Print copy. This is the print copy of an article whose version of record is a web page. That page, with the data and code behind every figure, is at
https://hertzforge.com/journal/2026-010/v1/
DOI of this version: 10.5281/zenodo.24311109

RESEARCH ARTICLE

When the Control Gauges Read Right and the Stress Is Wrong: A Closed-Loop Check of the Proposed Practice for Multi-Shaker Tests

Kevin Cross ORCID 0009-0002-6608-9426

Evidence	**Simulator.** Closed-loop runs of six control laws through Hertzforge's multiple-shaker random control service on its simulator, with measurement noise, on 20 generated frame structures of the Open Benchmark Set (version 0.2) in five setups, linear and at two sizes of nonlinearity, with the main findings repeated on ten structures held back from the study. No hardware tests.
Version	Version 1, published 2026-11-17. First published 2026-11-17. Change note: First published.
Latest version	10.5281/zenodo.24311108
Cite as	Cross, K. (2026). When the Control Gauges Read Right and the Stress Is Wrong: A Closed-Loop Check of the Proposed Practice for Multi-Shaker Tests. Hertzforge Journal 1, 2026-010, version 1. https://doi.org/10.5281/zenodo.24311109

Abstract

A multi-shaker vibration test is usually judged by its control accelerometers: if they read their targets, the test is called good. The first article of this series proposed a practice for judging such tests. Its starting point, from idealized simulations and a published hardware test, was that with as many shakers as control accelerometers a controller can make the controls read right while the article is stressed well above what it sees in service, and that accelerometers held out of the control are the ones that show it. That article named the checks that would settle the point. A simulator can make four of them: measurement noise, a frequency response measured on the shakers rather than known exactly, a closed control loop, and structures whose response changes with level, and it can use more structures than that article had. This article reports what those checks show, from a study of six control laws run closed loop through Hertzforge's control service on 20 generated frame structures in five setups, with each main finding repeated by an independent audit on ten structures held back from the study. On the same 12 control gauges, adding shakers until there were as many shakers as gauges lowered the control-gauge error of the module's default law by 1.2 dB and raised the stress error by 8.8 dB on all 20 structures (8.5 dB on all ten held-back structures). A setting that limits the drive in the weak directions left the control gauges where they were and lowered the stress error by 6.4 dB (6.3 dB held back). With as many shakers as control gauges, the held-out gauges tracked the stress error and the control gauges did not. A law that limits its drive in the weak directions, and a law that works on band averages, did not show the over-test. Everything here is simulation, and nothing has met hardware.

Contents

Summary . . . 3
1. The First Article . . . 4
2. The Test . . . 4
2.1 The Structure . . . 4
2.2 The Five Setups . . . 5
2.3 Driven, Measured, Judged . . . 6
2.4 Kinds of Law . . . 7
2.5 The Controller’s Checks . . . 7
2.6 The Runs . . . 8
2.7 The Independent Audit . . . 8
3. Linear Structures . . . 9
3.1 Shakers Added . . . 9
3.2 Fewer Control Gauges . . . 11
3.3 Tuned Settings . . . 12
3.4 Which Measurement Warns . . . 13
3.5 Report Card Lines . . . 16
4. Nonlinear Structures . . . 17
5. Measuring Again . . . 18
6. What the Checks Show . . . 18
7. Limits . . . 19
Contribution Statement . . . 19
How This Was Made . . . 19
Acknowledgments . . . 20
References . . . 20
Implemented In . . . 20
Version History . . . 20

Summary

A vibration test with several shakers is controlled from a set of accelerometers on the article, the control gauges. The controller adjusts the shaker drives until the control gauges read their target spectra. If they do, the test is usually called good.

The first article of this series (Cross and Zwink 2026) showed, in simulations where the stress in every part of the article was known, that this is not enough. With as many shakers as control gauges, a controller that inverts the measured response exactly made every control gauge read within a tenth of a decibel of its target, and the article was stressed about 11 dB more than in service. Gauges left out of the control, the held-out gauges, showed it; the control gauges did not. On hardware, Rohe, Nelson and Schultz (Rohe et al. 2019) had found that with as many shakers as control gauges the error at the control gauges went nearly to zero while the error at two gauges held out of the control grew large. That article's simulations were idealized: the controller knew the article's response exactly, there was no noise and no loop, and the article did not change with level. Its Section 8 listed what would settle the matter.

This article reports a study that makes the checks a simulator can make. Six control laws ran closed loop through Hertzforge's own control software, with measurement noise, each measuring the article's response on the shakers before the test, on 20 generated frame structures, in five arrangements of shakers and control gauges, on structures that do not change with level and on the same structures with joints that soften and a member that hardens. The stress in every beam element on the shakers was compared with its stress in service. An independent audit then ran the main comparisons again on ten structures that nobody had generated or seen while the findings were being written.

What it found, each with the audit's verdict:

- **More shakers on the same control gauges, more over-test.** On the same twelve control gauges, going from six shakers to twelve lowered the default law's control-gauge error by 1.2 dB and raised its stress error by 8.8 dB, on all 20 structures. Audit: holds; the stress error rose by 8.5 dB on all ten held-back structures.
- **Fewer control gauges, more over-test, comparing two setups.** Six shakers on six control gauges against six on twelve: the control-gauge error was 2.7 dB lower and the stress error 7.4 dB higher, on all 20 structures. The two setups place their shakers differently, so this compares two setups, not the count of control gauges alone. Audit: holds; 8.7 dB on all ten held-back structures.
- **Two tests alike at the control gauges can be far apart in stress.** With six shakers and six control gauges, a setting that ignores the weakest drive directions left the control-gauge error where it was and lowered the stress error by 6.4 dB, on all 10 structures run. Audit: holds; 6.3 dB on all ten, the control gauges again unchanged.
- **With as many shakers as control gauges, the held-out gauges track the stress and the control gauges do not.** Rank correlation with the stress error: +0.85 at the held-out gauges, +0.25 at the control gauges (held back: +0.81 and +0.27). Audit: holds.
- **Two lines of the report card did not separate good tests from bad.** The over-test alarm fired on every run. The exact-fit flag almost never fired, because in a closed loop with noise the control error never came near zero, while the over-test it was meant to catch happened in most runs. Audit: holds; the alarm fired on every held-back run and the flag on none.

The failure belongs to the control laws that invert the whole measured response. A law that limits its drive in the weak directions, and a law that works on band averages, did not over-test on these setups. The finding of the first article holds without an exact fit: the control gauges need only read about right, not exactly right, for the stress to be wrong.

1. The First Article

The first article (Cross and Zwink 2026) proposed a practice for judging multi-shaker tests. Two of its proposals are tested here. P2: never read a small control error as a good test; with as many shakers as control gauges, cap the inverse so that it ignores directions much weaker than the strongest, or free the cross terms, or use a band law; better, use fewer shakers than control gauges. P3: keep some gauges out of the control and use them only to judge the test, because they are the only gauges that see an exact-fit failure. Its report card puts these into lines: line 2 grades the pattern of the spectra at the held-out gauges and at the control gauges, line 3 raises an alarm when any band at any gauge is more than 6 dB over its target, and line 6 flags "a near-zero control error with shakers equal to controls".

The evidence was idealized. Every law computed its drives from the true frequency response of the article, with no noise, no measurement error and no loop. The article was linear, and most of its random results rest on ten structures, which its Section 8 calls thin. That section lists what would settle the matter. Four of the items can be checked in a simulator and are checked here: noise, identification error, a closed loop, and articles whose response changes with level. This study also uses more structures, from the same family. A second structural family, also on the list, is not attempted, and the other items need hardware.

This article reads one set of runs by the setup of the control problem: how many shakers, how many control gauges, and which gauges warn. The benchmark article (Cross 2026e) defines the problems and the scorecard and reads the runs by the effect of softening and hardening; the control law study (Cross 2026c) reads them law by law. How each law works is described in the primer on the six laws (Cross 2026d). The problems are published as a data set (Cross 2026b), and the study's scorecards with the audit's record as a second one (Cross 2026a). Nothing here grades a law.

2. The Test

2.1 The Structure

Each structure is a box chassis of square aluminum tube, about 0.6 by 0.4 by 0.3 m, with small components on cantilevered tubes (three on structure 0), lumped masses standing for boxes bolted to the chassis, and a hardening member: a small mass on two thin steel strips clamped at both ends, which stiffens as it moves (Figure 1 a). It is the frame family of the first article, reworked for the benchmark (Cross 2026e). The model is made of beam elements, 96 of them for structure 0, which has 76 flexible modes between 20 and 2000 Hz.

Every structure is generated from a whole number, its seed, which is its label in the data set. A different seed gives a different structure of the same kind: its dimensions, masses, damping, service loads, where its gauges and shakers go, how much its joints soften and where its member sits. The study used 20 structures, labeled by their seeds: 0 to 9, 12 to 15, 17 and 19, and 32 to 35. Four seeds in the first twenty (10, 11, 16 and 18) gave structures on which the member could not be placed, and the next unused seeds took their place, as the benchmark's rules require. Seeds 20 to 29 were held back for the audit (Section 2.7) and were not generated, run or read by the study.

In service the article is bolted at its four bottom corners to a larger host frame, and random forces act on the host. The targets are what a first unit of the design reads in service. The unit on the shakers is a second unit of the same design with slightly different stiffness, mass and damping, as when a test article is not the serial number the specification came from. On the shakers the article hangs on a soft suspension.

2.2 The Five Setups

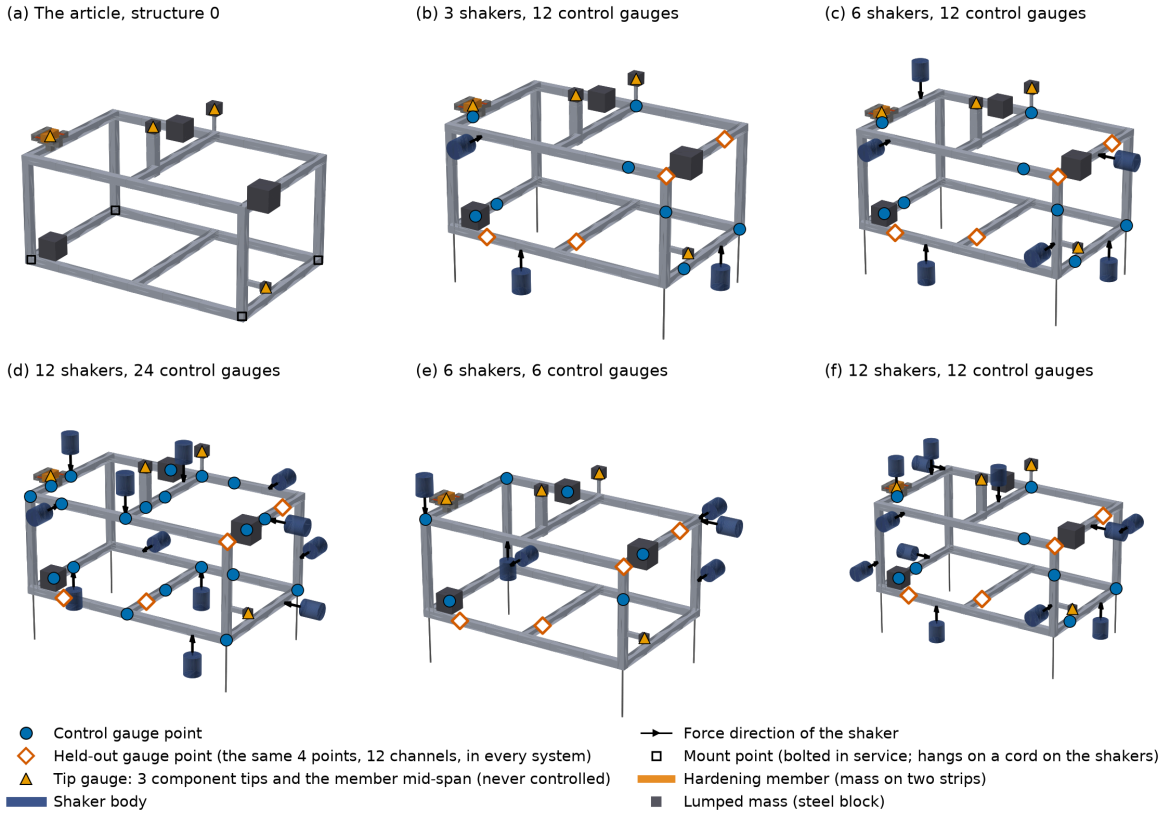


Figure 1: The article of structure 0 and the five systems, rendered from the beam model. (a) The article: the tube chassis, the three components on cantilevered tubes with their tip masses, the lumped masses, the hardening member, the four mount points and the tip gauges. (b) to (f) The article on its shakers in each system: each shaker body with its force direction, the control gauge points, the held-out gauge points and the tip gauges. Shakers and suspension cords show position and direction only. The stress is read in every beam element.

A **setup** here is one arrangement of shakers and control gauges on the article. The article on its shakers with its gauges in one setup, as the controller sees it, is a **system**; the first article of the series calls it the rig. Five setups were run on every structure (Figure 1): 3 shakers with 12 control gauges, 6 with 12, 12 with 24, 6 with 6, and 12 with 12. The first three have more control gauges than shakers; the last two have as many shakers as control gauges. Each gauge is one accelerometer channel, one point in one direction.

- **Control gauges** are read by the controller, which drives the shakers to bring them to their targets. They are chosen by a standard sensor-placement method (Effective Independence) on the chassis and the component bases.
- **Held-out gauges** are recorded and never given to the controller. There are 12, three directions at each of four chassis points, the same in every setup of a structure. They have targets of their own from service and are used only to judge the test.
- **Tip gauges**, 12 channels at the three component tips and the middle of the hardening member, are never controlled. They have a service response to compare with.
- **Shakers** push at chassis points in one direction each. Each is an ideal point force with no dynamics of its own. They are placed one at a time, each where it best keeps the shakers able to move the control gauges in independent ways.

How the setups relate matters for reading the results. On every structure, the setups with 3, 6 and 12 shakers on 12 control gauges share the same 12 control gauges, and their shakers are nested: the 3 are the first 3 of the 6, and the 6 are the first 6 of the 12. Going from 6 shakers to 12 on those gauges adds shakers and changes nothing else. The setup with 6 shakers and 6 control gauges is different: its shakers are placed anew for its own control gauges. Over the 20 structures it shares 0 to 2 of its 6 shaker positions and 1 to 4 of its control gauges with the setup of 6 shakers and 12 control gauges. The same holds between 12 shakers on 12 control gauges and 12 on 24 (1 to 5 shakers and 1 to 10 gauges shared). The held-out gauges are the same in every setup.

2.3 Driven, Measured, Judged

Each run is one control law on one structure in one setup, run through Hertzforge's multiple-shaker random control service, the code that runs a real test, on its simulator. The runs were made by the study's run program and its player, the examples `x4_runner` and `obs_play` of the crate `sds-control` in the Hertzforge repository (github.com/Hertzforge/Hertzforge), on commit `646c4ca4`. The same programs are on the repository's main branch at commit `268c3bfd`, where six cases rerun gave scorecards equal in every number. The run follows the problem's limits (Cross 2026e):

1. With the drives off, the noise floor is measured.
2. The system's frequency response, from each shaker's drive to each control gauge, is measured with uncorrelated random drive about 20 dB under the target level. This is **the measured frequency response**: the controller's only knowledge of the article. It carries the noise of the measurement, and on a structure that changes with level it is the article as it is at low level.
3. The law computes a first drive from the targets and the measured frequency response, and the level rises in steps of 3 dB from 12 dB under the target to full level, one loop at each step.
4. At full level the loop runs 30 times. In each loop the controller plays the drive, estimates the spectra at the control gauges, and the law corrects the drive. The last 15 loops are the record.

Measurement noise is 40 dB under the quietest control gauge's target. The analysis band is 20 to 2000 Hz in lines 0.5 Hz apart. The **target** at the control gauges is the full set of service spectra line by line, with the cross spectrum between every pair, except for the band-average law, which runs the banded specification: the service spectra as one-third-octave band values, with no cross terms, as it is designed to be. The comparisons between that law and the others are therefore comparisons of a law with its own form of target.

What is judged:

- **Control-gauge error**: at each control gauge and each one-third-octave band, how far the spectrum measured over the record is from the target, in decibels; then the RMS over gauges and bands. This is the per-gauge error of line 2 of the report card.
- **Held-out-gauge error**: the same at the 12 held-out gauges.
- **Summed error**: the same, but on the sum of the spectra over the gauges of a group, so that one gauge over and another under can cancel. This is the measure that line 2 grades.
- **Stress error**: for each beam element, its RMS stress on the shakers over its RMS stress in the same unit in service, in decibels, positive meaning over-test. The stress error of a run is the RMS of that over the elements, and the **stress mean** is its mean. The stress comes from the same model as the gauges, so it is a truth no real test has.

Lower is closer for every error. A change in an error is the error of one run minus that of another, on the same structure.

2.4 Kinds of Law

Hertzforge ships six laws for this test, described with their equations and defaults in the primer (Cross 2026d). This article names a law only where a result differs by the kind of law, and the kinds are three (Table 1).

Table 1: The three kinds of control law in the study, with the primer's sections.

Kind	What it does	Laws (code names)
Inverts the whole measured response	At each frequency line it inverts the measured response from shakers to control gauges and drives to make the control spectra and cross spectra match	`pinv_closed` (the module's default, primer Section 2.2), `pinv_open` (no loop correction, 2.3), `buzz_match` (2.5), `min_drive` (2.6)
Limits its drive in the weak directions	It drives only the directions in which the system responds most strongly on each line	`shape_svd` (2.4)
Works on band averages	It sets one flat drive per fractional-octave band, matched to the band-averaged control spectra	`band_average` (Brandon Zwink's law, 2.7)

Each law ran at its default settings on every structure and setup. On the first ten structures, in the two setups with six shakers, each law also ran at one tuned setting, chosen beforehand on two other structures (structures 30 and 31, used for nothing else) as the setting with the smallest held-out-gauge error. For the default law on the setups with as many shakers as control gauges, the tuned setting damps the weak directions (Tikhonov regularization, amount 0.1); for the open-loop form of the same law it ignores directions more than 10 times weaker than the strongest. Both limit the drive in the weak directions, which is what P2 of the first article proposes. The tuned settings were chosen on a build whose control-law, simulator and controller code is identical to the study's.

2.5 The Controller's Checks

The control module has checks that stop a run on its own (primer Section 3). The growing-error check stops a run when one control gauge's error stays at or above 3 dB without improving by 0.1 dB for 5 judged loops, the loops of the level ramp included. The drive-change check stops it when the cap on how fast a drive may change has held a drive back on 10 or more lines in 5 of the last 10 drives the law was allowed to make. The drive-voltage abort is the drive ceiling plus 1 dB: with the problem's 5.0 V ceiling, 5.61 V RMS. These defaults are conservative on purpose. They protect the shaker, the amplifier and the article, and widening them is the test engineer's decision for their own equipment and article.

A simulator has no hardware to protect, so the study widened the checks so that a run would complete: no limit on the count of lines outside abort lines, a stop only after 16 loops without improvement, and an overload fraction of 1.0. On the first five structures every case was also run with the checks at their defaults: 990 runs, of which 901 stopped (840 on the growing-error check, 61 on the drive-change check), most at the seventh loop and a few as late as the 34th, and all 901 completed once the checks were widened. On the ten held-back structures, 811 of 870 runs at the default checks stopped (792 and 19), at loops 6 to 35, and 810 completed once widened. A run stopped by a default check is a fact about the default, not a failure of the law.

The results here are each case's completed run with the checks widened where it exists, else its completed run at the default checks. For the linear structures at default settings that is 598 cases: 577 run with the checks widened and 21 that completed at the default checks. Three cases have no result: one law, on the setup with three shakers, reached the drive-voltage abort even with the checks widened, and the cases are left out (one such run on the held-back structures). Most runs at the default checks would have stopped early, so the results describe what each law does when it is let run, not what a laboratory running the defaults would see.

2.6 The Runs

Table 2: The study's runs this article uses. Each case was run once. The scorecards of every run are in the results record (Cross 2026a).

Runs	Structures	Setups	Stages	Cases
Each law at its default settings, the measured frequency response measured once and kept	20	all five	linear, realistic, strong	1,797
Each law at its tuned setting, the same	10 (structures 0 to 9)	6 with 6, 6 with 12	all three	360
Each law at its defaults, the frequency response measured again during the test	10	6 with 6, 6 with 12	realistic, strong	240

Stages. At stage 1 the structure is linear: it is the same at every level. At stage 2 the same structure changes with level: its bolted joints soften (stiffness falls as the motion grows, so those modes fall in frequency) and the member hardens (stiffness rises with its displacement, so its mode rises). Stage 2 comes in two sizes, realistic and strong. The benchmark article (Cross 2026e) describes both effects and their sizes.

Measured once and kept, or measured again. In most runs the measured frequency response is measured once before the test and kept. In one set of runs the law measures it again during the test and blends the new estimate with the old, so that it follows a structure that changes with level.

Spread. Every result is given as the median over structures with the first and third quartiles in brackets, the number of structures that moved up or down by more than 0.2 dB, and where it matters the structure with the largest change. Each case was run once, so the spread from one run to the next of the same case is not measured. Differences under 0.2 dB are not claimed.

2.7 The Independent Audit

Five comparisons (the five findings of the Summary) were written down, with the study's numbers and the test each would be held to, before ten further structures were generated. The audit then ran them on those ten held-back structures (seeds 20 to 23, 26 to 28, 50, 51 and 52; 50 and 51 replace 25 and 29, on which the member could not be placed), as a separate agent with its own code, using the set's own generator, the study's runner build and the scorer. Structure 24 was refused by the simulator at the strong size (one mode could rise past its 4,096 Hz bound) and was replaced by structure 52; its completed runs count in no verdict. Every verdict was computed by code that first reproduces the study's numbers from its scorecards exactly. The audit's scorecards, its held-back problems and its record are part of the results record (Cross 2026a).

The audit's rule: a result **holds** when the held-back structures move the same way and their median falls inside the study's quartiles; where the direction is the same but the median falls outside, the size differed, and the text says which way. A result opposite in direction, or under 0.2 dB on the held-back structures, does not hold, except where the study itself found no clear change, which such a result agrees with. The audit did not run the cases that measure the frequency response again during the test, and one run per case leaves run-to-run noise unmeasured there too.

3. Linear Structures

Table 3 sets the scene: each kind of law on each setup, at its defaults.

Table 3: Control-gauge, held-out-gauge and stress errors (dB, median over 20 structures) for the default law, which inverts the whole measured response, the law that limits its drive in the weak directions, and the band-average law (on its banded specification). Linear structures, default settings, frequency response measured once and kept. “Weak directions” is the law that limits its drive in the weak directions; “band averages”, the band-average law. “6 x 12” is 6 shakers and 12 control gauges, and so on. Every law and the quartiles: `data/by_setup_linear.csv`.

Law and error	3 x 12	6 x 12	12 x 24	6 x 6	12 x 12
Default law: control	7.86	6.01	5.74	3.10	4.72
Default law: held-out	8.06	6.50	5.43	10.07	11.18
Default law: stress	4.87	3.40	2.71	11.13	12.81
Weak directions: stress	5.56	3.83	3.40	3.21	3.58
Band averages: control	6.81	4.11	3.14	2.00	2.19
Band averages: held-out	6.76	5.43	3.80	6.34	4.72
Band averages: stress	4.29	3.13	2.20	3.12	2.61

On the setups with more control gauges than shakers, the three kinds sit within about 2 dB of one another in stress, and lower control error goes with lower stress. On the two setups with as many shakers as control gauges, the default law reads lower at its control gauges than it did with twelve control gauges and six shakers, and its stress error is 11 to 13 dB, the article over-tested by about 10 to 12 dB on average. The other two kinds stay near 3 dB.

3.1 Shakers Added

The cleanest view of the setup is the one that changes only the number of shakers. On the same 12 control gauges, Figure 2 (a) follows each structure from 3 shakers to 6 to 12.

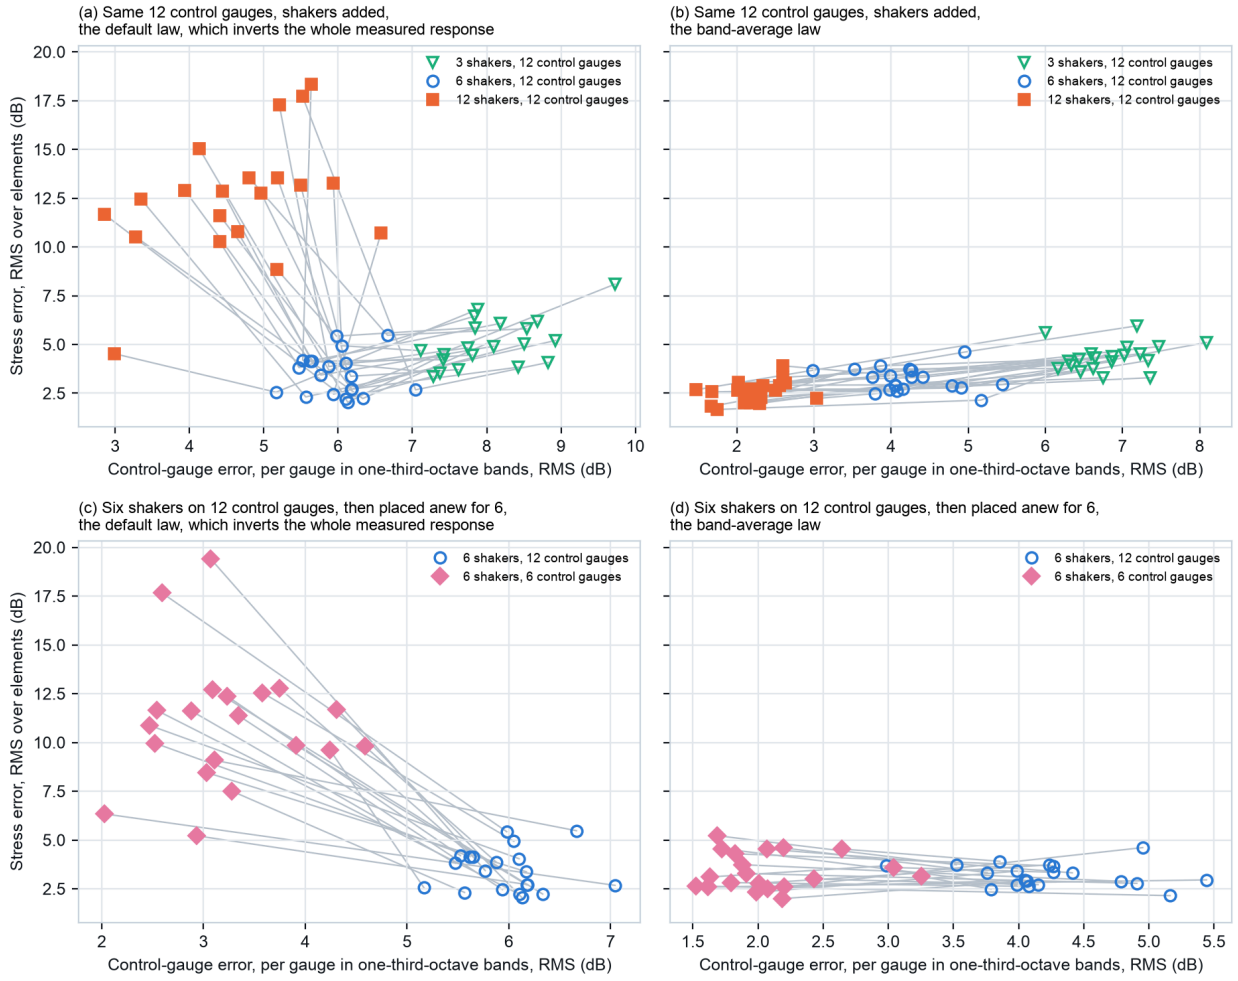


Figure 2: Control-gauge error against stress error, structure by structure, as the setup changes. Each gray line joins one structure's runs. Top: the same 12 control gauges with 3, 6 and 12 shakers. Bottom: 6 shakers on 12 control gauges, and 6 shakers placed anew for 6 control gauges. Left: the module's default law, which inverts the whole measured response. Right: the band-average law, on its banded specification. Linear structures, default settings.

Table 4: Twelve shakers minus six, on the same 12 control gauges, same structure and law (dB, median [quartiles] over 20 structures). Linear structures, default settings. For the other two laws that invert the whole response the range of their medians is given; each is in `data/setup_analysis.txt`. The columns are the control-gauge, held-out-gauge and stress errors, and the structures up and down in stress.

Kind of law	Control	Held-out	Stress	Up / down
Default law	−1.21 [−1.90, −0.77]	+4.78 [+3.97, +5.18]	+8.82 [+8.20, +10.33]	20 / 0
Open-loop form of the default law	−1.07 [−1.76, +0.09]	+4.83 [+3.68, +5.61]	+9.10 [+7.86, +11.37]	20 / 0
Other two laws that invert the whole response	−0.74 to −0.61	+1.78 to +5.11	+2.74 to +9.82	16 to 20 / 0 to 2
Limits its drive in the weak directions	−0.43 [−0.57, −0.14]	−0.52 [−0.67, −0.21]	−0.21 [−0.61, +0.08]	4 / 10
Works on band averages	−1.79 [−2.50, −1.56]	−0.88 [−1.13, −0.50]	−0.46 [−0.94, −0.24]	1 / 16

Going from 3 shakers to 6 lowered every error for every kind of law: control gauges, held-out gauges and stress fell together (the default law by 1.9, 1.3 and 1.3 dB). Going from 6 to 12 split the kinds (Table 4). For the laws that invert the whole measured response, the control gauges read closer to their targets and the article was stressed far more: the default law's stress error rose by 8.8 dB on all 20 structures (largest, structure 0, +15.0 dB), and its stress mean, the average over-test, by 10.7 dB. The held-out gauges rose by 4.8 dB and showed it. The open-loop form, which computes its drive afresh each loop with no correction from the last, moved the same way by about the same amounts (stress +9.1 dB, 20 of 20). For the law that limits its drive in the weak directions and for the band-average law, twelve shakers did slightly better than six on every measure.

Audit: holds. On the ten held-back structures, twelve shakers against six raised the default law's stress error by 8.53 dB (study 8.82), on all ten. Its control-gauge error fell by 0.95 dB (study 1.21), on 8 of 10. Its held-out-gauge error rose on all ten, by less than on the study's structures: 3.50 dB against 4.78. The band-average law's stress error fell by 0.74 dB (study 0.46), on 8 of 10. The law that limits its weak directions moved by under 0.2 dB (-0.09 dB), which agrees with the study's small change.

3.2 Fewer Control Gauges

The second comparison fixed before the audit is six shakers on six control gauges against six shakers on twelve (Figure 2 c and d), and twelve on twelve against twelve on twenty-four. These change more than the count of control gauges: the shakers are placed anew for the new control gauges (Section 2.2). Read them as two setups of the same structure, one with as many shakers as control gauges and one with twice as many control gauges as shakers; they do not isolate the number of control gauges.

Table 5: Six shakers on six control gauges minus six shakers on twelve, same structure and law (dB, median [quartiles] over 20 structures, and structures up or down in stress by more than 0.2 dB). Linear structures, default settings. The first four laws invert the whole response, `shape_svd` limits its drive in the weak directions and `band_average` works on band averages (Table 1). The columns are the control-gauge, held-out-gauge and stress errors, and the structures up and down in stress.

Law	Control	Held-out	Stress	Up / down
`pinv_closed` (default)	−2.65 [−3.29, −2.28]	+3.67 [+2.35, +5.02]	+7.39 [+5.65, +9.05]	20 / 0
`pinv_open`	−0.84 [−1.31, +0.28]	+6.15 [+4.30, +7.20]	+11.50 [+10.04, +13.14]	20 / 0
`buzz_match`	−2.24 [−2.90, −1.75]	+3.87 [+2.68, +5.13]	+8.15 [+5.38, +8.92]	20 / 0
`min_drive`	−1.63 [−2.13, −1.20]	+2.63 [+0.85, +3.15]	+2.61 [+2.19, +3.78]	18 / 1
`shape_svd`	−0.84 [−1.13, −0.64]	+0.39 [−0.27, +0.87]	−0.13 [−0.89, +0.22]	5 / 10
`band_average`	−2.12 [−2.55, −1.68]	+1.10 [+0.34, +1.83]	+0.39 [−0.42, +0.70]	12 / 7

With six shakers, the setup with six control gauges read 2.65 dB lower at the control gauges than the setup with twelve, and 3.67 dB higher at the held-out gauges and 7.39 dB higher in stress, on all 20 structures (Table 5). The largest rise in stress was on structure 8, +16.0 dB. The other three laws that invert the whole response moved the same way in stress, by +2.6 to +11.5 dB; the open-loop form by the most. For the law that limits its drive in the weak directions (-0.13 dB) and the band-average law (+0.39 dB, quartiles −0.42 to +0.70) the stress did not move clearly.

With twelve shakers, twelve control gauges against twenty-four: the stress error rose by +9.40 [+7.59, +10.57] for the default law, +10.26, +9.80 and +2.71 for the other three that invert the whole response, each on 18 to 20 of 20 structures, and by +0.28 and +0.12 for the law that limits its weak directions and the band-average law.

Audit: holds. On the ten held-back structures the default law's six-gauge setup read 2.40 dB lower at the control gauges (study 2.65) and 8.71 dB higher in stress (study 7.39), on all ten. Its held-out error rose on all ten by more than on the study's structures: 5.29 dB against 3.67 (the buzz-test law likewise, 5.22 against 3.87). The open-loop form's stress rose by 10.26 dB (study 11.50), on all ten. With twelve shakers the default law's stress rose by 9.35 dB (study 9.40), on all ten. For the law that limits its weak directions and the band-average law the held-back structures agree with "no clear move": −0.54 and +0.41 dB with six shakers (6 and 7 of 10 in the study's direction), −0.24 and +0.07 dB with twelve.

3.3 Tuned Settings

On six shakers and six control gauges, the first ten structures were also run with each law at its tuned setting. For the two forms of the pseudo-inverse, the laws that invert the whole response with and without the loop correction, the tuned setting limits the drive in the weak directions (Section 2.4). It was chosen on two other structures by the held-out-gauge error, without the stress.

Table 6: Tuned setting minus default, six shakers on six control gauges, same structure (dB, median [quartiles] over 10 structures). Linear structures. The columns are the control-gauge, held-out-gauge and stress errors, and the structures down in stress.

Law	Control	Held-out	Stress	Down in stress
Default law (`pinv_closed`)	+0.23 [−0.29, +0.38]	−3.05 [−3.32, −2.57]	−6.39 [−8.41, −4.50]	10 of 10
Open-loop form (`pinv_open`)	−0.88 [−1.25, +0.64]	−5.03 [−5.50, −3.16]	−11.57 [−12.32, −9.26]	10 of 10

For the default law, the two tests are alike at the control gauges: the control-gauge error changed by +0.23 dB, 5 structures up and 4 down. In stress they are far apart: the tuned setting lowered the stress error by 6.39 dB on all ten structures (largest, structure 8, −12.3 dB). The held-out gauges fell by 3.05 dB and showed the difference. An engineer comparing the two runs by their control gauges alone would have had no reason to prefer one. For the open-loop form the stress fell by 11.57 dB on all ten.

On six shakers and twelve control gauges, the tuned setting chosen for that setup changed the default law's stress error by +0.15 dB, within the spread: the gain is on the setup with as many shakers as control gauges. For the other four laws the tuned setting moved the stress error by less than 0.8 dB in median on either setup.

Audit: holds. On the ten held-back structures the tuned setting lowered the default law's stress error by 6.30 dB (study 6.39), on all ten, and left its control-gauge error unchanged (-0.07 dB, 5 of 10 each way), as on the study's structures. Its held-out error fell on all ten, by more than on the study's structures: 3.41 dB against 3.05. For the open-loop form the stress fell on all ten by less than on the study's structures: 9.09 dB against 11.57.

3.4 Which Measurement Warns

An engineer has no stress truth. What could have told them a test was over-stressing the article? Figure 3 plots each run's stress error against its control-gauge error and its held-out-gauge error, and Table 7 gives the rank correlations.

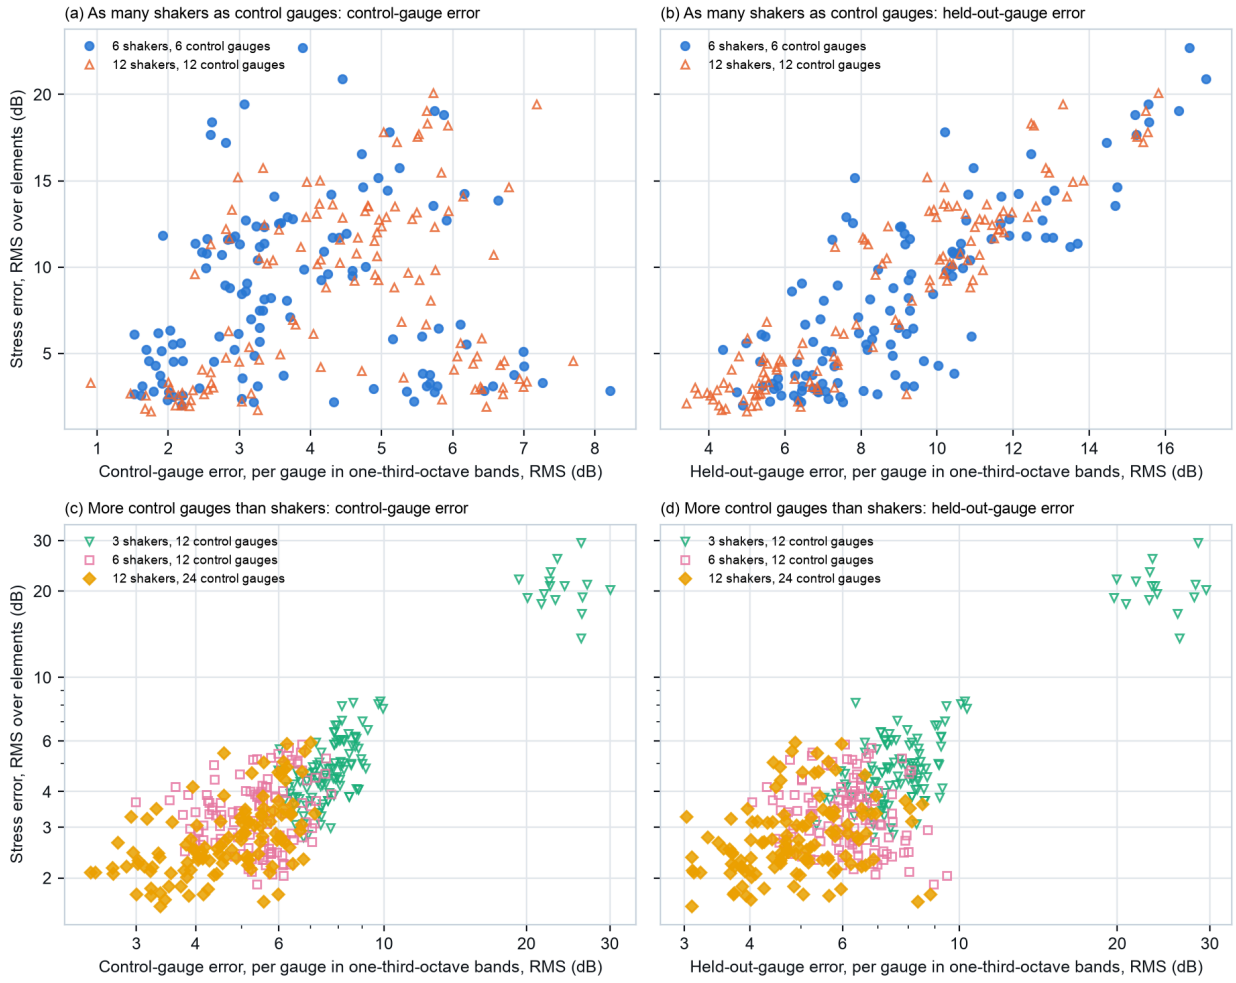


Figure 3: Stress error against the control-gauge error (left) and the held-out-gauge error (right), one marker per run, six laws pooled. Top: the setups with as many shakers as control gauges (240 runs). Bottom: the setups with more control gauges than shakers (358 runs), on logarithmic axes. Linear structures, default settings.

Table 7: Spearman rank correlation of each gauge error with the stress error, over runs, six laws pooled, on the study's structures and on the held-back structures. "Equal" is the setups with as many shakers as control gauges, "more gauges" those with more control gauges than shakers. Linear structures, default settings.

Setups	Runs	Control, per gauge	Held-out, per gauge	Control, summed	Held-out, summed
Equal: study	240	+0.25	+0.85	−0.09	+0.91
Equal: held back	120	+0.27	+0.81	−0.16	+0.88
More gauges: study	358	+0.72	+0.55	+0.67	+0.63
More gauges: held back	179	+0.63	+0.71	+0.64	+0.68

With as many shakers as control gauges, the held-out gauges tracked the stress error (+0.85 per gauge, +0.91 summed) and the control gauges did not (+0.25 per gauge, −0.09 summed). With more control gauges than shakers, the two groups tracked alike: the control gauges a little better on the study's structures, the held-out gauges a little better on the held-back ones. These are the patterns the first article found under idealized

control (its Section 4.3, for the summed spectra: 0.81 at the controls and 0.79 held out with unequal counts, and −0.08 to −0.16 at the controls under the inverse law once square systems were pooled in).

Audit: holds. The test fixed beforehand was that, with as many shakers as control gauges, the held-out gauges' correlation exceed the control gauges' by at least 0.30 on the held-back structures. It did, by 0.54 per gauge (+0.81 against +0.27), against 0.60 on the study's structures.

Pooled correlations mostly separate setups and laws. Within one setup they are weaker: on 6 shakers and 6 control gauges +0.17 at the control gauges and +0.78 at the held-out gauges, on 12 and 12 +0.31 and +0.89, but on 6 shakers and 12 control gauges +0.14 and −0.17, where every law is within a few decibels of the truth and no gauge sorts the runs. The first article said the same of its own correlations within one setup.

The rank correlation does not say whether a gauge warns in the right direction. Table 8 counts, for each sign an engineer could read, how many of the over-testing runs it caught and how many other runs it also flagged. An over-testing run here is one whose stress mean is more than 3 dB over service.

Table 8: Which signs caught the over-testing runs on the setups with as many shakers as control gauges: 240 runs, 150 of them with the stress mean more than 3 dB over service. Linear structures, default settings, six laws pooled. Each sign is read from the gauges.

Sign	Over-tests caught	Others flagged
Control pattern (line 2) graded C or worse: summed control error over 4.5 dB	1 of 150	1 of 90
Held-out pattern (line 2) graded C or worse: summed held-out error over 4.5 dB	144 of 150	22 of 90
Level (line 1) outside 1.5 dB: mean error over the control and held-out gauges and bands	141 of 150	12 of 90
Control gauges read more than 1.5 dB high on average	73 of 150	1 of 90
Held-out gauges read more than 1.5 dB high on average	148 of 150	16 of 90
Tip gauges: response spectrum at Q of 10 more than 1.5 dB high on average	146 of 150	13 of 90
Over-test alarm (line 3): any band at any gauge more than 6 dB over	150 of 150	90 of 90

On these setups the held-out gauges, read for their pattern or for their average level, caught nearly every over-test, as did the gauges at the component tips. The control gauges read high on average in half of the over-testing runs, and their pattern grade caught one. The alarm caught every over-test and flagged every other run as well (Section 3.5).

With more control gauges than shakers, 21 of 358 runs over-tested by more than 3 dB on average. Sixteen of them are one law on the setup with three shakers: the law that chooses the cross terms needing the least drive. In those runs the control and held-out gauges read 4 to 21 dB low on average while the stress was 11 to 29 dB high: the law put its drive into narrow lines that over-stressed the article while the band averages at the gauges read low. Every gauge measure flagged those runs as far from their targets, but the average level pointed the wrong way. The control law study reports that law's runs on that setup.

The counts in Table 8 and the within-setup correlations were not among the comparisons fixed before the audit, and the audit did not test them.

3.5 Report Card Lines

The scorecard of each run is written by the benchmark's scorer (Cross 2026e), which applies the report card of the first article. Before reading its lines as evidence about the proposed practice, the scorer was checked against that article's text.

- **Line 1, the level.** The scorer takes the mean decibel error over the control and held-out gauges and their bands; the first article says "all gauges", and the scorer leaves out the tip gauges.
- **Line 2, the pattern.** The first article grades the RMS decibel error of the summed spectra in one-third-octave bands, at the held-out gauges and at the control gauges, A to 3 dB, B to 4.5 dB, C to 6 dB. The scorer computes the same.
- **Line 3, the over-test alarm.** The first article raises an alarm at +6 dB in a one-third-octave band at any control or held-out gauge, and reports the share of bands over it. The scorer computes the same.
- **Line 6, the exact-fit flag.** The first article flags "a near-zero control error with shakers equal to controls" and gives no number. The scorer flags a run when there are at least as many shakers as control gauges and the per-gauge control error is under 1.0 dB. The 1.0 dB is the benchmark article's choice, not the first article's. The scorer therefore applies this line with a threshold the practice does not state, and its count of flags says nothing about the practice.

Table 9: Report card lines as the scorer applied them: the pattern grades of line 2 at the control and held-out gauges, stress within 3 dB RMS, the exact-fit flag (at the scorer's 1.0 dB) and the over-test alarm. Linear structures, default settings, six laws pooled, on the study's structures and on the held-back structures. "Equal" is the setups with as many shakers as control gauges, "more gauges" those with more control gauges than shakers.

	Equal: study	Equal: held back	More gauges: study	More gauges: held back
Runs	240	120	358	179
Control A or B	238	120	304	159
Held-out A or B	74	43	253	138
Stress within 3 dB	41	17	129	57
Exact-fit flag	1	0	0	0
Alarm	240	120	358	179

With as many shakers as control gauges, the control pattern graded A or B in 238 of 240 runs, while the stress error was within 3 dB in 41. The held-out pattern graded A or B in 74.

What can be said of line 6 without the scorer's threshold is what the control error was. On the setups with as many shakers as control gauges, the per-gauge control error was at least 0.91 dB in every run, with a median of 3.90 dB; it was under 2 dB in 20 of 240 runs and under 3 dB in 73. The summed control error had a median of 1.72 dB. In a closed loop with noise, the control error did not come near zero, while the over-test that an exact fit causes happened in 150 of the 240 runs. A flag keyed to a near-zero control error has little to fire on in such runs, whatever its threshold.

The over-test alarm fired in all 598 runs on the linear structures, and in all 1,797 runs at default settings over the three stages. The worst band of a run was at least 9.4 dB over its target, with a median of 22.8 dB. It fired on every one of the 170 linear runs whose stress error was within 3 dB. The first article foresaw part of this: it did not propose a hard stop on every gauge, because the worst control band ran well over even when the article as a whole was close. The share of bands over the alarm, which the line also reports, was higher on the setups with as many shakers as control gauges (median 20.7 percent) than on the others (10.2 percent).

Audit: holds. The test fixed beforehand was that the flag fire in under 5 percent of the runs with as many shakers as control gauges and the alarm in over 95 percent of all runs. On the held-back structures the flag fired in none of 120 runs and the alarm in all 299, and the control pattern graded A or B in all 120 runs with as many shakers as control gauges while the stress was within 3 dB in 17.

4. Nonlinear Structures

The same comparisons were made at stage 2, at both sizes. None of them was among the comparisons fixed before the audit, and the audit did not test them.

Table 10: The setup comparisons at stage 2, default law (dB, median over 20 structures, or 10 for the tuned setting; quartiles for the stress). The columns are the default law's control-gauge, held-out-gauge and stress errors, and the structures up and down in stress. "6 x 12" is 6 shakers and 12 control gauges, and so on. Every law: `data/setup_analysis.txt`.

Comparison	Stage	Control	Held-out	Stress	Up / down
6 x 6 minus 6 x 12	realistic	−1.28	+2.52	+5.48 [+3.69, +6.91]	20 / 0
	strong	−1.16	+1.87	+4.78 [+3.06, +7.92]	20 / 0
12 x 12 minus 12 x 24	realistic	−0.28	+4.30	+6.69 [+5.81, +8.38]	20 / 0
	strong	+0.78	+3.54	+9.29 [+6.89, +10.82]	20 / 0
12 x 12 minus 6 x 12	realistic	−0.33	+3.39	+6.90 [+5.15, +8.66]	20 / 0
	strong	+0.28	+2.41	+8.42 [+5.99, +10.11]	20 / 0
6 x 6, tuned minus default	realistic	−0.74	−2.20	−3.65 [−6.18, −2.96]	0 / 10
	strong	−0.67	−1.46	−3.11 [−3.76, −2.48]	0 / 10

The setup effect held at both sizes for the laws that invert the whole response: every one of their comparisons raised the stress error on 19 or 20 of 20 structures. With nonlinearity, the control-gauge error no longer always fell as the stress rose: for twelve shakers on twelve control gauges at the strong size it rose a little (+0.78 dB), and for the open-loop form of the default law it rose by 1.2 to 3.9 dB in every comparison. For the law that limits its weak directions and the band-average law, the stress changes were within about 0.9 dB in median, the largest being the band-average law on six shakers and six control gauges at the strong size (+0.84 [+0.08, +2.24] dB, 14 structures up and 4 down).

Table 11: Rank correlation of the per-gauge error with the stress error, and report card counts, at stage 2: rank correlations in the rows for the control and held-out gauges, counts of runs in the others. "Equal" is the setups with as many shakers as control gauges, "more gauges" those with more control gauges than shakers. Six laws pooled, default settings.

	Equal: realistic	Equal: strong	More gauges: realistic	More gauges: strong
Runs	240	240	359	360
Control gauges	+0.50	+0.53	+0.67	+0.23
Held-out gauges	+0.83	+0.75	+0.55	+0.27
Control A or B	222	194	305	306
Stress within 3 dB	26	20	128	42
Exact-fit flag	0	0	0	0

With as many shakers as control gauges the held-out gauges still tracked the stress better than the control gauges, by a smaller margin than on the linear structures (0.33 and 0.22, against 0.60). With more control gauges

than shakers at the strong size, neither group tracked the stress well (+0.23 and +0.27). The over-test alarm fired on every run at both sizes.

5. Measuring Again

A structure that changes with level moves away from the frequency response measured at low level. Measuring it again during the test is one answer. On the first ten structures, at both stage 2 sizes, each law ran with the response measured again and blended into the old one each loop.

Table 12: Frequency response measured again during the test minus measured once and kept, default law, same structure (dB, median over 10 structures; quartiles for the stress). The columns are the default law's control-gauge, held-out-gauge and stress errors, and the structures up and down in stress. "6 x 12" is 6 shakers and 12 control gauges, and so on. Every law: `data/setup_analysis.txt`.

Setup	Stage	Control	Held-out	Stress	Up / down
6 x 6	realistic	+1.73	+2.16	+4.85 [+4.63, +5.14]	10 / 0
	strong	+1.61	+1.81	+3.75 [+2.39, +4.89]	10 / 0
6 x 12	realistic	+0.10	+0.11	+0.23 [+0.07, +0.45]	5 / 0
	strong	+0.15	+0.43	+1.00 [+0.46, +1.64]	9 / 0

For the default law, measuring again made the test worse, and much worse on the setup with as many shakers as control gauges: +4.85 dB of stress error at the realistic size, on all ten structures, against +0.23 dB on the setup with twelve control gauges. The same was true of the open-loop form and the buzz-test form of the inverse. Here the control gauges did show it: their error rose by 1.6 to 1.7 dB with the stress. For the law that limits its weak directions and the band-average law, the change was within 0.2 dB in median on both setups. The audit did not run these cases.

6. What the Checks Show

The first article's Section 8 asked for noise, identification error, a closed loop and articles that change with level, and called its ten structures thin. Within one family of frame structures, this study made those checks, and the findings read against the first article as follows.

- **The failure does not need an exact fit.** Under idealized control the first article's controls matched to 0.09 dB. In a closed loop with noise the control gauges read 3 to 5 dB from their targets on the setups with as many shakers as control gauges, far from an exact fit, and the stress was still about 10 to 13 dB over on average for three of the four laws that invert the whole response (4 to 5 dB for the fourth). The comparisons of setups that show it held on the held-back structures.
- **The setup decides it.** Adding shakers to the same control gauges until there were as many of each raised the stress error by 8.8 dB for the default law (8.5 dB held back); a setting that limits the drive in the weak directions removed most of it while leaving the control gauges where they were (6.4 dB, 6.3 dB held back).
- **The held-out gauges are the measurement that warns.** With as many shakers as control gauges they tracked the stress, and the control gauges did not, on the study's structures and on the held-back ones.
- **It depends on the kind of law.** The laws that limit their drive in the weak directions or work on band averages did not over-test on these setups; capping the inverse or using a band law is what P2 of the first article proposes for this case.
- **Two report card lines, as the scorer applies them, did not separate good runs from bad on these structures,** on the study's structures or the held-back ones.

What this suggests for the next version of the proposed practice is not settled here. The practice is revised in its own article.

7. Limits

- **Simulation only.** Every run is inside one computer. The shakers are ideal point forces with no dynamics of their own, and the gauges are perfect apart from the added noise. Nothing here has met hardware.
- **The truth comes from the same model as the gauges.** The stress truth is computed from the same modal model that produces the gauge readings, so the gauges and the truth share every simplification of the model.
- **One family of structures.** All 30 structures, study and held back, are frames of one design family, loaded in service through their mounts. A second family, or an article loaded by distributed pressure, could behave differently.
- **The number of structures.** Twenty for the comparisons at default settings and ten held back; ten for the tuned setting and for measuring the response again. Quartiles over ten structures are coarse.
- **One run per case.** Each case was run once with one noise record. How much a result moves from one run to the next of the same case is not measured, so the spread given is over structures only.
- **The default checks.** Most runs would have stopped early at the controller's default checks. The results describe each law let run to the end with the checks widened, not what a laboratory running the defaults would see.
- **Different targets.** The band-average law runs the banded specification and the other laws the full line-by-line spectra and cross spectra, so a comparison between them is of each law with its own form of target.
- **Setups that change more than one thing.** The comparisons of six against twelve control gauges, and of twelve against twenty-four, also move the shakers. Only the comparison of 3, 6 and 12 shakers on the same 12 control gauges changes one thing.
- **The default law's loop.** The findings on adding shakers and on tuning rest on the default law, whose shared closed loop has two known limits (primer Section 6): where the measured frequency response is off near a mode, a gauge can be held high, because a cross term may not grow while its gauge reads high; and a term at zero stays at zero. A line-by-line port of the published correction algorithm holds a gauge high in the same way when the target has cross terms. The open-loop form, which has no such loop, showed the same over-test (stress +9.1 dB with twelve shakers against six, +11.5 dB with six control gauges against twelve, and its tuned setting lowered the stress by 11.6 dB), so these findings do not rest on those limits; no run was rerun for them.
- **What the audit did not check.** The audit tested the five comparisons of the Summary. It did not run the cases that measure the frequency response again, and it did not test the results at stage 2, the counts of warning signs or the correlations within one setup.
- **Small differences.** Differences under 0.2 dB are not claimed.

Contribution Statement

Kevin Cross led the work, set its direction and is responsible for the article.

How This Was Made

The study code, the analysis, the audit, the figures and the drafting of the text were done with Claude Code, an AI coding assistant running Anthropic's Claude Opus models, under Kevin Cross's direction. The author is responsible for every word, number and line of code.

Acknowledgments

The observation this article checks was first made on hardware by Rohe, Nelson and Schultz of Sandia National Laboratories (Rohe et al. 2019): “A smaller error on the control gauge responses does not necessarily indicate a better test.” The regularization of the inverse (Schultz (Schultz 2019)) and the case for trading “a minor amount of control response accuracy” for it (Schultz and Rohe (Schultz and Rohe 2026)) are published work of Sandia National Laboratories, credited here to Sandia National Laboratories. The band-average law is Brandon Zwink’s.

References

Cross, Kevin. 2026a. *Hertzforge Open Benchmark Set, MIMO Random Suite: Results of the Control Law Study and Its Independent Audit, Version 0.2*. Data set, Zenodo. https://doi.org/10.5281/zenodo.24310746.

Cross, Kevin. 2026b. *Hertzforge Open Benchmark Set: MIMO Random Suite, Version 0.2*. Data set, Zenodo. https://doi.org/10.5281/zenodo.24310742.

Cross, Kevin. 2026c. “Six Control Laws for Multi-Shaker Random Vibration: How Each Scores on the Open Benchmark Set.” *Hertzforge Journal* 1: 2026–09. https://doi.org/10.5281/zenodo.24311107.

Cross, Kevin. 2026d. “Six Control Laws for Multi-Shaker Random Vibration: What Each Does and How It Is Built.” *Hertzforge Journal* 1: 2026–07. https://doi.org/10.5281/zenodo.24291981.

Cross, Kevin. 2026e. “The Open Benchmark Set: Public Test Problems, a Fair Run and a Scorecard for Multi-Shaker Control Laws.” *Hertzforge Journal* 1: 2026–08. https://doi.org/10.5281/zenodo.24311105.

Cross, Kevin, and Brandon Zwink. 2026. “What Is a Good Multi-Shaker Test? A Proposed Practice for Judging MIMO Vibration and Shock Tests on Flexible Articles.” *Hertzforge Journal* 1: 2026–01. https://doi.org/10.5281/zenodo.24252422.

Rohe, Daniel P., Garrett D. Nelson, and Ryan A. Schultz. 2019. “Strategies for Shaker Placement for Impedance-Matched Multi-Axis Testing.” *Sensors and Instrumentation, Aircraft/Aerospace, Energy Harvesting and Dynamic Environments Testing, Volume 7, Proceedings of IMAC XXXVII*, 195–212. https://doi.org/10.1007/978-3-030-12676-6_18.

Schultz, Ryan A. 2019. “A Demonstration of Force Estimation and Regularization Methods for Multi-Shaker Testing.” *Sensors and Instrumentation, Aircraft/Aerospace, Energy Harvesting and Dynamic Environments Testing, Volume 7, Proceedings of IMAC XXXVII*, 229–43. https://doi.org/10.1007/978-3-030-12676-6_21.

Schultz, Ryan, and Daniel P. Rohe. 2026. “Understanding 6DOF Test Methods and Their Impacts on Results.” *IMAC XLIV, International Modal Analysis Conference, Palm Springs, California*. https://www.osti.gov/biblio/17086461.

Implemented In

Nothing in this article is implemented in a Hertzforge release yet.

Version History

Version 1, 2026-11-17. DOI 10.5281/zenodo.24311109. First published.

Reviews. This article is reviewed in public. Its reviews, its review status and any later version are at https://hertzforge.com/journal/2026-010/

Copyright 2026 Kevin Cross. Text, figures and data CC BY 4.0; code MIT.

Methods are published for research. Anyone who applies one to a test is responsible for protecting the equipment and the article under test.