavlja namjerno usku verziju šireg pitanja: može li neslaganje biti artefakt sistematske greške u kalibraciji supernova tipa Ia kao standardnih svjetiljki? Standardna analiza grešaka odgovara na to procjenom svake sistematske greške, njenim propagiranjem unaprijed i izvještavanjem o ukupnoj nesigurnosti — postupak koji, po svojoj konstrukciji, ne može reći jesu li sistematske greške dovoljno velike da budu značajne. Ovaj članak uvodi drugačiji postupak. Predlažem inverzno budžetiranje sistematskih grešaka: umjesto propagiranja grešaka unaprijed, metoda polazi od uočene napetosti i izračunava, za svaku sistematsku grešku kandidata, veličinu koju bi morala imati da bi u potpunosti zatvorila procjep — njen zahtjev zatvaranja — te zatim testira taj zahtjev naspram veličine koju podaci stvarno dopuštaju — njenu empirijsku ovojniciu. Sistematska greška kvalifikuje se kao održiv kandidat za zatvaranje samo kada njen zahtjev zatvaranja padne unutar empirijske ovojnice. Primjena metode na osam identifikovanih sistematskih grešaka cefeida i supernova, uz oslanjanje na sedamnaest SCOPUS-indeksiranih studija i sedam institucionalnih izvora, daje jasan rezultat: napetost zahtijeva koherentni pomak od približno 0,17 magnituda u nultoj tački ljestvice udaljenosti, dok nijedna pojedinačna sistematska greška nema empirijsku ovojnicu veću od 0,06 magnituda, a ovojnice sabrane kvadratno dosežu tek 0,11 magnituda. Nijedna pojedinačna kalibraciona greška ne može zatvoriti napetost; samo malo vjerovatno usklađivanje svih osam sistematskih grešaka istog predznaka pri njihovim ekstremnim dopuštenim vrijednostima moglo bi to učiniti. Degeneracija zakona prašine u galaksiji domaćinu izdvaja se kao sistematska greška s najširoom ovojnicom i stoga najvišom rezidualnom održivošću zatvaranja, što je identifikuje kao prioritetni cilj daljnjeg opservacijskog ograničavanja.

Ključne riječi: *Hubble konstanta, Hubble napetost, supernove tipa Ia, kalibracija cefeida, sistematske greške, ljestvica udaljenosti, inverzno budžetiranje grešaka.*