

Verification by Analysis of a Flown Deployable Truss: What the MODE Ground-to-Orbit Frequency Shift Is Not

Marcin W. Kasiak — PhD, IWE, PE

PSE-TM-01 · Revision A · 15 October 2026 · postsingularity.engineer · mkasiak@postsingularity.engineer

Independent research, carried out on personal time using public data and a personally held software licence. Views are the author's own and not those of his employer. Not prepared under, and not offered subject to, a professional seal.

Abstract

The Middeck 0-Gravity Dynamics Experiment (MODE) flew a jointed deployable truss on Space Shuttle mission STS-48 in September 1991, producing one of the few public datasets in which one structure was modal-tested both on a laboratory suspension and in orbit. The measurements show the baseline truss softening between 1 g and 0 g: torsion by 1.4 %, bending by 0.8 %, shearing by 0.7 %, against a reassembly scatter of 0.54 %. Neither the finite element analysis published with the experiment nor its 1994 successor predicted that shift. Both moved by 0.1 % or less, and both authors attributed the discrepancy to gravity acting on nonlinear joints without demonstrating it.

This note reports an independent re-analysis. The test article was rebuilt from the original dimensioned sources, verified against a second, independently written solver to better than 0.5 %, and reproduced the 1993 predictions to within 0.8 % without calibration. Three quantities were then predicted and compared against measurements not used in building the model: the suspension sway frequency, 0.284 Hz predicted against 0.28 Hz documented; the stiffening produced by a fivefold increase in suspension stiffness, +2.53 % predicted against +2.58 % measured, agreement to 0.09 standard deviations; and the preload trend across the three tested cable tensions, -2.34 % predicted against -2.07 % measured, with cable slack emerging unprompted at the lowest preload.

Seven candidate mechanisms for the ground-to-orbit shift were then tested, each as a stiffness perturbation with its gravity driver quantified from converged static states. All seven are eliminated. The result is a structural asymmetry: the shift is torsion-selective, and only cable axial stiffness has the right modal selectivity, but gravity changes mean cable tension by 0.08 %. Every quantity gravity does change substantially, such as longeron compression at 6.3 %, drives bending instead. Damping is excluded separately, being two to three orders of magnitude too small to move the peak.

Two documented source values were also found unreliable: a module mass figure arithmetically incompatible with the stated total, and a pair of preload values that cannot both hold given the confirmed cable topology. The contribution is not the mechanism but a quantified map of where it is not, with the reference uncertainty within which any future candidate must be judged.

1. Introduction

A large deployable space structure cannot be tested at full scale in the environment where it must work. It is packaged into a launch fairing, deployed into a shape far larger than any ground rig can hold without gravity distorting it, and then expected to meet stiffness, shape and pointing requirements that were never demonstrated as a whole. What certifies it is analysis.

That situation is not unique to spaceflight. A long-span bridge is never proof-loaded once complete. It is erected through partial configurations that are individually unstable, its final geometry is reached by controlling each stage, and its capacity is certified analytically under professional liability. The discipline built around that problem — staged construction analysis, geometry control, connection behaviour, and documented verification by analysis — is mature, and it addresses the same question a deployable aperture poses: how much can an analysis be trusted when the test cannot be run.

This note is an attempt to test that transfer against measurements rather than to assert it. The test case is deliberately unforgiving: a structure that was flown, measured in both environments, and analysed twice by the group that built it, leaving a documented discrepancy that neither analysis resolved.

The question addressed here is narrow and answerable: what causes a jointed deployable truss to be measurably softer in orbit than on the ground, and how tightly can a modern analysis, built only from public documents, constrain the answer.

2. The experiment and the open question

The MODE Structural Test Article was a scaled jointed truss: two four-bay deployable modules of 8-inch cubic bays, joined by an erectable centre bay, with a rigid appendage at each end. Longerons were polycarbonate rod with an over-centre latching knee joint at mid-bay; each bay face carried a crossed pair of pretensioned steel cables. One bay had adjustable cable tension, tested at three settings. The article was excited by a proof-mass actuator at three force levels and measured by eleven accelerometers.

It was tested on a soft four-point suspension in the laboratory, with interchangeable springs giving nominal plunge frequencies of 1, 2 and 5 Hz, and again in the Shuttle middeck on slack tethers. Three nominally identical articles were built, and the reassembly and unit-to-unit scatter was measured: 0.54 % in frequency, 0.22 in damping ratio.

The baseline results are the subject here. Comparing the 1 Hz suspension case with orbit, at the lowest force level and highest preload:

MODE	GROUND, 1 HZ SUSPENSION	ORBIT	MEASURED SHIFT	1993 ANALYSIS	1994 ANALYSIS
Torsion	7.74 Hz	7.63 Hz	−1.42 %	−0.1 %	+0.1 %
Bending	20.43 Hz	20.27 Hz	−0.78 %	+0.1 %	≤ 0.1 %
Shearing	29.42 Hz	29.22 Hz	−0.68 %	0.0 %	≤ 0.1 %

The measured softening is small but consistent, and at 2.6 standard deviations for torsion it is well clear of scatter. Both analyses of record predicted essentially nothing.

The authors were explicit about the failure. The 1992 thesis reports that gravity stiffening was included on the flexible elements but that the true stiffening effect on the joints was not captured. The 1994 thesis concludes that better zero-gravity accuracy can only be achieved if all the nonlinear stiffness and dissipative effects of nonlinear structural elements are included. Both name the mechanism; neither demonstrates it. That is the open question this re-analysis addresses.

3. Method

The method is borrowed from bridge certification practice, where an analysis that cannot be checked against a proof load has to carry its own evidence. Four rules were applied throughout.

Every input is tagged by provenance. Each parameter is recorded as documented, with its source page; inferred, derived from documented facts and with the derivation shown; or assumed, named as such. The tags are carried in the model input itself, not in a separate note, so a reader of the model sees which numbers are evidence and which are choices. Where two sources conflict, the conflict is recorded rather than resolved silently.

Nothing is tuned to the quantity being predicted. Two calibrations were made, both against documented values rather than against the measurement of interest: cable pretension was scaled so the converged tension matched the instrumented 25 lbf, and a single parameter in the one nonlinear mechanism trial was set on ground data before being frozen. No frequency was ever adjusted to approach a measured frequency.

The prediction is held out. Orbital results were never used in building or calibrating the model. Where a prediction is reported below against measured data, that data played no part in producing it.

Results are cross-checked against a second solver. The model was built in a commercial finite element code and independently in a small purpose-written finite element program, from the same parameter set but with separate element formulations and mass matrices. Agreement between them proves that the input represents the intended model; it says nothing about whether the model represents the hardware, and the two claims are kept apart.

Scoring uses the measured scatter as its unit. A prediction is reported in standard deviations of the 0.54 % reassembly scatter, so that agreement and disagreement are both judged against the repeatability of the experiment rather than against an arbitrary tolerance.

4. The model and its verification

Geometry came from the 1992 thesis, whose appendix contains the complete input deck of the original analysis, and from its dimensioned figures. One figure proved to be a scale drawing, which allowed the longeron assembly to be measured directly: at 600 dpi the scale is 390 pixels per inch, calibrated against the polycarbonate rod whose diameter is separately tabulated. Measured segment envelopes agree with the deck's segmentation to within 0.03 inch, and every section resolves to an exact stock size, which is strong evidence the tables were read correctly.

The model comprises 239 nodes and 348 elements. Longerons are split into seven segments — fitting, lug, rod, knee joint, rod, lug, fitting — as in the original. Diagonals are tension-only cable elements with prestress. The suspension is modelled explicitly as four hangers whose axial stiffness reproduces the documented

plunge frequency, anchored above the article.

Verification proceeded in two steps, both against things other than the measurement of interest.

Against the second solver, the two implementations agree within 0.5 % across six modes, a systematic offset traceable to consistent versus lumped mass formulation rather than scatter. The input file therefore represents the intended model.

Against the 1993 analysis, reproducing its configuration gives:

MODE	1993 PREDICTION	THIS MODEL	DIFFERENCE
Torsion	7.68 Hz	7.74 Hz	+0.76 %
Bending	19.26 Hz	19.36 Hz	+0.54 %
Shearing	28.27 Hz	28.37 Hz	+0.34 %
Breathing	30.15 Hz	30.30 Hz	+0.49 %

This was achieved without calibration, and mode identity was confirmed from the computed shapes rather than assumed from ordering — two local modes concentrated in the outer bays fall between torsion and bending, and the original author excludes them for the same reason.

The model therefore reproduces the analysis of record. What follows is an addition to it, not a replacement.

5. Two documented values that do not survive

Both were found by reconciling documented numbers against each other, not by chasing a frequency. Both had propagated into results before they were caught.

5.1 The module mass

The summary report states a deployable module mass of approximately 8 lb. The original input deck, itemised, gives 2.27 lb of modelled elements per module, and the natural reading is that joints, hinges and fittings were omitted from the model. Measuring those bodies from the scale drawing refutes it: filling every drawn envelope with solid aluminium — an upper bound that ignores the clevis slots — adds about 0.2 lb per module against a 5.7 lb deficit.

The deficit is instead in the 8 lb figure. The thesis states that the largest configuration tested weighed approximately 47 lb. Reconciling that down to the baseline gives 42.2 lb, against the deck's itemised 42.1 lb, a residual of 0.1 lb. At 8 lb per module the baseline alone would exceed the stated largest configuration by 6 lb.

The consequence is quantitative. Rebuilt on the 47 lb basis, the linear model's stiffness deficit against measured ground frequencies falls from a mode-dependent 1.03 to 1.46 times to a near-uniform 1.01 to 1.12. Most of what appeared to be a mode-dependent stiffness error was one soft figure in a summary report, carried downstream.

5.2 The preload pair

Cable tension at lock is stated as 25 lbf and longeron preload as typically 28 lb, the latter described as 50 % of the estimated buckling load. With the confirmed topology, each end-batten corner receives two cable ends at 45 degrees, so equilibrium forces a ratio of 1.414. The documented pair implies 1.12. They cannot both hold, and no internal stiffness redistributes it: a sweep of corner-body stiffness over four orders of magnitude moves the ratio by 2 %.

The two values do not rest on the same evidence. The cable tension was instrumented, with strain-gauge tabs described and drawn. The longeron figure appears once, qualified twice, as a design estimate. The model is therefore calibrated on the cable. An independent check supports the choice: the design criterion that local longeron string modes remain above 40 Hz is satisfied with margin at the calibrated preload.

A third conflict is recorded rather than resolved. The two theses give different adjustable-bay preloads — 25, 8 and 0.5 lb in one, 24, 13 and 7 lb in the other. The first is used here, being the primary source, stated twice, and consistent with the fixed-bay tension; the discrepancy is flagged in the dataset.

6. Held-out predictions

Three quantities were predicted against measurements that took no part in building the model.

QUANTITY	PREDICTED	MEASURED	ERROR
Suspension sway frequency	0.284 Hz	0.28 Hz	within reported precision
Stiffening, 1 Hz to 5 Hz suspension, torsion	+2.53 %	+2.58 %	0.09 σ
Stiffening, 1 Hz to 2 Hz suspension, torsion	+0.32 %	+0.39 %	0.12 σ
Preload trend PL1 to PL3, torsion	-2.34 %	-2.07 %	0.5 σ

The suspension sensitivity is the most informative of these. The original author notes that the difference between the nominal and stiffest suspensions tended to exceed the difference between the nominal ground tests and the orbital tests. The boundary effect is therefore larger than the effect under investigation, and predicting it to 0.09 standard deviations establishes two things: the ground reference is sound, and the model responds correctly to a genuine stiffness change of the relevant magnitude. A model that could not do this would be unable to rule anything out.

The preload trend is a weaker but still useful check, and it produced one unprompted result. With tension-only cable elements and calibrated pretension, the lowest preload setting produces slack cables in the model without being asked to. The original testing reports slackening of the tensioning cables as an observed nonlinearity at that setting, and explicitly reports that stiffening due to joint deadbands is not in evidence. The model reproduces the behaviour the experiment saw.

The bending and shearing preload trends keep the right sign but are five to six times too small. That is noted rather than explained.

7. Mechanism elimination

Each candidate was treated the same way. The mechanism was applied as a stiffness perturbation and its modal selectivity computed — which modes it moves, and in what proportion. Separately, the physical quantity that drives it was measured between the converged 1 g and 0 g states. A mechanism survives only if it has both the right selectivity and a driver large enough to produce the measured shift.

The target is torsion -1.42% , bending -0.78% , shearing -0.68% , with torsion dominant in the ratio 1 : 0.55 to bending.

MECHANISM	SELECTIVITY	GRAVITY DRIVER	OUTCOME
Rotational free-play, 32 knees and 64 lugs	-0.075% torsion at a free hinge	not reached	Short by 19 to 55 times even at zero joint stiffness
Longeron axial compliance	bending-to-torsion 8 to 126	+6.3 % longeron compression	Wrong selectivity: drives bending, measurement is torsion
Cable termination contact, Hertzian	correct by construction	+0.08 % cable tension	No driver: 0.0004 % maximum, and the wrong sign
Mean cable tension	torsion-dominant, 1 : 0.31	+0.08 % cable tension	Needs 3 to 5 % stiffness loss; gravity supplies 0.03 %
Cable tension distribution	torsion-dominant	spread widens, mean unchanged	Reaches 32 % of torsion, none of bending or shearing
Cable slackening	fixes the preload trend	2 of 8 cables slack at 1 g	Wrong sign for the gravity shift, by 5.6 σ
Batten-frame load path	not reached	not reached	Closed on statics: ratio immovable over 10,000 times stiffness

Two entries deserve comment.

Rotational free-play was the mechanism both original authors implicated, and it fails for a structural reason rather than a numerical one. A truss carries global bending and torsion as longeron axial force; the bending stiffness of a longeron barely enters those modes. Rotational freedom at the knee therefore acts on a degree of freedom the global modes do not use. The documented 2-degree over-centre latch geometry gives a real deadband, but on the wrong axis.

Cable termination contact was the most promising candidate and was retracted during the study. A 3 to 5 % loss of cable axial stiffness reproduces torsion and shearing to within a factor of 1.02, and the implied seating travel of 11 to 17 micrometres is ordinary for a steel ball in a spherical receptacle. But the contact force at a termination is the cable tension, and gravity changes mean cable tension by 0.08 %. Because contact stiffness varies as the cube root of load, no parameter value in the family produces more than 0.0004 %, and the concavity gives the wrong sign. A sufficient perturbation with the right modal signature is not the same as a mechanism gravity can drive; the distinction is the reason the candidate closed.

8. Discussion

The eliminations share a structure, and stating it is the main result of this note.

The shift is torsion-selective, and that is what closes the field. Torsion and shearing store much of their strain energy in the diagonal cables; bending strains the longerons. A mechanism is therefore constrained not only to be large enough but to load the right members. Of the candidates tested, only cable axial stiffness has the correct dominant mode, at a torsion-to-shearing ratio of 1 : 0.31 against a measured 1 : 0.48. And gravity moves cable tension by 0.08 %. Conversely, the one quantity gravity moves substantially — longeron compression, by 6.3 % — drives bending by a factor of over a hundred relative to torsion, which is the inverse of what is measured.

Right selectivity, no driver. Right driver, wrong selectivity. Every mechanism tested falls on one side or the other.

The shift is stiffness, not damping. Measured damping roughly doubles between ground and orbit, and one might ask whether the frequency change is a damped-natural-frequency artefact. It is not: at these damping levels the peak moves by 0.0005 to 0.006 %, short by two to three orders of magnitude. Even the most extreme damping change in the dataset gives 0.065 %.

The reference has a floor. Any candidate mechanism must clear the uncertainty in the 1 g reference before agreement means anything. Hanger placement, which the sources do not fix, is worth 0.71 standard deviations on torsion. The choice between the softest suspensions is worth another 0.72. Against an effect of 2.6 standard deviations, roughly 0.7 is irreducible without better documentation of the test rig.

What this leaves. The measured shift is real, is stiffness, is torsion-selective, and is not produced by any of the seven linear mechanisms available in a model built from the documented geometry. That is a narrower statement than either original study reached, and it constrains the next attempt.

9. Limitations

Every test was linear. Each mechanism was applied as a stiffness perturbation and evaluated by eigenvalue analysis about a converged static state. The claim supported is therefore that no linear mechanism of the kinds tested reproduces the shift. It is not a claim about nonlinear response.

That matters because the measurements are amplitude-dependent. Frequency falls and damping rises with excitation level in every mode, the shift itself varies with force level — torsion holds near -1.4 % while bending falls from -0.78 % to -0.20 % — and the orbital data show abrupt magnitude changes with no counterpart on the ground. Testing against that behaviour requires nonlinear forced response, which would resolve amplitude and damping together. It was not attempted here, and it is the natural next study.

Two geometry items remain assumed. The suspension attachment stations are not documented in usable form; the only statement available is that the article hung from four top corners. Which module carried the adjustable bay is also unstated. Neither affects the free-free results, and the hanger sensitivity is quantified above.

One structure, one flight. This is a scaled laboratory truss with polycarbonate members, not a flight aperture. What is demonstrated is a method and its evidentiary standard; the numbers belong to this article alone.

The 1 g reference could be better known. Since the suspension effect exceeds the gravity effect, an experiment designed to isolate the latter would need the rig documented to a standard the original report does not reach. That is a finding about experiment design as much as about analysis.

10. Conclusions

1. A model of the MODE truss built only from public documents reproduces the 1993 analysis of record to within 0.8 % without calibration, and agrees with an independent solver to within 0.5 %.
2. It predicts three held-out quantities, the strongest to 0.09 standard deviations of the experiment's own repeatability. The ground reference is therefore sound and the model responds correctly to real stiffness changes.
3. Seven candidate mechanisms for the ground-to-orbit frequency shift are eliminated, each with its modal selectivity and its gravity driver quantified.
4. The eliminations have a common structure: only cable axial stiffness has the right modal selectivity, and gravity does not drive it; every quantity gravity drives strongly has the wrong selectivity.
5. The shift is stiffness rather than damping, by two to three orders of magnitude.
6. Two documented source values do not survive scrutiny: a module mass incompatible with the stated total, and a preload pair incompatible with the confirmed topology. The first accounted for most of an apparent mode-dependent stiffness deficit.
7. The reference uncertainty floor is about 0.7 standard deviations against a 2.6 standard deviation effect, set by undocumented suspension details.

The mechanism remains unidentified. What has changed is that the space of candidates is now mapped, with quantified reasons for each exclusion, and the next study — nonlinear forced response, resolving amplitude and damping together — has a defined target and a verified model to start from.

11. Sources and reproducibility

All sources are public. The experiment summary report is a US Government work cleared for public use; the two theses are freely available from their institution's repository and carry no distribution limitation. No controlled or export-restricted material was used.

Primary sources

- Crawley, Van Schoor and Bokhour, *The Middeck 0-Gravity Dynamics Experiment — Summary Report*, NASA-CR-4500, March 1993. Configuration description, suspension and actuator details, tabulated ground and orbital modal parameters, scatter statistics.
- Barlow, *Modeling and ground modal identification of space structures*, MIT S.M. thesis, 1992. Member geometry, dimensioned drawings, the original input deck, ground test results and nonlinearity observations.
- Bobronnikov, *Modeling and modal identification of jointed space structures in one- and zero-gravity environments*, MIT S.M. thesis, 1994. Second-generation models and their orbital predictions.

Reproducibility. The model is generated from a single parameter file in which every value carries its provenance tag and source page, so the input can be audited against the sources without reading the solver input. The transcribed measurement dataset, the parameter file, the generator and the result-reading scripts are available on request.

Tools. SOFiSTiK 2022, on a personally held licence, with an independently written finite element program in Python as cross-check. All work was performed on personal equipment and personal time, using public data only.