

A shape test for a threshold transition

Luc Côté — Zenodo deposit, DOI 10.5281/zenodo.22801278 *Theory of the Intersection of Fields*

First publication of this protocol, before any execution: version of 14 September 2026, DOI 10.5281/zenodo.22715013. That version fixes the protocol's priority over its results. Later versions of the deposit, including this one, do not alter this date.

Protocol, to be frozen before looking at the velocities

Status: frozen 12 September 2026. Both lists of cuts are cited section by section, both population assumptions are settled, and the two blemishes noted after the first freeze are corrected in this same pass. **Item 8c was corrected on 12 September, forced by the data:** the three-dimensional separation does not exist in the files, the line moves to the projected separation, and the count is a lower bound. The protocol is complete. **No modification after this date.**

Why this protocol rather than another

Wide-binary gravity tests estimate a parameter: by how much the observed gravity exceeds the Newtonian one. That is parameter estimation, and it cannot distinguish laws that predict the same mean by different routes.

A single global factor together with a free semi-major axis per system does not belong to this test: that model class does not contain laws whose ratio varies with separation, and the detail of that verification is treated elsewhere. The protocol below deliberately departs from it.

It measures no amplitude. It measures where a transition occurs, and whether it occurs at the same place for stars of different masses.

The idea in one page

A threshold law places its transition at a radius that depends on the mass of the pair. So does a continuous interpolation: both scales, the threshold radius and the MOND radius, grow as the square root of the mass. **That dependence therefore separates nothing.**

What separates them is the coefficient. A threshold transition places the step considerably earlier — at 0.32 times the MOND radius in the case examined here.

Hence the variable of the protocol:

$$\xi = s / \sqrt{M_{\text{tot}}} \quad [s \text{ in AU, } M_{\text{tot}} \text{ in solar masses}]$$

Every pair, whatever its mass, then lands on the same axis. The unit of ξ is the astronomical unit per square root of a solar mass.

And the value of the knee is predicted, not sought:

law	predicted ξ_*
threshold transition (front)	2 266 AU·M$\odot$^(-1/2)
continuous interpolation (MOND radius)	7 030 AU·M$\odot$^(-1/2)

The two are a factor 3.10 apart. **These two numbers are written before the computation and must not be adjusted afterwards.** The acceptance window is set here at $\pm 20\%$: an empirical knee outside [1 813 ; 2 719] does not recover the threshold transition, and outside [5 624 ; 8 436] does not recover the continuous interpolation.

The prediction is therefore: if the transition is a threshold, the profiles of the three mass bins superpose in ξ and the knee falls at 2 266. If it is a continuous rise, the knee falls at 7 030. If the profiles superpose as well in s as in ξ , neither is supported.

This is a test of location, not of amplitude. No scale factor is estimated.

But that does not make it unassailable. If mass correlates with separation — which catalogue selection can produce — a variation with separation of the unresolved-companion fraction or of the eccentricity distribution can manufacture a knee in ξ . Hence a mandatory cross-check: **plot a multiplicity indicator — RUWE, photocentric offset, astrometric goodness of fit — against ξ and against s .** If that indicator also shows a knee at ξ_* , gravity has not been isolated and the test does not conclude.

The observable

For each pair:

$$\tilde{v} = v_p / \sqrt{(G M_{\text{tot}} / s)}$$

the ratio of the projected relative velocity to the Newtonian circular velocity at the projected separation. It is dimensionless and requires no orbit model.

Warning: its Newtonian value is not 1. Projection and eccentricities lower it — for a thermal distribution and isotropic viewing, the median is **0.55**, the mean 0.565, the 75th percentile 0.76 and the 90th 0.94. These values depend on the cuts and on the eccentricity distribution adopted. The published quality framework normalises its Newtonian prediction to 0.5, in agreement with this.

All comparisons are therefore made against the Newtonian level of the pipeline itself, measured in the control window $\xi < 1\,500$ on the same draw and with the same cuts, and never against 1. Normalising against 1 would falsify the whole control zone.

The five steps

1. Freeze the cuts before looking

Two sets of cuts exist in the literature and reach opposite conclusions on the Newtonian regime: Chae's, and those of the quality framework of Cookson et al. **Both must be run in parallel and published together.**

If a transition appears only under one of the two sets, that is not a result about gravity — it is a result about the cuts. This must be sayable.

The cuts are frozen, with the code and a timestamp, **before** examining the velocities of $\xi > 1500$.

2. Rescale in ξ , over at least three mass bins

Split the sample into at least three bins of total mass, and plot the profile of $\tilde{v}$ against ξ for each.

Threshold prediction: the three profiles superpose, and the knee falls at the same ξ^* .

Continuous-interpolation prediction: no common knee; the profiles superpose as well, or as poorly, in s as in ξ .

Mandatory control: repeat the analysis with the masses mixed. The structure must dilute. If it survives mixing, it does not come from the law — this is the control that protects against an artefact of binning.

3. Describe first, confront second — two distinct steps

3a. Description, with no law. Plot the observed profiles of $\tilde{v}$ against ξ , by mass bin, and the same against s . Do they superpose better in ξ than in s , and where does the knee fall? This step uses no model and answers only the question “is there a common knee”.

Do not compare the median alone: compare **the full empirical distribution function** of $\tilde{v}$ in each ξ interval. **And look at the high quantiles — P75, P90 — before the median:** a step shows there sooner, because the most radial orbits cross the threshold first and populate the tail before shifting the centre.

3b. Confrontation with the laws. The common knee, if it exists, does not say which of the laws produces it. That requires **synthetic populations**, one per law confronted, passing through the same projection, the same cuts, the same eccentricity distribution and the same companion model. The observed empirical distributions are then compared with those, by a Kolmogorov–Smirnov test or a χ^2 on the quantiles.

Only two laws enter at equal standing — Newton and the threshold transition — for the reason given at item 9: the numerical solution of the two-body problem is beyond our reach, and the continuous interpolation enters through its one-body form, with an indicative verdict.

Without those populations, step 3a answers “is there a knee”, not “which law”. The two questions must not be conflated.

3bis. Two sets of population assumptions, and a second stopping rule

The eccentricity and companion distributions are not measured. Two sets are therefore frozen, and **both are run**:

	eccentricity distribution	unresolved-companion fraction
set A	thermal, $p(e) = 2e$	constant with separation — the implicit assumption of those who remove companions by cuts rather than modelling them
set B	$p(e) \propto e^\alpha$, α calibrated within $\xi < 1\,500$ only	$f = \min(1; 0.2 + 0.5 \log_{10}(s / 0.2 \text{ kAU}))$, the form published by Chae (§IV.2)

Both sets apply to both sets of cuts, that is four combinations, all run and all reported.

Why both, and not a compromise. The treatment of unresolved companions is precisely where the field divides: those concluding in favour of Newtonian behaviour remove them by cuts and model none; the one maintaining the anomaly models them with a fraction growing with separation. **It is that choice, not the data, that separates the two published conclusions.** A protocol retaining only one would settle the controversy by decree. Both are therefore run, and the stopping rule decides whether the result survives.

Neither assumption is borrowed from the article the set B cuts come from: that work removes companions by selection and publishes no fraction, no period distribution and no mass ratio. The varying form is cited from Chae, for both sets of cuts.

Second stopping rule, as binding as the one on cuts: if the ranking of the three laws changes between set A and set B, the protocol stops and the result is declared inconclusive. The ranking must survive both sets of assumptions to be reported.

4. The angle between separation and velocity

The angle between projected separation and projected velocity constrains the eccentricity distribution almost without passing through gravity. Almost: the time density along the orbit depends on the force, and a law that makes the orbit precess by tens of degrees per revolution deforms that distribution differently from a mere rescaling.

Condition for this test to be clean: calibrate the eccentricity distribution in the control window, $\xi < 1\,500$, and forbid it from changing in the threshold window. That is an assumption about formation, and it must be declared as such.

If the angle changes where $\xi = \xi_*$ while the distribution calibrated at small separations forbids it, this is the observable closest to a geometric measurement. If it does not change, nothing is lost.

Three windows in ξ , not in separation

A zone defined in raw separation separates nothing. The front falls at 1 603 AU for $0.5 M_{\odot}$ and at 3 583 AU for $2.5 M_{\odot}$: a 1–3 kAU window is already beyond the threshold for light pairs and still inside it for heavy ones. It mixes the two regimes precisely where they are meant to be told apart.

The three windows of the test are therefore defined in ξ :

window	bounds in ξ ($\text{AU} \cdot M_{\odot}^{-1/2}$)	what is expected there
control	$\xi < 1\,500$	inside the threshold for every mass. The ratio sits at the pipeline's Newtonian level under a threshold transition, and 3 to 8 % above it under a continuous interpolation.
discriminant	$1\,800 < \xi < 3\,500$	the threshold transition makes its rise here; the continuous interpolation has barely begun.
late	$5\,500 < \xi < 8\,500$	the continuous interpolation makes its rise here; the threshold transition is already on its plateau.

The control window is where the three hypotheses are not collinear, and it holds the most pairs:

observation in the control window	favours
excess of 3 to 8 % above the pipeline's Newtonian level	continuous interpolation
flatness at the pipeline's Newtonian level	Newton or threshold transition, against the interpolation

The pipeline's Newtonian level is measured **in that same control window**, on the same cut sample and with the same eccentricity distribution. It is of order 0.55 in the median, not 1.

A fifth observable, which tests no shape but requires no adjustment

A pair whose relative velocity exceeds the Newtonian escape velocity at its separation is, under Newton, either a contaminant or a system coming apart. Under enhanced gravity it can be bound.

The count of such systems depends directly on the plateau height. A law predicting 1.82 binds more of them than a law predicting 1.40: the escape line moves as the square root of the enhancement, that is 35 % higher in velocity in the first case against 18 % in the second.

This is a **count**, not an estimate. No factor is adjusted, no orbit distribution enters: one counts the pairs above a line, and the line is set by the law.

What it tests and does not test. It tests the height of the plateau, not the shape of the transition — so it does not replace what precedes, it adds to it. And it remains vulnerable to contamination: an unresolved companion produces exactly the same velocity excess as enhanced gravity. The count is therefore informative only if the contaminant fraction is bounded independently, which returns to the unsettled disagreement of the field.

What nevertheless makes it useful: it pre-registers. The prediction is written before looking — so many pairs expected above the line under Newton, so many under a continuous interpolation, so many under a threshold transition — and DR4 checks it. In the thirty-six-system sample, four pairs exceed the Newtonian line: that is not a statistic, it is a quantity to freeze now so that it may become one later.

What the protocol will not do

It will not arbitrate the controversy. The analyses concluding in favour of Newtonian behaviour and those maintaining an anomaly work on the same data with different cuts, and their disagreement is unsettled. The result will therefore be conditional: *under such a set of cuts, such a hypothesis is excluded*.

It will not give a three-sigma verdict on shape. Separating heights takes of order 150 to 300 clean pairs in a narrow window; separating shapes takes more, with the mass resolved and the companion fraction under control in the same window. Current catalogues, after severe cuts, do not have that count.

It will not neutralise unresolved companions. The disagreement over how their fraction varies with separation is itself unsettled, and any step between 2 and 5 kAU has a non-gravitational explanation available as long as it lasts.

And orbital periods will remain out of reach, before as after DR4. This test will remain a population test.

What must not be done

Do not project with Keplerian orbits. A threshold law makes the orbit precess by tens of degrees per revolution; the classical orbital elements do not represent it. Draw the radii with the time density of the actual potential — $dt \propto dr/|v_r|$ — and integrate.

And for the comparison interpolation, use the numerical solution of the two-body problem rather than a one-body formula: the gap between the two reaches 28 % for equal masses. Without that, one compares three caricatures.

On the thirty-six-system sample

It cannot carry this test, and that must be said.

Too few objects on either side of their own threshold, which depends on their mass. The semi-major axis is not observed. Four systems already exceed the Newtonian escape velocity, so that a higher law binds them and a lower law makes them contaminants: any likelihood comparison there mixes gravity with membership.

What it still serves for: classifying each system, at its mass, as inside or outside its threshold — a classification, not an estimate. And counting the systems above the Newtonian escape line, which tests a height, not a shape. Four out of thirty-six is not a statistic; it is a quantity to pre-register for DR4.

The price of moving to the large sample: one loses the three-dimensional kinematics, hence any individual constraint on inclination and radial velocity. One gains two orders of magnitude in number — the only thing that allows a step to be seen rather than averaged away.

Executable specification

What precedes fixes the approach; what follows fixes the numbers. **Everything written here is frozen and timestamped before opening the velocities of $\xi > 1\,500$** , and must not move afterwards. Anything left blank would be improvised during the computation, that is, chosen on seeing the result.

1. The two samples. Common parent catalogue: El-Badry, Rix & Heintz (2021), Gaia eDR3, with its chance-alignment probability R estimated per pair. Both sets start from there.

	set A — Chae	set B — quality framework
reference	Chae 2023, ApJ 952, 128 (arXiv 2305.04613), §II.1 for all cuts below	Cookson et al. 2026, MNRAS 547, stag342 (arXiv 2602.24035)
pair type	BinType = MSMS	main sequence, plus HR filter
heliocentric distance	d < 200 pc (extended sample; reference sample at d < 80 pc)	d < 130 pc (illustrative figures at 150 pc)
chance alignment	R < 0.01	$\tilde{v} < 2.5$ and $\Delta v_{\text{sky}} < 10$ km/s
projected separation	0.2 < s < 30 kAU	1 < s < 30 kAU
relative proper-motion error	< 0.01 (strict variant 0.005)	full propagation, cut on the uncertainty of $\tilde{v}$
absolute magnitude	4 < M_G < 14 (strict variant 4 < M_G < 12)	HR filter on the Belokurov and Penoyre region, least affected by unresolved

	set A — Chae	set B — quality framework
		binaries
RUWE	explicit cut, sample median ≈ 1.03	RUWE < 1.25 on both components
astrometric multiplicity degrouping	—	ipd_frac_multi_peak = 0
	—	yes, multiple systems separated before analysis
radial velocity	not required (main sample)	maximum difference imposed where available

Verification carried out 12 September 2026, section by section, in Cookson et al. 2026 (arXiv 2602.24035):

cut	section	what the article says
$\tilde{v} < 2.5$	§3.1.1	the Newtonian limit $\sqrt{2}$ rises only to ≈ 2.5 for a general class of non-Newtonian theories; the cut therefore loses no signal
maximum radial-velocity difference	§3.1.2	—
separation bounds	§3.1.3	—
degrouping	§3.2.1	multiple systems separated before analysis
null values, noisy stars	§3.2.2	—
RUWE < 1.25	§3.2.3	recommended threshold for DR3; the optimum varies from 1.15 to 1.37 with sky position, mean 1.25
HR filter	§3.2.4	region least affected by unresolved binaries
heliocentric distance	§3.2.5	counts per distance bin stop rising around 60–70 pc and stay flat out to 150
ipd_frac_multi_peak = 0	§3.3	quality check
spherical projection correction	§4.1.6	see below — an analysis step, not a selection

The set A values are read in Chae 2023, **§II.1**, which carries all the cuts above. That article imposes no explicit RUWE cut; it reports the sample median, ≈ 1.03 . **Both sets of cuts are therefore now cited.**

An analysis step the protocol overlooked, and which is not a cut. The quality framework imposes a **spherical projection correction** (§4.1.6) and shows that it becomes important **precisely when crossing the transition radius** — where the test is decided. It concludes that a mishandled or omitted correction can manufacture a signal where the data are Newtonian. **This correction is therefore mandatory in both sets**, applied identically, and its omission is to be reported as a control variant, not as an option.

And a confirmation of the Newtonian level: the article recalls that the median of $\tilde{v}$ is **about 0.5**. The independent calculation reported above gives 0.55 for a thermal distribution. The two agree, and confirm that normalising against 1 would be wrong.

Record the fingerprint of the file obtained for each of the two. Without that, two divergent catalogues will be recreated askew, and a difference of query will be attributed to gravity.

Expected sizes, and why this matters: set B is markedly more restrictive — 130 pc against 200, minimum separation 1 kAU against 0.2, RUWE capped, astrometric multiplicity excluded. Its sample will be smaller by a substantial factor. This is expected, and it is precisely the reason to run both: they differ not only in number, **they differ in conclusion on the Newtonian regime**, one finding the anomaly and the other not.

2. Formulas and their units. ξ , $\tilde{v}$, the reference Newtonian velocity, with units written out — separation in AU, mass in solar masses, ξ in $\text{AU}\cdot\text{M}\odot^{(-1/2)}$.

3. The predicted values and their tolerances.

	predicted $\xi_\star$	acceptance window ($\pm 20\%$)
threshold transition	2 266 $\text{AU}\cdot\text{M}\odot^{(-1/2)}$	1 813 – 2 719
continuous interpolation	7 030 $\text{AU}\cdot\text{M}\odot^{(-1/2)}$	5 624 – 8 436

The two windows do not overlap: their ratio is 3.10 for a 20 % tolerance. An empirical knee therefore falls in one, in the other, or in neither — never in both.

4. The binning, frozen.

Mass bins — at the **tertiles of the total mass of the cut sample**, computed before opening the velocities of $\xi > 1\,500$, and reported with their numerical bounds. **The tertiles are computed separately on each of the two sets of cuts:** two triplets of bounds, not one, since the two samples do not share the same mass distribution. Choosing tertiles rather than absolute bounds guarantees three comparable counts whatever the set of cuts; it is arbitrary, which is why it is frozen.

ξ *intervals* — constant width in logarithm, **twelve intervals** covering 500 to 20 000 $\text{AU}\cdot\text{M}\odot^{(-1/2)}$, that is a factor 1.36 per interval. This range brackets both predicted values, 2 266 and 7 030, with four intervals on either side of each.

Minimum count — **100 pairs per cell** (ξ interval $\times$ mass bin). Below that, the cell is reported as empty and enters no test. This threshold is the one below which a 90th-percentile quantile ceases to mean anything.

Merging rule, written in advance — under the more restrictive set of cuts, many cells will be empty. **Two adjacent ξ intervals both below the minimum count are merged, at most once**, and the merge is reported. If after one merge the cell remains below 100, **that set of cuts is declared insufficient for the shape test** and only its control window is reported. This rule is written now precisely so that it is not chosen on seeing the counts.

3bis-ter. The knee detector, as a single formula.

“The knee falls at 2 266” means nothing to a program until the knee is defined. One estimator is retained, and it is written here:

ξ_{emp} is the abscissa minimising the χ^2 of a **two-plateau profile with a single node** — low level free, high level free, node position free in [500; 20 000] — fitted to the median of $\tilde{v}$ per ξ interval, weighted by the cell count.

Verdict: threshold transition if $\xi_{\text{emp}} \in [1\,813; 2\,719]$; continuous interpolation if $\xi_{\text{emp}} \in [5\,624; 8\,436]$; neither otherwise.

One estimator only is tried. Testing several and keeping the one that gives the cleanest result is forbidden — it is the commonest way of manufacturing a knee.

5. The two sets of population assumptions, as defined in step 3bis. For each, the companion-model recipe is written out: fraction, distributions of periods and of mass ratios of the hidden companion, and how it contributes to the measured velocity. Set A holds the fraction constant with separation; set B lets it vary, following the published form. **Both are run and reported together.**

6. The three orbit generators. Time density of the declared potential for the threshold transition, numerical two-body solution for the continuous interpolation, Kepler for Newton. None of the three may be replaced by a one-body formula.

7. The expected outputs, listed in advance: profiles of $\tilde{v}$ against ξ by mass bin; the same against s ; empirical distributions per cell; table of tests against the synthetic populations; mass-mixing control; cross-check of the multiplicity indicator; the control window ($\xi < 1\,500$) reported against the pipeline’s Newtonian level.

8. DATA APPENDIX — to be pasted before the first run, with no other modification. The three pieces below are missing, and they cannot be reconstructed: the articles must be opened and copied. **Until they are pasted: no code on $\xi > 1\,500$, no tertiles computed, no figure.** And once pasted, **no further modification of the protocol is permitted** — neither here nor elsewhere in this file.

The rest of this document is frozen at the date given at the head. These three pieces are data, not choices of method: pasting them reopens nothing. They cannot be filled from a reconstruction; the articles must be opened.

a) *The two lists of cuts. Settled 12 September.* Set B is cited section by section (Cookson et al. 2026, §3.1.1 to §4.1.6); set A is too (Chae 2023, §II.1, which carries all of its cuts).

b) *The companion model. Settled 12 September, and it was a choice of method, not a pasting of data.* The article the set B cuts come from publishes no fraction: it removes companions by selection. The “varying following the published form” of the previous version therefore had no source. It is replaced by the **two explicit assumptions of step 3bis** — constant fraction, and Chae’s §IV.2 form — applied to **both** sets of cuts, that is four combinations.

c) *The thirty-six Gaia identifiers, and the escape-velocity formula.* The count of pairs beyond the Newtonian line is a pre-registration only if it can be replayed. The list of identifiers published with the sample is therefore needed, and the exact line.

CORRECTION OF 12 SEPTEMBER, forced by the data and declared as such. The previous version fixed the line at the **three-dimensional** separation, asserting it was available for this sample. **It is not.** Checked on the files themselves: they carry only the **projected** separation and a parameter inferred by deprojection, which is not data. Reconstructed from the parallaxes, the radial component gives a median offset of 21 589 AU for a median uncertainty of 29 648 — **consistent with pure noise**, with twelve systems out of thirty-six having an offset smaller than its own error bar. The formula as written was therefore not computable.

The line is therefore fixed at the projected separation:

$$v_{\text{esc}} = \sqrt{(2 \text{ G M}_{\text{tot}} / s_{\text{proj}})}, \text{ compared with the measured relative velocity.}$$

This is not an arbitrary fallback, and the sign of the error is known. The true separation is always greater than or equal to the projected one, so the escape velocity computed on the projected separation is **overestimated** — by a median factor 1.075 for isotropic viewing (the median of $\sin \theta$ is $\sqrt{3}/2$; the mean, $\pi/4$, would give 1.128). The line sits too high, fewer systems cross it, and **the resulting count is a lower bound on the true count**. That is conservative, and it is reported as such.

Under gravity enhanced by a factor γ , that line is multiplied by $\sqrt{\gamma}$: **1.183 for a continuous interpolation at 1.40, and 1.348 for a threshold transition at 1.82.**

The pre-registration, with its numbers. Two series exist and **they do not combine**:

line	recomputed from our files (32 systems)	published column (36 systems)
Newtonian	5	6
× 1.183 (continuous interpolation at 1.40)	2	2
× 1.348 (threshold transition at 1.82)	0	0

The series retained is the left one: 5, 2, 0 over thirty-two systems. The criterion for a pre-registration is that it can be replayed, and this is the only one of the two entirely reconstructible from the files of the run. The published column uses masses we do not have —

the two series differ by a median ratio of 0.923 with a maximum offset of 0.286, which is not rounding noise.

Four systems are excluded, having no mass in that file. The one-unit difference on the first line comes from a single one of them, at 1.12 in the published column; the other three are 0.70, 0.45 and 0.45.

The zero must be read plainly: under the plateau at 1.82, no system in this sample becomes bound that would not already be bound under Newton. This is not a refutation — thirty-two objects, and a lower bound — but it is not favourable either, and it is written before DR4.

And the “four out of thirty-six” of the previous versions was wrong. It corresponds to no file of the run. It is replaced by the three numbers above, which are replayable.

9. Two limitations to accept or circumvent in advance.

The generator for the continuous interpolation — choice settled. The numerical solution of the two-body problem is not an evening's script. **The option retained, and it is frozen: the confrontation of step 3b covers Newton and the threshold transition.** The continuous interpolation enters through its analytic one-body form, **with the mandatory statement that this form overestimates the enhancement by up to 28 % for equal masses**, and its verdict is reported as indicative, not as a comparison of equal standing. No other option will be considered at computation time.

The angle test. Step 4 remains a comment until its statistic is written. It is written here: **a Kolmogorov–Smirnov test of the angle distribution per ξ interval against the one calibrated in the control window ($\xi < 1\,500$), with the same projection correction.**

10. The stopping rules, written as true-or-false conditions. First: the result differs between the two sets of cuts. Second: the ranking of the laws changes between the two sets of population assumptions. Third: the multiplicity indicator shows a knee at the same ξ^* . If any of the three is true, the protocol stops and the result is declared inconclusive — this is written before, not after.

Two biases of the protocol itself, to be named before reading the results

Set B is conditioned on the test observable. Its $\tilde{v} < 2.5$ cut acts on the very quantity being measured. It is faithful to the article it comes from, and justified there — but it is not neutral here: it truncates the tail of the distribution, precisely where a threshold transition and companion contamination produce similar effects. **Consequence to be written in advance: the 90th percentile of set B is structurally lower than that of set A.** A truncated tail must not be read as a failure of the threshold transition.

The two-plateau detector is harsh on the threshold transition. After projection, that law produces not a step function but a spread rise with a shoulder. The χ^2 node may therefore fall between 2 266 and 7 030 — verdict “neither” — even though the rise begins at the right ξ .

That verdict is accepted in advance. Adding, on the day of the computation, a second estimator of the “onset of slope” kind that would make the threshold transition easier to recognise is forbidden. The prohibition is already written above; it is repeated here because this is the moment when it will be most tempting.

What goes into the Letter, and what does not

Goes in: the abscissa ξ , the two predicted values of ξ^* , the two sets of cuts run in parallel, the stopping rules, and the fact that the thirty-six-system sample does not carry this test.

Does not go in: the detail of the injection into a global-factor model, which reopens a debate this protocol departs from; and any formulation claiming that nothing can absorb the effect, since selection and multiplicity can.