

Applying the Preisach Principle: A Field Protocol


📅 Sep 24, 2026 · @Danny

How to use this protocol

This protocol turns a new case into a candidate instance of the Preisach principle. It then tests whether the coupling in that instance earns its place against the independent null. It is a companion to *From Independent to Coupled Hysterons: The Preisach Principle* (the formal paper) and to the *Sticky Switches* illustration set, and assumes the paper's vocabulary and sign conventions throughout.

The input is any group of parties under a shared pressure. The output is one of three verdicts:

- **Not admissible.** A checkpoint in Step 1 or Step 2 failed. The case does not map onto the four elements of §3.4 without an extension, and no diagnostic is run.
- **Admissible, coupling not supported.** The case maps cleanly, but the signatures do not beat the null of §5.2 by the margin fixed in advance.
- **Admissible, coupling supported in the model.** The signatures beat the null with margin. This is still a statement about a hand-estimated instance, not about the world, until the operational definitions and independent coupling measurements of §5.2 exist.



Step 7 is numbered last because it is easiest to write once Steps 1–4 have fixed the vocabulary of the case. It must be written down, with its numbers, before Step 5 is run. A falsification condition drafted after seeing the transition graph is not a falsification condition.

Step	Output	Can stop the protocol?
1. Parties and binary state	Party list, operational state rule per party, coarsening log	Yes
2. Scalar drive	U with units, source, resolution, polarity σ_i per party	Yes
3. Thresholds	u_i^+ , u_i^- with uncertainty ranges; naive switching order	No, but flags inadmissible windows
4. Coupling	Signed, asymmetric C with an evidence grade per entry; frustration check	No
5. Diagnostic	Stability intervals, t-graph, four signature readouts	No
6. Null comparison	Signature table, coupled vs null, with margins	Decides the verdict
7. Falsification	Pre-registered failure conditions and pass margins	Decides the verdict

Step 1. Identify the parties and test the binary state

List every party that could plausibly commit under the pressure, then write, for each one, an **operational state rule**: the observable record that would show $s_i = 1$ rather than $s_i = 0$. "Has adopted the position" is not a rule. "Has published the policy on its website" is.

This step is a checkpoint, not a formality. Each party must pass all four tests below. A party that fails one either gets reformulated or takes the whole case out of scope.

1. **Two committed states, not a continuum.** Does the party's behaviour actually cluster into two states, or does it vary smoothly? If the natural variable is continuous (how much, how often), binarising it throws information away. That can be acceptable when the question is genuinely "whether" rather than "how much", but write down what was discarded in a coarsening log.

2. **Latching.** Once switched, does the party stay switched while the drive sits inside some range, without continued pushing? Look for a reversal cost: sunk investment, public commitment, reputational exposure. Without latching, $u_i^+ \approx u_i^-$ and the party is a threshold unit, not a hysteron. It can stay in the model, but it contributes no memory.
3. **Exactly two states.** Is there a third option (partial adoption, a workaround, a quiet exit) that parties actually use? If so, either fold it explicitly into one of the two states and log the choice, or stop: a three-state element needs a multi-state hysteron extension that the paper does not cover.
4. **Unit identity.** Is the party a single decider, or an aggregate whose members can split? An aggregate that splits in practice should become several hysterons, or be treated as a continuous sub-density (§3.1), not as one relay.

Outputs. A party table (index, name, state rule, which way it commits), a verdict per party (pass, pass with coarsening, fail), and the coarsening log.

Stop condition. If any party that matters to the question fails test 3 or test 4 and cannot be reformulated, the case is not admissible. Do not proceed by quietly treating it as binary.

Step 2. Identify the single scalar drive

Name the drive U , its units, where it would be measured, and its **resolution**: the smallest change in U that the record could actually distinguish. Resolution matters later, in Step 6, because a signature narrower than it cannot be observed.

Fix a direction convention. Parties that commit under decreasing U get polarity $\sigma_i = -1$ and use the local field $\sigma_i U$ in the stability condition, as in §4.1.

Incommensurability check. List every pressure the parties themselves cite, then ask:

1. **Do the pressures co-vary?** If, over the period of interest, they rise and fall together, a one-dimensional projection loses little. If they move independently, it loses a lot.
2. **Is there one exchange rate?** Try to say how much of pressure X is worth one unit of pressure Y . If that rate differs by party, there is no common scalar: each party is responding to a different projection of a vector drive.
3. **Does any party respond to a pressure the others ignore?** A private pressure cannot be folded into a shared U .

If two or three pressures survive this check as genuinely separate, the model does not apply directly. There are three honest options:

- **Vector drive.** Replace each window by a switching boundary in a two- or three-dimensional drive space. This is a real extension, with its own memory properties, and it is not covered by the paper.
- **Slow background.** If one pressure moves much more slowly than the others, treat it as a parameter that shifts thresholds between analysis windows, not as a drive. State the window over which it is assumed frozen.
- **Split the case.** Analyse sub-groups that do share a single drive separately.

Also check the quasistatic condition of §3.3: parties must switch fast compared with how fast U moves. If decisions take longer than the drive takes to cross a window, rate effects are outside the model.

Stop condition. More than one incommensurate pressure that cannot be frozen or split out means the case is not admissible without extension. Record which pressures failed and why.

Step 3. Estimate the thresholds

For each party, estimate a bare up threshold u_i^+ and a bare down threshold u_i^- in units of U , with $u_i^+ > u_i^-$. **This is a rough first pass, not fitted data.** Every value carries an uncertainty range, and the ranges are used again in Step 6.

Useful sources, roughly in order of strength:

- observed switch points in the party's own history;
- stated rules or commitments ("we review prices when costs rise 25%");
- analogous parties in comparable cases;
- judgement, labelled as such.

The bare-versus-effective trap. An observed switch point is an *effective* field $u_{i\pm}(S)$, which already includes whatever coupling was active at the time. To recover the bare field, record the configuration S at the moment of each observation. Leave the correction until Step 4 has estimated C , and do not use the same observation to estimate both the bare field and the coupling. That is the circularity §5.2 warns about.

Checks.

- Window width $u_i^+ - u_i^-$ is positive for every party, including at the ends of its uncertainty range. A range that lets the window close is a Step 1 latching problem in disguise.

- Record the **naive switching order**: the up order sorted by u_i^+ and the down order sorted by u_i^- . This is the prediction of the uncoupled model and the baseline for scrambling in Step 5.

Output. A threshold table: party, u_i^+ , u_i^- , range, source, and the configuration each estimate was taken in.

Step 4. Estimate the coupling matrix

Fill in C one ordered pair at a time. Row i is the party whose thresholds move; column j is the party whose commitment moves them. Keep the paper's convention from §2.1:

$$u_i^\pm(S) = u_i^\pm + \sum_{j \neq i} c_{ij} s_j$$

$c_{ij} < 0$ is **cooperative**: j being up makes i easier to switch up and harder to switch down. $c_{ij} > 0$ is **antagonistic**.

For each pair (i, j) , record four things:

- Whether** j 's commitment changes i 's thresholds at all. The default is zero. Every non-zero entry needs a stated mechanism, not just a plausible story.
- Direction.** c_{ij} and c_{ji} are estimated separately. **The matrix may be asymmetric**, and often is: a market leader moves a follower much more than the reverse. Never fill c_{ji} by copying c_{ij} .
- Sign**, cooperative or antagonistic, with the mechanism in one sentence.
- Magnitude** in units of U , with a range, and an **evidence grade**: M (mechanism only), E (observed episode), P (independent perturbation, the §5.2 standard).

Checks.

- Frustration.** Relabel any $\sigma = -1$ parties so that every party responds to $+U$, which flips the sign of every coupling that crosses the polarity boundary. Then look for directed loops with an odd number of antagonistic links. With none, all couplings are effectively cooperative, no-passing holds, and return-point memory is guaranteed (§2.3). In that case RPM violation cannot be one of your signatures, and Step 7 must say so.

- **Equal shift.** The form above shifts both of i's thresholds by the same amount. If the mechanism plausibly moves only one of them, the window width becomes state-dependent and admissibility condition 1 of §3.3 can fail. Flag it rather than forcing it into the equal-shift form.
- **Circularity.** No entry may be estimated from the same transitions that Step 5 will use as evidence of coupling.

Output. The matrix, a table of non-zero entries (i, j, value, range, sign, mechanism, grade), and the frustration verdict.

Step 5. Run the diagnostic

With the thresholds and C fixed, compute the transition graph and read off four signatures. First state the **update rule** for simultaneous instabilities (the paper's default is most-unstable-first) and the driving assumption (quasistatic). Both go in the write-up, because the t-graph can depend on them (§3.3).

5a. Stability intervals. For each of the 2^N configurations S, compute

$$U^-(S) = \max_{i: s_i=1} \sigma_i u_i^-(S), \quad U^+(S) = \min_{i: s_i=0} \sigma_i u_i^+(S)$$

(with the σ_i handling of §4.1 for reversed-polarity parties). S is stable for U in $(U^-(S), U^+(S))$ when that interval is non-empty. Tabulate all of them, including stable configurations that turn out to be unreachable.

5b. Transition graph. Start from saturation (all parties at the state they take for very low U) and follow the up-edge and down-edge out of each reached configuration. At each edge, relax under the update rule and record every flip in order. For N up to about 12, enumerate; beyond that, sample.

5c. Read off the four signatures.

1. **Avalanches.** Count transitions that flip more than one party at a single U. Also note near-avalanches: separate transitions that coupling has pulled within one resolution step of each other.
2. **Scrambled order.** Compare the order in which parties switch, on the way up and on the way down, with the naive order from Step 3. The strong form is a party whose switch point depends on which other parties are up, so that it switches at different U on different histories.

3. **Coupling-only mixed states.** List reachable configurations that the $C = 0$ transition graph never reaches from saturation.
4. **Return-point memory.** Run reversal histories: saturation, then two or three alternating nested extrema, then an inner sub-loop, then a return to each earlier turning point. RPM holds if every return lands on the configuration recorded at the first visit. Record the grid step, the range of U , the history depth and the number of histories. A coarse grid can miss violations that live in narrow windows.

Output. The interval table, the t -graph (a diagram for small N), and a signature table with counts. Nothing in this step is a verdict yet.

Step 6. Compare against the null model

This is the null-model test of §5.2 applied to one instance. The claim of coupling is only as strong as its margin over the independent Preisach model, so run two nulls, not one.

- **Same-threshold null.** Rerun Step 5 with $C = 0$ and the Step 3 thresholds unchanged. This shows what the coupling adds to *these* parameters.
- **Refit null.** Ask whether an independent model with *different* thresholds could reproduce the same observations. This is the fairer test. An uncoupled model has one up order and one down order, but the two can differ, so a single sweep's release order can be matched by re-sorting the u_i^- . Nearly coincident thresholds can also imitate an avalanche to within resolution. What no refit null can produce is a party that switches at **different U on different histories**, a violation of return-point memory, or non-congruent minor loops (§1.3).

A signature counts toward the coupling claim only if it clears every one of these margins:

1. **Absent in the null.** Present with C , absent at $C = 0$, and not reproducible by a refit null.
2. **Robust to the estimates.** Present in at least a pre-set share of draws (for example 80%) when thresholds and couplings are jittered across their Step 3 and Step 4 ranges.
3. **Wider than resolution.** The window of U in which the signature lives is at least twice the resolution of U from Step 2.
4. **Insensitive to the update rule.** It survives a switch from most-unstable-first to sequential or random updating (§5.1).

Where there is real data, fit both models to the same transitions and compare them by a criterion that penalises the extra parameters in C. Without data, as in every first pass, this step can only say whether the instance *would* be distinguishable from the null, and at what measurement resolution.

Output. A table with one row per signature: null count, coupled count, jitter survival, window width against resolution, and a pass or fail against each margin.

Step 7. State the falsification conditions up front

Write these down, with numbers, after Step 4 and before Step 5 is run. They say which results would mean the coupling story does not fit this case. Fill in the bracketed values for the case at hand.

- **F1. No avalanches.** No transition flips more than one party at a single U, in more than [20]% of jitter draws. Coupling is too weak to change behaviour, and the instance is the classical model with extra parameters.
- **F2. No scrambling.** Every party switches at the same U, in the same order, on every tested history, as the refit null predicts. There is no state-dependence for coupling to explain.
- **F3. Memory holds perfectly.** Return-point memory holds on every history at grid step [$\leq$ half the resolution of U], across the jitter range, in an instance Step 4 found frustrated. The antagonistic entries are then doing no observable work. Together with F1 and F2, this means the coupling story does not fit. On its own, it removes the antagonistic part of the story. It does not remove the cooperative part, since cooperative coupling preserves RPM anyway.
- **F4. Signatures below resolution or fragile.** Every signature that passes F1 to F3 lives in a window of U narrower than [$2 \times$ resolution], or vanishes in more than [20]% of jitter draws, or flips with the update rule. The coupling may exist in the model but would never be seen.
- **F5. Admissibility.** Any checkpoint failure in Step 1 or Step 2 that was waved through. This is not a finding about coupling. It means the instance was never valid.

With data, add a sixth condition: **F6.** An observed transition that the fitted t-graph forbids, such as a party switching when its effective threshold was not crossed.

The verdict rule is fixed in advance too. Coupling is supported only if at least one signature that no refit null can reproduce (history-dependent switch points, or an RPM violation) survives F3 and F4. Avalanches and single-sweep scrambling alone are not

enough.

Worked example: three cafés and rising input costs

The protocol ends in **admissible, coupling not supported** for this case. The coupling produces an avalanche, a scrambled release order and genuine return-point-memory violations. But every signature an independent model cannot copy lives in a window of drive five times narrower than the drive can be measured. The case is invented and the numbers are hand-set, as in §4 of the paper. Every result below was computed from them, not narrated.

Three cafés share one street. Wholesale coffee and milk costs are rising, and each café must decide whether to raise its menu prices.

Step 1. Parties and binary state

i	Party	State rule: $s_i = 1$ when	Verdict
A	Chain franchise; prices set by its regional office	The posted price of a standard latte is above its baseline	Pass. The decider is the regional office, not the shop; logged.
B	Independent bakery-café	Same rule	Pass with coarsening. Raises vary from 20 to 40 cents; only "whether" is kept.
C	Discount coffee kiosk	Same rule, or a smaller cup at the same price	Pass with coarsening. Shrinking the cup is a third option, folded into $s = 1$ as an effective price rise; logged.

Latching is plausible for all three. Reprinting menus costs money, and customers punish visible flip-flopping, so a café that has raised does not cut again at the first dip in costs. Had C's cup-shrinking been a separate, commonly used state rather than a disguised price rise, test 3 would have failed and the case would have stopped here.

Step 2. Scalar drive

The cafés cite four pressures: bean costs, milk costs, rent and wages.

- **Beans and milk** arrive on the same supplier invoice and co-move closely. They are combined into one input-cost index.
- **Rent** matters only to B, whose lease renews next year. It is a private pressure, so it is frozen as slow background and assumed constant over the analysis window.
- **Wages** rise on one fixed date for all three cafés. That is a step change in background, not a drive, and the analysis window is chosen to avoid it.

The drive is **U = the rise in the input-cost index since baseline, in units of 10%**, so $U = 2.5$ means costs are up 25%. The supplier index moves in 5% steps, so the **resolution is 0.5**. All three cafés commit under rising U , so $\sigma_i = +1$ for all. Prices are decided within about a week, and the index moves monthly, so the quasistatic condition holds. The checkpoint passes, but only because two pressures were frozen by assumption. That goes in the limitations.

Step 3. Thresholds (rough first pass, not fitted)

i	u_i^+	u_i^-	Range	Source
A	2.5	1.0	± 0.15	Stated regional policy: review prices at +25% costs
B	1.5	0.5	± 0.15	Owner's estimate of margin headroom
C	3.5	2.0	± 0.15	Judgement: the kiosk's model is to hold the lowest price

All windows stay open across their ranges. The naive order from the uncoupled model is **up B, A, C** and **down C, A, B**: last in, first out.

Step 4. Coupling

Entry	Value	Sign	Mechanism	Grade
c_BA	-0.6	Cooperative	The chain raising gives the bakery cover to follow	M
c_CA	-1.2	Cooperative	Price umbrella: once the chain is dearer, the kiosk has room to raise	M
c_AB	-0.3	Cooperative	The bakery raising gives the chain a little cover	M
c_AC	+0.8	Antagonistic	Once the kiosk raises, the chain's office runs promotions to win its customers	M

All other entries are zero, and every magnitude carries a $\pm 30\%$ range. The matrix is asymmetric: c_CA and c_AC have opposite signs. The loop $A \rightarrow C \rightarrow A$ has one antagonistic link, so **the instance is frustrated**, and return-point memory is not guaranteed. Every entry is grade M. None has been measured independently.

Step 7, filed before Step 5


The conditions were fixed before anything was computed. Jitter uses 200 draws, with couplings varied by $\pm 30\%$ and thresholds by ± 0.15 . A signature must survive in at least 80% of draws and occupy a window of at least 1.0 (twice the 0.5 resolution). RPM is tested on a grid no coarser than 0.25. The verdict rule is the one in Step 7: at least one history-dependent signature must clear F3 and F4.

Step 5. Diagnostic

Update rule: most-unstable-first, with quasistatic driving. States are written s_A s_B s_C.

S	Stable for U in, coupled	Stable for U in, C = 0
000	$(-\infty, 1.5)$	$(-\infty, 1.5)$
010	(0.5, 2.2)	(0.5, 2.5)
110	(0.7, 2.3)	(1.0, 3.5)
111	$(1.5, \infty)$	$(2.0, \infty)$
011	(2.0, 3.0), unreachable	(2.0, 2.5), unreachable
100	empty	(1.0, 1.5), unreachable

001, 101	empty	empty
----------	-------	-------



Parenteses list the parties that flip, in order. For comparison, the $C = 0$ graph is a plain ladder: up B at 1.5, A at 2.5, C at 3.5, and down C at 2.0, A at 1.0, B at 0.5.

- 1. Avalanches: present.** From all-raised, lowering costs to $U = 1.5$ drops A first. The chain losing its antagonistic push from the kiosk is what tips it. With A down, the price umbrella is gone and C's down field jumps from 0.8 to 2.0, so C follows at the same U. That is 1 of 6 transitions, against 0 of 6 at $C = 0$. There is also a near-avalanche on the way up: C's up field falls from 3.5 to 2.3, so C now raises 0.1 after A. That gap is below resolution.
- 2. Scrambled order: present.** Uncoupled, the release order is C, then A, then B. Coupled, A releases *before* C. The strong form is present too: A's down field is 1.5 when C is up and 0.7 when C is down. A therefore releases at different U on different histories. Sweeping c_{AC} shows the reversal switches on at $c_{AC} = 0.1$ exactly, where A's effective down field in 111 ($0.7 + c_{AC}$) overtakes C's (0.8).
- 3. Coupling-only mixed states: absent.** Both graphs reach the same four configurations: 000, 010, 110 and 111. Coupling removes 100 from the stable set but adds nothing new. This signature fails outright.
- 4. Return-point memory: violated.** Take saturation, then $U = 2.5$, then 0.75, then 2.25, then back to 0.75. The first visit to 0.75 lands in 010. The return lands in 110. The reason is that on the way up, 110 exists only for U in $[2.2, 2.3)$. The avalanche on the way down skips 110 entirely. So a sub-loop that turns inside that window comes back to a configuration the main loop never passed through. Counts: 44 of 4,600 return checks fail on a 0.25 grid, and 224 of 71,980 on a 0.1 grid. At $C = 0$, none fail.

No configuration ever has two parties unstable at once, so the update rule plays no role in this t-graph. That was checked by inspection.

Step 6. Null comparison

The **refit null** is an independent model with A at (2.2, 1.5), B at (1.5, 0.5) and C at (2.3, 1.45). It reproduces the full up-and-down sweep to within resolution, avalanche and reversed release order included. Across all 287,980 histories of the form $\text{saturation} \rightarrow e_1 \rightarrow e_2 \rightarrow e_3 \rightarrow e_2 \rightarrow e_1$, on a 0.05 grid, the refit null and the coupled model disagree only when a turning point falls in [2.2, 2.3). That is the real coupling signature. They also disagree in [1.45, 1.5], which is an artefact of the refit's own offset.

Signature	C = 0	Coupled	Refit null copies it?	Jitter survival	Window vs 1.0 needed	Counts toward coupling?
Avalanche	0 of 6	1 of 6	Yes, to within resolution	200 of 200	Single sweep	No
Reversed release order (single sweep)	C, A, B	A, C, B	Yes	200 of 200	Single sweep	No
History-dependent release point for A	No	1.5 vs 0.7	No	Tied to the 110 window	0.1	No: fails F4
Coupling-only mixed state	None	None	Not applicable	Not applicable	Not applicable	No: absent
RPM violation	0	224 of 71,980	No	157 of 200 (79%)	0.1	No: fails F4

Verdict

F1, F2 and F3 are not triggered. F4 is. The coupling is not idle: it produces an avalanche, a reversed release order and real RPM violations. But the two signatures a refit null cannot copy both live in a 0.1-wide window of U, which is 1% of input cost. The record resolves only 5% steps. RPM violation also falls just short of the 80% jitter bar.

Under the rule filed in advance, the verdict is **admissible, coupling not supported**. With monthly invoices, the three cafés would look exactly like three independent hystérons with slightly different thresholds.

Two things could change the verdict. The first is a drive measured at 0.05 resolution (0.5% cost steps), which would make the window observable. The second, which is better, is an independent measurement of c_AC. For example, check whether the chain's promotions actually track the kiosk's price changes, which would upgrade the entry from grade M to grade P.

Limitations

The protocol makes a fitting narrative easy to produce. Steps 1 to 4 take any case with a few parties and a pressure, and turn it into a matrix, a t-graph and a set of signatures. Each step invites a mechanism sentence, and a mechanism sentence can almost always be written. Four grade-M couplings and three judgement thresholds were enough to produce avalanches and memory violations in the worked example. The Step 6 margins and the Step 7 conditions filed in advance are the only parts of the protocol that push back. They work only if they are written before the diagnostic and honoured afterwards.

Admissibility is partly assumed. In the example, two of four pressures were frozen by assumption, and a third behaviour (cup-shrinking) was folded into a binary state by decision. Each such choice is defensible, and each one quietly narrows what the result can mean.

The numerical checks are limited. They use one update rule, grids of 0.25, 0.1 and 0.05, histories with at most three extrema, and 200 jitter draws. As in §5.3 of the paper, finer grids or longer histories could change the counts.

Moving from illustrative to real takes three things that no amount of narrative can supply:

- operational definitions of each state and of U, fixed before the data are examined;
- a measured drive with a stated resolution;
- couplings measured independently of the transitions they are meant to explain, ideally by perturbing one party and recording the threshold shift in another (§5.2).

Until those exist, an application of this protocol is an illustration of what the Preisach principle *would* say about a case. It is not evidence that the case is described by it. We offer the protocol as a working procedure to be tightened by use, not as a validated method.