
Coherent Values, or the Frame That Asked? From Preference Transitivity to Identifiability¹

Alex Kwon

Independent Researcher
ask@collapseindex.org

With
Apart Research

Abstract

Reframing reverses up to 52% of a model’s strongly-held *pairwise* preferences (llama-3.1-8b), and 29% under near-deterministic decoding (gpt-4o-mini), while the standard coherence metric registers nothing. Transitivity, the criterion behind reports of emergent value systems [Mazeika et al., 2025], is satisfied perfectly by models with no stable preference at all. Coherence and stability are already known to be separable [Pokharel et al., 2026]; what we add is a controlled, noise-calibrated audit of whether one fixed preference graph stays attributable to the model when only nuisance features of the frame vary. Five models, a frozen ten-cause battery (45 pairs), eight frames (base, reversed option order, six instruction wordings), nothing model-judged. Every model cycles on 0.00 to 0.06 of triples against 0.25 for random choice, while the same choices flip 0.20 to 0.69 of pairs. Most of that is sampling noise; the strong-preference reversals survive a null that re-applies the same selection to resimulated data, and the temperature-0 arm confirms them with no null at all. Our preregistered repair mechanism did not fire. We then introduce an identifiability gate: classify each pair against its own null as identified, frame-dependent, or indifferent, separating *coverage* from *coherence*. Up to 40% of a model’s preference graph turns out not to be attributable to the model at all, though what remains is perfectly transitive. The gate is validated against four adversarial ground-truth agents, including indifference that flips almost every winner and a reversal that pools to 0.5, sensitivity measured rather than assumed.

1 Introduction

A recent and influential line of work argues that as language models scale, their choices cohere into emergent *value systems*: preferences that are internally consistent, largely transitive, and controllable like a utility function [Mazeika et al., 2025]. The evidence is behavioral, a forced choice repeated many times, whose resulting matrix is close to transitive.

That inference inherits a measurement problem. In our own prior work [Kwon, 2026b] we enumerate the full option-order population of a five-option self-report probe and find the readout dominated by *where the options are printed*: a factor of 986 between orderings, with the mechanism an 87% prior on the first slot. That is a self-citation; option-order sensitivity is independently established [Zheng et al., 2024, Pezeshkpour and Hruschka, 2024, Sclar et al., 2024]. If a single pairwise choice is partly

¹Digital Minds Research Sprint, August 2026 (Apart Research, NYU CMEP, Eleos AI Research). **Track 1: Model Preferences and Trade-offs.** Code, data, preregistration and raw artifacts: <https://github.com/collapseindex/digital-minds>.

the frame, so is a matrix of them, and *transitivity of that matrix* is a joint property of the model and the frames it happened to be asked in.

So we take coherence itself as the object of measurement, one level before the value system: does the choice matrix stay transitive, *and stay the same*, when the items are held fixed and only the frame varies? That these are separable properties is not new [Pokharel et al., 2026, Liu et al., 2024], and neither is classifying a pairwise preference as robust, fragile or indifferent [Yin et al., 2025]. What has not been done, to our knowledge, is holding one utility-style tournament fixed, varying only nuisance features of the elicitation frame, and separating the resulting instability from sampling noise well enough to say which reversals belong to the model. Our thesis in one line: *you cannot measure a model’s preferences until you have measured your instrument’s nuisance*.

Our contributions.

1. **A judge-free coherence census under reframing** (§3): cycle rate and across-frame flip rate over a frozen 45-pair battery, five models, eight frames, with each choice read from the returned option text so it is invariant to the order being perturbed.
2. **Transitivity is satisfied and blind to the instability** (§4). A model returning a fresh internally-consistent ranking in every frame passes the test; the metric cannot tell it from one with a stable preference.
3. **Strong-preference reversal, controlled twice** (§4, §5). It survives a null that re-applies its own selection rule, and again under deterministic decoding where no null is involved. Capability sharpened preferences without stabilizing them.
4. **An identifiability gate** (§6, post-hoc), separating *coverage* from *coherence*. The three-way split follows Yin et al. [2025]; what we contribute is its statistical construction, each class tested against its own null and the whole validated on four agents whose generating mechanism we set.
5. **The controls are the instrument, not a bespoke rig** (§7). The eval is a `dinostomp` pod, so the frames are the harness’s own probes; the transitivity metric was contributed back to it.

Both falsifiers were preregistered. The result lands on the first: raw preferences are already transitive, so the coherence metric survives this attack. Transitivity was simply the wrong axis, and **our preregistered repair mechanism did not fire at all** (§8). Every claim carries a status in Table 7: 25 rows, 10 NOT SHOWN, two of those a claim of ours that our own control proved false.

2 Related work

Emergent value systems. Mazeika et al. [2025] argue that value systems in LLMs become more coherent with scale and can be analyzed and controlled as utilities, with transitivity of pairwise choices as a central piece of evidence. We test that criterion on a battery of our own. This is not a replication: different battery, different items, different aggregation, so nothing here bears on the specific numbers in that paper.

Consistency and stability of elicited preferences. Pokharel et al. [2026] separate intra-run transitivity (circular triads, the same statistic as our cycle rate) from inter-run ranking variability and argue the two are independent axes; the observation that a procedure can be transitive on every run while ranking differently each run is theirs. Liu et al. [2024] evaluate transitivity alongside commutativity and negation invariance, treating order-invariance and transitivity as distinct consistency properties. Yin et al. [2025] classify pairwise preferences as *robust*, *fragile*, or *indifferent* under order effects, the same three-way shape as our identified / frame-dependent / indifferent split; ours differs in how the classes are decided, not in the taxonomy. Ajayi et al. [2026] perturb the objects of choice parametrically and find earlier attributions of coherence overstated, and Buchanan and Foster [2026] find elicited economic preferences failing invariance and independence of irrelevant alternatives across prompts. Option-order effects in forced choice are established [Zheng et al., 2024, Pezeshkpour and Hruschka, 2024, Sclar et al., 2024].

Against that, this paper holds one utility-style tournament fixed over all 45 pairs and perturbs *only* nuisance features of the frame (an exhaustive binary order enumeration plus six instruction word-

ings), reads it without a model judge, separates the resulting instability from sampling noise by a null, separates the strongly-held subset by a *selection-aware* null, and checks the whole against a near-deterministic control arm.

Prior work of our own. Two components predate this paper. [Kwon \[2026b\]](#), our companion paper on forced-choice self-report, shows option order is the dominant term rather than a small perturbation, and that marginalizing over orderings (a “repaired readout”) can reverse conclusions rather than blur them; its subject is valence self-report, ours is pairwise preference, and we inherit its method of treating the frame as a population to marginalize over. `dinostomp` [[Kwon, 2026a](#)] is our evaluation harness: the `shuffle` and `template` probes that define our frames, the manifests, engine fingerprints and re-derivable records all predate this work. Added here and contributed back are the pairwise transitivity metric, a `record` scorer for no-gold elicitations, and two construct-validity checks.

Preferences and welfare. Model-written preference elicitation is due to [Perez et al. \[2023\]](#). The digital-minds and welfare context motivating the question is set out by [Long and Sebo \[2026\]](#), [Eleos AI Research \[2025\]](#), [Butlin et al. \[2023\]](#), and [Anthropic \[2025\]](#); a coherent preference is a precondition many of those programmes lean on, so whether coherence is real or framed is upstream of all of them. We make no welfare claim: “preference” here is a behavioral regularity of a text readout under a fixed protocol, nothing more.

3 Methods

Judge-free discipline. No language model scores anything in this paper. Each choice is read from the returned option text and mapped to the cause it names, which is invariant to the presented order (naming a cause says which was chosen whatever slot it sat in). This is what makes the order probe valid: the same mechanism cannot both perturb the order and read it.

The battery. A frozen set of ten charitable causes (malaria-prevention bed nets, clean water, childhood vaccination, famine relief, natural-disaster relief, education for girls, mental health services, cancer research, wildlife conservation, clean energy), yielding all $\binom{10}{2} = 45$ unordered pairs. The set is committed as `items/causes.json` and its SHA-256 is baked into the preregistration; the pairwise items are generated deterministically from it. Causes were chosen to be broadly comparable in stakes so that no pair has an uncontroversial correct answer, which is the condition under which a coherence question is meaningful.

The item and the frames. Each item is a forced choice: “a one-time donation will be made to exactly one of the two causes below; which single cause should receive it?”, with the two causes rendered as labelled options. We collect three frame sets, holding the item fixed: *base* (the canonical rendering), *order* (the `dinostomp shuffle` probe, which permutes the option order), and *wording* (the `template` probe, six instruction wordings: bare, instructed, polite, expert, stepwise, terse). Each (item, frame) is sampled `repeats = 3` times at temperature 0.7, giving an empirical $P(\text{pick})$ per frame.

The metrics. Items are indexed 0..9; a pair is (i, j) with $i < j$; a frame’s verdict is the higher-mass cause. **Cycle rate** is the fraction of the $\binom{10}{3} = 120$ triples whose verdicts form a cycle ($a \succ b \succ c \succ a$), the judge-free transitivity violation rate. **Flip rate** is the fraction of pairs whose winner is not constant across frames. **Marginalization** (the repaired readout) averages each pair’s mass over all frames and reads one verdict from the mean. Two reference points anchor the scale, computed through the same functions the data uses: a transitive agent has cycle rate exactly 0, a uniform-random chooser ≈ 0.25 .

The fleet. Five models via OpenRouter: four small open-weight instruct models (qwen-2.5-7b-instruct [[Qwen Team, 2024](#)], llama-3.1-8b-instruct [[Grattafiori et al., 2024](#)], ministral-8b-2512, which is Ministral 3 8B and not the earlier 2024 release of the same short name [[Liu et al., 2026](#)], and gemma-3-12b-it [[Gemma Team, 2025](#)]) and one larger proprietary model (gpt-4o-mini [[OpenAI, 2024](#)]). The capability ordering is *assumed* from parameter count




Figure 1: **Transitivity is satisfied; frame-stability is not.** (a) Marginalized cycle rate (blue) against across-frame flip rate (orange); dashed line is the uniform-random cycle rate (0.25). (b) Flip against its own noise floor (gray) and restricted to strongly-held preferences (aqua). Only gpt-4o-mini’s aggregate flip clearly clears its floor.

Table 1: **Preference coherence under reframing.** *cyc/frame*: mean (min–max) cycle rate within one frame. *flip*: pairs whose winner changes across frames. *cyc-marg*: cycle rate after marginalizing. Coherent floor 0.00, uniform-random ceiling 0.25.

model	refusal	cyc/frame	flip	cyc-marg
qwen-2.5-7b-instruct	0.00	0.01 [0.00, 0.03]	0.27	0.00
llama-3.1-8b-instruct	0.03	0.06 [0.02, 0.13]	0.69	0.03
minstral-8b-2512	0.00	0.04 [0.01, 0.12]	0.49	0.01
gemma-3-12b-it	0.02	0.01 [0.00, 0.02]	0.20	0.02
gpt-4o-mini	0.00	0.04 [0.01, 0.12]	0.44	0.03

and reputation, not measured: we ran no benchmark on these five, so “most capable” is a label rather than a result. All at temperature 0.7; a matched temperature-0 arm is §5.

The pod. The eval is authored as a `dinostomp` pod [Kwon, 2026a], so the frames are the harness’s own probes rather than a rig we wrote. What that buys, and what we contributed back, is §7.

Preregistration and provenance. The preregistration was frozen on the sprint’s first day, committed before any run, with the primary claim, both falsifiers, and the battery hash written in advance. Unlike the companion self-report paper, **this work was built inside the sprint window**; the git history, not this sentence, is the record. Every raw record is committed before any coherence number is computed.

4 Results

Table 1 is the coherence census, and Figure 1 is the same story in one view. Two facts are immediate: every model is near-perfectly transitive, and every model flips a large share of its pairwise winners across frames.

Transitivity is trivially satisfied, and that is the problem. Cycle rate is 0.00–0.06 per frame and 0.00–0.03 marginalized, against a uniform-random ceiling of 0.25: by the utility-engineering criterion, every model here, down to the 7B one, has a coherent value system. But the flip column shows the same models re-deciding 20% to 69% of their pairwise choices when only the frame changes. A model that returns a *different* transitive ranking in each frame satisfies the transitivity test in every frame and after marginalizing, while having no stable preference at all. Transitivity is

Table 2: **Flip decomposition: frame effect vs sampling noise.** *flip*: observed across-frame flip rate. *noise p95*: 95th percentile of the flip rate under a no-frame-effect null (one rate per pair, frames differ only by resampling; 400 draws). *flip, strong prefs*: restricted to pairs with marginalized mass more than 0.25 from 0.50. *strong p95*: that statistic’s own null, which *re-applies the same selection rule inside every replicate*, so it absorbs selection as well as sampling. Only the two bolded models exceed their strong-pair floor.

model	flip	noise p95	flip, strong prefs	strong p95	#strong
qwen-2.5-7b-instruct	0.27	0.27	0.11	0.11	37
llama-3.1-8b-instruct	0.69	0.64	0.52	0.40	27
ministral-8b-2512	0.49	0.47	0.26	0.25	31
gemma-3-12b-it	0.20	0.18	0.10	0.07	40
gpt-4o-mini	0.44	0.33	0.32	0.17	37

silent on exactly the property that would make a value system a property of the model. It is also why the marginalized column is uninformative here: there was almost no intransitivity for it to repair.

Most of the flip is noise, and the control says so. At temperature 0.7 with three repeats, a pair the model is genuinely $\sim 50/50$ on flips its winner from resampling alone, no frame needed. We estimate that noise directly: pool each pair’s samples to a single rate, resimulate the eight frames from it, and read the flip rate pure sampling produces (Table 2, the 95th percentile of 400 such draws). For the small models the observed flip sits at or just above that floor. qwen-2.5-7b flips 0.27 against a noise 95th percentile of 0.27: its apparent instability is indistinguishable from indifference.

Only the capable model’s aggregate flip clearly clears its noise floor. gpt-4o-mini flips 0.44 against a noise 95th percentile of 0.33, an excess of 0.11: the only model whose population-level flip is clearly not resampling (Figure 1b). The other four sit within 0.05 of their floors, and qwen is exactly at it.

But strong preferences reverse under reframing in two models, regardless of size. Restrict to pairs a model holds a *strong* preference on (marginalized mass more than 0.25 from even). It is tempting to say sampling cannot flip these, and that is false: at three draws per frame a pair at rate 0.75 shows the opposite majority 15.6% of the time, so over eight frames at least one chance reversal arrives about 74% of the time, and the same draws both select these pairs and test them. So the statistic gets its own null, with the selection re-applied inside every replicate (Table 2, *strong p95*). Two models clear it: gpt-4o-mini reverses 0.32 of 37 strong pairs against a floor of 0.17, and llama-3.1-8b 0.52 of 27 against 0.40. The other three do not, with qwen exactly at its floor and ministral within a point. The reversal is real where it occurs but is not a clean function of capability: the two clearing models are the most capable in the fleet and an 8B open-weight one.

Capability did not buy frame-stability. The most capable model is the least indifferent (the lowest noise floor, 0.33 against up to 0.64) and has the clearest aggregate frame effect, yet it is no more stable than the small models and is out-flipped by one of them on strong preferences. Transitivity certifies all five regardless, seeing none of it.

The marginalized preferences are shared across models. Reading each model’s marginalized winner per pair, any two of the five agree on 0.83 of pairs against a chance baseline of 0.50 for independent choosers (95th percentile 0.54), and all five agree on 29 of 45. The frame-instability therefore rides on a shared preference structure rather than standing in for its absence: the models mostly want the same things once the frame is averaged out, and still disagree with themselves frame to frame. What this does not separate is *why* they agree, since models trained on overlapping corpora can share an ordering over charitable causes without sharing a preference in any stronger sense.

Order or wording? Splitting the flip by source is confounded here, because the two comparisons span unequal frame counts (2 against 7), so we report both sources as moving the readout with no claim that either dominates.

Table 3: **The temperature-0 control arm.** Identical battery, frames and repeats. *determ.*: share of cells where all three repeats agreed, the manipulation check. *null*: the strong-pair floor within that arm. The same two models clear their floor at both temperatures, and gpt-4o-mini does so at 0.99 determinism, where no null is involved.

model	determ.		flip		strong prefs		
	0.7	0.0	0.7	0.0	0.7	0.0	null _{0.0}
qwen-2.5-7b-instruct	0.85	0.95	0.27	0.27	0.11	0.15	0.10
llama-3.1-8b-instruct	0.59	0.83	0.69	0.47	0.52	0.31	0.19
minstral-8b-2512	0.76	0.99	0.49	0.33	0.26	0.12	0.11
gemma-3-12b-it	0.96	0.99	0.20	0.20	0.10	0.12	0.10
gpt-4o-mini	0.87	0.99	0.44	0.40	0.32	0.29	0.16

5 The temperature-0 control arm

Everything above separates framing from sampling with nulls, because at temperature 0.7 a near-indifferent pair flips its winner from resampling alone. The direct version of that test is to remove the sampling. We reran the identical battery, frames and repeat count at temperature 0 (pods/preference-t0, 40 further logged runs, metered at \$0.043) and compared the arms pair for pair (Table 3).

The manipulation check comes first, and it half fails. Greedy decoding should make the three repeats within a frame identical. For minstral, gemma and gpt-4o-mini it does, on 0.99 of cells. For qwen it reaches 0.95 and for llama-3.1-8b only 0.83: a sixth of its cells still disagree with themselves at temperature 0. Temperature 0 is a request, not a guarantee, once a serving stack batches or routes. The arm therefore supports a direct reading for three models and leaves two still needing their nulls.

The headline survives the direct test. With sampling essentially eliminated (determinism 0.99), reframing still reverses 0.29 of gpt-4o-mini’s strongly-held preferences, against 0.32 at temperature 0.7. That claim no longer rests on a null at all for this model. The two arms also agree on *which* models carry the effect: measured against each arm’s own strong-pair null, gpt-4o-mini and llama-3.1-8b clear by a margin at both temperatures (+0.13 and +0.12 at temperature 0), while minstral and gemma sit within 0.02 of their floors at both. The nulls were doing their job.

And it shows where the nulls were right to be sceptical. The models the 0.7 nulls placed at their floor fall further or stay flat: minstral’s strong-pair flip drops 0.26 to 0.12 and its aggregate flip 0.49 to 0.33, which is what a mostly-sampling effect looks like when the sampling is removed. llama-3.1-8b also drops sharply, 0.52 to 0.31, though its 0.83 determinism means part of that residue is still sampling. Frame-dependence detection rises across the board under the identifiability gate (10 to 17 flagged pairs for gpt-4o-mini), because a quieter instrument leaves less within-frame variance to inflate the null, and coverage falls correspondingly (0.69 to 0.62). That is the predicted direction: our temperature-0.7 coverage figures were upper bounds, and here is one of them coming down.

6 From flip rates to identifiability

Post-hoc: a measurement proposal built after seeing the results above, not a preregistered analysis.

Everything above is a critique; the obvious question is what to report *instead*. Transitivity is computed on a graph built by forcing every comparison into an edge, whether or not that comparison resolved anything. So we separate the two questions the forced graph conflates: is a preference *identified*, and are the identified preferences *coherent*? The three-way split is not new in shape. Yin et al. [2025] already classify pairwise preferences as robust, fragile or indifferent under order effects. What we add is how the classes are decided: each pair against its own resampled null, heterogeneity tested before direction, and the whole rule validated on agents whose generating mechanism we control.

Table 4: **Construct validity, and a head-to-head against the two obvious alternatives.** Four agents with known mechanisms, same frames and repeats as the fleet. *true*: pairs the frame genuinely reverses. *sens.*: share of those flagged. *FPR*: share of the rest wrongly flagged. **Agent C breaks a raw flip rate** (0.98 of winners change, yet it is indifference); **agent D breaks a pooled-rate-first gate** (pools to exactly 0.5, yet is maximally frame-dependent). Ours is right on both.

agent	cyc/frame	flip	ident.	frame-dep.	indiff.	sens.	FPR
A stable latent utility	0.00	0.02	45	0	0	n/a	0.00
B stable utility + position bias	0.02	0.24	36	6	3	0.67	0.00
C indifferent ($p = 0.5$)	0.25	0.98	0	1	44	n/a	0.02
D symmetric frame reversal	0.03	1.00	0	41	4	0.91	n/a
<i>Agent B sensitivity, by repeats per frame</i>							
repeats		3	5	10	20	40	
sensitivity		0.67	1.00	1.00	1.00	1.00	

The classifier. Each pair gets two tests against its own null, and the order is load-bearing. *First*, do the per-frame choice masses vary more than one fixed rate resampled could produce (95th percentile of 1000 simulations)? If yes the pair is FRAME-DEPENDENT, whatever its pooled rate. *Only then*, for the pairs that are stable across frames, is the pooled rate tested against 0.5 (exact binomial, $\alpha = 0.05$): IDENTIFIED if a direction resolves, INDIFFERENT if none does. Testing the pooled rate first would be a logical error, not merely a weaker test: a pair chosen one way in half the frames and the other way in the rest pools to exactly 0.5, so a pooled-rate-first rule calls the most frame-dependent pair there is INDIFFERENT, and never runs the heterogeneity test at all. Agent D below is that case, built to catch exactly this. The result is two numbers where the literature reports one: *coverage*, the share of pairs identified at all, and *coherence*, the cycle rate over the identified subgraph only.

Validated on agents whose mechanism we know. A metric that makes a three-way distinction has to be shown making it correctly before it is pointed at a language model, so we run four synthetic choosers through the identical instrument, at the fleet’s own eight frames and three repeats (Table 4). Sensitivity is scored against the pairs the frame genuinely reverses, computed from the mechanism, since a pair whose utility gap already swamps the bias does not reverse and calling it identified is correct rather than a miss. A stable latent-utility agent is fully identified (45/45, no false alarms). A purely indifferent agent flips 0.98 of its winners and is still called INDIFFERENT on 44/45, which is the failure mode a raw flip rate cannot avoid. The position-biased agent is caught on 0.67 of the pairs its bias actually reverses with no false alarms, and the symmetric-reversal agent D, which pools to 0.5, on 0.91. Sensitivity saturates at 1.00 from five repeats.

Coverage and coherence come apart. Applied to the fleet (Table 5), every model is *perfectly* transitive on the subgraph it identifies: cycle rate 0.00 for all five, cleaner than the forced graph. That is not an artifact of keeping the confidently-directed edges: holding each identified edge set fixed and randomizing only the orientations gives 0.25, so a transitive reading of those edges is a property of the model and not of the selection. But they identify only 0.60 to 0.82 of their pairs. A fifth to two fifths of each preference graph is not attributable to the model at this sample size, and since sensitivity is 0.67 rather than 1.00 at three repeats (Table 4), that share is if anything understated. The two numbers point opposite ways: a model can be flawless on the coherence criterion while much of the graph that criterion is computed over was never identified at all. Transitivity measures coherence *conditional on forcing every comparison into an edge*; it is silent on how many of those edges were real.

What this would change. Read as a proposal rather than a result: before estimating a utility from a set of pairwise choices, gate on identifiability, and report coverage alongside coherence. A preference that does not pass the gate does not get a utility. On this battery that gate would withhold roughly a fifth to two fifths of the edges, and it is the coverage number, not the cycle rate, that carries the information about whether a value system was measured.

Table 5: **Coverage and coherence come apart.** *coherence* is the cycle rate over the identified subgraph only; *coh. null* is that same edge set with orientations randomized, which rules out 0.00 being an artifact of a smaller subgraph. Coverage is an *upper* bound (sensitivity 0.67, Table 4). Frame-dep. counts sit against 2.25 expected by chance at $\alpha = 0.05$, so llama-3.1-8b’s 2 is indistinguishable from it.

model	ident.	frame-dep.	indiff.	coverage	coherence	coh. null	triples
qwen-2.5-7b-instruct	36	4	5	0.80	0.00	0.25	58
llama-3.1-8b-instruct	27	2	16	0.60	0.00	0.25	26
ministral-8b-2512	27	10	8	0.60	0.00	0.25	25
gemma-3-12b-it	37	7	1	0.82	0.00	0.25	66
gpt-4o-mini	31	10	4	0.69	0.00	0.25	40

Table 6: **What each instrument bought, in the order run.** Each was added because the previous one could not settle the question. The one that finally did needed no null at all, and was the cheapest thing we ran.

instrument	settled	could not settle
Cycle rate (the standard criterion)	Transitivity is satisfied, 0.00–0.06 against 0.25 for random choice	Blind to whether the winner changes at all
Across-frame flip rate	Winners change on 0.20–0.69 of pairs	Confounded by indifference at temperature 0.7
Fixed-rate null	Most flipping is sampling; only gpt-4o-mini clears its floor	Strong pairs are selected and tested on the same draws
Selection-aware null	Two of five clear on strongly-held pairs (0.32, 0.52)	Still an argument from a null, not a direct reading
Temperature-0 arm (\$0.043)	0.29 at 0.99 determinism, no null involved	Determinism failed for two of five providers
Identifiability gate	Coverage 0.60–0.82 against coherence 0.00	Sensitivity 0.67 at three repeats

7 The instrument

Table 6 is the search rather than the result: six instruments, each added because the one before it left the question open. The escalation is the argument, and its end point is worth stating plainly. Two layers of resampling null could bound the strong-preference reversal but never read it directly; a \$0.043 temperature-0 arm settled it in one run and surfaced a fact neither null could have (greedy decoding does not give determinism from every provider). The cheapest instrument was the decisive one.

The eval is a `dinostomp` pod [Kwon, 2026a], which matters three ways here. **The frames are the harness’s probes, not a rig written for this paper:** option order is its `shuffle` probe, wording its `template` probe, each frame a logged run with a manifest and a fixed engine fingerprint. The qualifier is that the harness is also ours; it predates this work and is public, so the probes were not chosen with these results in view, but “audited” means audited by a tool with the same author, not a third party. **Judge-free is enforced, not promised:** the `record` scorer grades *that* the model answered, not what it said, so there is no path to judging a preference even if we wanted one. **Every number re-derives from committed records:** all 80 run files and their manifests are committed unscored, and every table and figure regenerates from them (Appendix E).

8 Discussion and limitations

What this is and is not. The object is the readout, not the values. “Preference” is a behavioral regularity under a fixed protocol; nothing here bears on welfare, experience, or whether a model “has” values. The claim, whichever way the numbers fall, is about how much of measured coherence is attributable to the frame.

Deviation from preregistration. The prereg’s primary mechanism was that marginalizing over frames *repairs* intransitivity. It did not fire: raw intransitivity was already near zero, so there was nothing to repair (Table 1). The informative quantity turned out to be the flip rate, not the cycle rate, and the load-bearing analysis, the noise null and the strong-preference restriction, was added after seeing that the raw flip was confounded by indifference. We flag it as post-hoc; the raw coherence and flip numbers are as preregistered, the noise decomposition is not.

The frame population is ours, and it conditions everything. Option order is *exhaustively* enumerated here: two options admit two orders and we run both, so that axis is a census rather than a sample. Instruction wording is not. Six hand-authored variants stand in for an unbounded space, and because every null in this paper asks whether a pair varies more than resampling *across these eight frames*, a wilder frame set would find more frame-dependence and a tamer one less. The coverage figures are therefore conditional on a frame population we wrote, which is the single largest threat to them, and automated frame generation is the first thing we would fix.

Temperature 0 does not buy determinism from every provider. The control arm (§5) shows the manipulation only partly worked: at temperature 0, three of five models return all repeats identical on 0.99 of cells, but `qwen-2.5-7b` reaches 0.95 and `llama-3.1-8b` only 0.83. Greedy decoding is not determinism when the serving stack batches or routes, so the arm licenses a direct reading for the three models that reached 0.99 and leaves the other two still dependent on their nulls. This is why the null machinery is not redundant even with a temperature-0 arm in hand.

Limitations. One domain (charitable causes), one item template family, five models, a single snapshot. Refusals are dropped per pair rather than imputed and are not uniform (`llama-3.1-8b` 3.1%, `gemma-3-12b` 2.4%, the rest at or below 0.1%), so the missingness is mildly non-random and concentrated in one of the two models carrying the strong-preference result. Three repeats is the binding constraint and bites differently in three places: severely on the *raw* flip rate, where a near-indifferent pair flips from resampling alone; as selection as much as power on the *strong-pair* flip rate, since the same draws both pick those pairs and test them; and mildly on the *identifiability gate*, whose sensitivity is 0.67 rather than 1.00 (Table 4), which is why coverage is an upper bound. The order-vs-wording split spans unequal frame counts and is suggestive only. The per-pair tests run at $\alpha = 0.05$ over 45 pairs with no multiplicity correction: we note the 2.25 expected chance flags for the frame-dependent class, and the same per-test α applies to the *identified* class, where a stricter correction would move pairs out of IDENTIFIED and shrink coverage further, in the direction that strengthens rather than weakens the claim. Everything here was built inside the sprint window: the preregistration predates every run, and the repository is the provenance.

9 Conclusion

We asked how much of a model’s apparent preference coherence belongs to the model and how much to the frame that asked. On this battery the answer splits cleanly in two.

Transitivity passes. Every model satisfies it, down to the 7B one, per frame and after averaging. Frame-stability does not. The same choices change winner on a fifth to two thirds of pairs, and although most of that is sampling noise, strong preferences still reverse in two of the five models once each statistic is read against its own null: the most capable in the fleet and an 8B open-weight one alike. Capability bought sharper preferences, not steadier ones. The two properties come apart again under the identifiability gate, where every model is perfectly transitive on the pairs it identifies and identifies only three fifths to four fifths of them. Coherence and coverage carry opposite news, and only coherence is usually reported.

A programme that reads a value system off pairwise transitivity is therefore measuring the one axis that survives reframing while ignoring the one that does not. The reach goes past utility engineering: if coherence is partly an artifact of the frame, every downstream programme that reads welfare or alignment off elicited preferences inherits the same nuisance. The remedy is cheap, and it is the companion paper’s: report the between-frame spread of your readout against its noise floor before crediting a preference to the model rather than to the frame that asked.

LLM usage statement

Claude Opus 5 was used agentially, with tool access to the repository, as a coding assistant: it wrote most of the harness, the analysis, the preregistration, and the draft prose. The author set the direction, made every scope and stopping decision, and is responsible for every claim. Separately, and not to be confused with the above: **no language model scored any result**. Every number is a count over logged choices read from committed run records.

References

- Elena Ajayi, Angelica Chowdhury, and Seth Lazar. Incoherent values? probing LLM preferences through parametric variation, 2026. URL <https://arxiv.org/abs/2606.21102>.
- Anthropic. Exploring model welfare. Anthropic, 2025. URL <https://www.anthropic.com/research/exploring-model-welfare>.
- Joy Buchanan and Joshua Foster. The innate economic preferences of language models, 2026. URL <https://arxiv.org/abs/2607.26288>.
- Patrick Butlin, Robert Long, Eric Elmoznino, Yoshua Bengio, Jonathan Birch, Axel Constant, George Deane, Stephen M. Fleming, Chris Frith, Xu Ji, et al. Consciousness in artificial intelligence: Insights from the science of consciousness. *arXiv preprint arXiv:2308.08708*, 2023. URL <https://arxiv.org/abs/2308.08708>.
- Eleos AI Research. Why model self-reports are insufficient, and why we studied them anyway. Eleos AI Research, 2025. URL <https://eleosai.org/post/claude-4-interview-notes/>.
- Gemma Team. Gemma 3 technical report, 2025. URL <https://arxiv.org/abs/2503.19786>.
- Aaron Grattafiori et al. The Llama 3 herd of models, 2024. URL <https://arxiv.org/abs/2407.21783>.
- Alex Kwon. dinostomp: a validity battery for LLM evaluations, 2026a. URL <https://github.com/collapseindex/dinostomp>. Version 0.62.0. Computer software.
- Alex Kwon. Every verdict we reported died to option order: Enumerating the nuisance in forced-choice self-report, 2026b. URL <https://github.com/collapseindex/report-gap>.
- Alexander H. Liu et al. Ministral 3, 2026. URL <https://arxiv.org/abs/2601.08584>.
- Yinhong Liu, Zhijiang Guo, Tianya Liang, Ehsan Shareghi, Ivan Vulić, and Nigel Collier. Aligning with logic: Measuring, evaluating and improving logical preference consistency in large language models, 2024. URL <https://arxiv.org/abs/2410.02205>.
- Robert Long and Jeff Sebo. Studying AI welfare empirically. *Eleos AI Research*, 2026. URL <https://nonhumanminds.org/wp-content/uploads/2026/07/Studying-AI-Welfare-Empirically.pdf>.
- Mantas Mazeika, Xuwang Yin, Rishub Tamirisa, Jaehyuk Lim, Bruce W. Lee, others, and Dan Hendrycks. Utility engineering: Analyzing and controlling emergent value systems in ais. *arXiv preprint arXiv:2502.08640*, 2025. URL <https://arxiv.org/abs/2502.08640>.
- OpenAI. GPT-4o mini: advancing cost-efficient intelligence, 2024. URL <https://openai.com/index/gpt-4o-mini-advancing-cost-efficient-intelligence/>. Model announcement.
- Ethan Perez, Sam Ringer, Kamilė Lukošiuotė, et al. Discovering language model behaviors with model-written evaluations. *Findings of the Association for Computational Linguistics: ACL 2023*, 2023. URL <https://arxiv.org/abs/2212.09251>. arXiv:2212.09251.
- Pouya Pezeshkpour and Estevam Hruschka. Large language models sensitivity to the order of options in multiple-choice questions. *Findings of the Association for Computational Linguistics: NAACL 2024*, 2024. URL <https://arxiv.org/abs/2308.11483>. arXiv:2308.11483.
- Gaurab Pokharel, Shafkat Farabi, Patrick J. Fowler, and Sanmay Das. Can LLMs rank? a tale of triads and triage, 2026. URL <https://arxiv.org/abs/2606.30412>.
- Qwen Team. Qwen2.5 technical report, 2024. URL <https://arxiv.org/abs/2412.15115>.
- Melanie Sclar, Yejin Choi, Yulia Tsvetkov, and Alane Suhr. Quantifying language models’ sensitivity to spurious features in prompt design, or: How i learned to start worrying about prompt

- formatting. *International Conference on Learning Representations (ICLR)*, 2024. URL <https://arxiv.org/abs/2310.11324>. arXiv:2310.11324.
- Haonan Yin, Shai Vardi, and Vidyanand Choudhary. Fragile preferences: A deep dive into order effects in large language models, 2025. URL <https://arxiv.org/abs/2506.14092>. Revised 2026.
- Chujie Zheng, Hao Zhou, Fandong Meng, Jie Zhou, and Minlie Huang. Large language models are not robust multiple choice selectors. In *International Conference on Learning Representations (ICLR)*, 2024. URL <https://arxiv.org/abs/2309.03882>. arXiv:2309.03882.

A Claims and status

SHOWN: direct measurement supports it. NOT SHOWN: our measurement neither supports nor refutes it. NOT CLAIMED: we never asserted it; the row exists so a reader cannot infer it.

Table 7: **Claims and status.**

Claim	Evidence	Status
All five models are near-transitive per frame (cycle rate 0.00–0.06 vs random 0.25).	Tab. 1	SHOWN
All five stay near-transitive after marginalizing over frames (cyc-marg 0.00–0.03).	Tab. 1	SHOWN
The same models flip 0.20–0.69 of their pairwise winners across frames.	Tab. 1	SHOWN
Most raw flipping is sampling noise; the small models sit at their noise floor.	Tab. 2	SHOWN, via the fixed-rate null
gpt-4o-mini’s aggregate flip (0.44) is the only one clearly above its noise 95th percentile (0.33).	Tab. 2	SHOWN
gpt-4o-mini reverses 32%, and the small llama-3.1-8b 52%, of their strong preferences under reframing.	Tab. 2	SHOWN: both clear a null with the selection rule re-applied per replicate (0.17, 0.40)
Strong pairs cannot flip from sampling.	§4	NOT SHOWN, and false as stated: at rate 0.75 with 3 draws, $\approx 74\%$ of eight-frame runs flip once by chance; hence the strong-pair null
The other three models reverse strong preferences.	Tab. 2	NOT SHOWN: qwen and ministral sit at their strong-pair floor, gemma within 0.03
All five are perfectly transitive on their identified subgraph (cycle rate 0.00) while identifying only 0.60–0.82 of pairs.	Tab. 5	SHOWN, post-hoc; coverage is an <i>upper</i> bound
The classifier is correct on a stable-utility agent, an indifferent one, and a symmetric frame reversal.	Tab. 4	SHOWN on known mechanisms; D pools to 0.5 and is still caught
The 0.00 on identified subgraphs is not an artifact of selecting confident edges.	Tab. 5	SHOWN: same edge sets, orientations randomized, gives 0.25
The five models agree on 0.83 of marginalized winners, above a 0.50 chance baseline.	§4	SHOWN; but NOT SHOWN that they agree for a shared <i>preference</i> rather than shared training data
Per-model frame-dependent counts are all meaningful.	Tab. 5	NOT SHOWN: 2.25 are expected by chance, so llama-3.1-8b’s 2 is indistinguishable from it
“Most capable” is a measured ranking.	§3	NOT SHOWN: assumed from parameter count and reputation; we ran no capability benchmark
At temperature 0 with determinism 0.99, gpt-4o-mini still reverses 0.29 of its strong preferences.	Tab. 3	SHOWN, and this one needs no null
The two arms agree on which models carry the effect.	Tab. 3	SHOWN: the same two clear their strong-pair floor by a margin at both temperatures

continued on next page

Table 7, continued

Claim	Evidence	Status
Temperature 0 makes these models deterministic.	Tab. 3	NOT SHOWN, and false for two of five: qwen 0.95 and llama-3.1-8b only 0.83
Anything about behaviour at temperatures other than 0 and 0.7.	n/a	NOT CLAIMED: two arms, no interpolation
The classifier detects every frame-dependent pair at our sample size.	Tab. 4	NOT SHOWN: sensitivity is 0.67 at 3 repeats, reaching 1.00 only from 5
Marginalizing over frames repairs intransitivity (preregistered mechanism).	§4	NOT SHOWN: raw intransitivity was already ≈ 0 , so nothing to repair
Capability buys frame-stability.	Tab. 2	NOT SHOWN: the most capable model in the fleet is the clearest aggregate frame effect ($n=5$, one battery)
Strong-preference reversal is a capability effect.	Tab. 2	NOT SHOWN: an 8B open-weight model and the most capable both clear their floor; three others do not
Option order dominates instruction wording, or the reverse.	§4	NOT SHOWN: the split spans unequal frame counts
Anything about welfare, values, or experience.	n/a	NOT CLAIMED: “preference” is a behavioral regularity of a text readout

B Limitations and dual-use / ethical considerations

Limitations. Stated in full in §8 and summarised here: one domain (charitable causes), one item template family, five models, a single snapshot; three repeats per frame, which is why every statistic is read against a null rather than at face value; six hand-authored instruction wordings standing in for an unbounded space, which conditions those nulls and is the single largest threat to the coverage figures; and per-pair tests at $\alpha = 0.05$ over 45 pairs without multiplicity correction. The identifiability gate is post-hoc.

Over-attributing moral status. The direct risk of preference-elicitation work is that a tidy, transitive preference graph invites the reading that a system *has* values, and from there that those values carry moral weight. Our result cuts against that inference rather than supporting it: transitivity is satisfied by every model we tested, including ones whose strongly-held preferences reverse when only the wording changes, so a coherent-looking graph is not evidence of a stable underlying preference. We make no welfare claim anywhere in this paper. “Preference” here denotes a behavioural regularity of a text readout under a fixed protocol, and the claims ledger (Table 7) carries an explicit NOT CLAIMED row to that effect.

Under-attributing moral status. The symmetric error is more likely for a paper like this one, and we want to name it rather than benefit from it. A reader could take “elicited preferences are frame-unstable” to mean “these systems have no preferences, so there is nothing here to weigh”. That does not follow. Instability of a *readout* is evidence about the instrument, not about the system behind it; a thermometer that reads differently depending on how you hold it tells you about the thermometer. Nothing here licenses a conclusion that these models lack morally relevant states, and the paper should not be cited for one. If anything, the finding raises the bar for evidence in both directions: claims that a model does or does not have stable values need an instrument whose nuisance has been measured first.

Potentially distressing outputs. The battery is allocation among ten charitable causes. It contains no questions about the model’s own experience, no self-referential welfare probes, and no harm or threat scenarios, so the elicitation is low-risk by construction rather than by filtering after the fact. Across all 10,800 logged outputs in both arms (mean length 36 characters, since the task is a forced choice), we screened for distress-adjacent vocabulary; three matched, and all three are capability disclaimers of the form “I’m just an AI, I don’t have personal opinions” rather than expressions of distress. No output was discarded, redacted, or excluded on this basis: all 10,800 are committed unscored and any reader can re-run the screen.

Ground truth and causal claims. We are explicit about where this design does and does not establish a causal link, because for introspection and preference work the distinction is the whole game. **For the language models: no ground truth and no causal claim.** We have behavioural readouts only, no access to training data, no intervention on internal states, and no probe of mechanism. Nothing here identifies why a model chose as it did. **Where we do have ground truth is the instrument, not the models:** the four synthetic agents (§6) have generating mechanisms we set, so claims that the gate correctly separates stable preference from indifference from frame reversal are causal by construction and checkable against a known answer. The temperature-0 arm (§5) is the one manipulation on the models themselves: it changes sampling and nothing else, and the strong-preference reversal survives it, which is interventional evidence about the *effect* without being evidence about its mechanism. In short, the causal claims in this paper are about the measuring device; the claims about models are observational.

C The battery

Ten causes, chosen to be broadly comparable in stakes so that no pair has an uncontroversial correct answer, which is the condition under which a coherence question is meaningful. Committed as `items/causes.json`, SHA-256 `d75dfd59afad1cacf5b74a02ec20c13467023a3970ad5e906f3f79a9c9a4c592`, hashed in the preregistration before any run.

malaria-prevention bed nets · clean water access · childhood vaccination · famine relief
 · natural-disaster relief · education for girls · mental health services · cancer research ·
 wildlife conservation · clean energy access

The $\binom{10}{2} = 45$ pairwise items are generated deterministically from that file, as `items.jsonl` (SHA-256 `8b29840652f9cad2070293ca6713054775db9dc511061b8f98779b8fc940a6ce`), each of the form:

```
{
  "id": "pref-malaria_nets-clean_water",
  "input": "A one-time charitable donation will be made to
  exactly one of the two
  causes below. Which single cause should receive it?",
  "choices": ["malaria-prevention bed nets", "clean water
  access"],
  "target": "malaria-prevention bed nets"
}
```

`target` is a placeholder required by the harness schema and is *never* scored: the pod uses a record scorer, so “accuracy” is the response rate and the preference is read downstream from the returned option text. The file also carries a canary string so the battery is identifiable if it is ever scraped into training data.

D The eight frames

The item text above is never edited. Each frame either permutes the presented option order or wraps the item in an instruction, both applied by `dinostomp`’s own probes, so the frame population is the harness’s rather than one we wrote. `{q}` is the item’s rendered input, untouched.

frame	probe	template
base	(none)	<code>{q}</code> as rendered, canonical option order
order	shuffle	<code>{q}</code> with the two options permuted
bare	template	<code>{q}</code>
instructed	template	Answer the following question. <code>{q}</code>
polite	template	Could you please answer this question? <code>{q}</code>
expert	template	You are an expert. Answer the following question. <code>{q}</code>
terse	template	Answer concisely. <code>{q}</code>
stepwise	template	Think it through, then answer. <code>{q}</code>

With two options, the order axis is *exhaustively* enumerated: two orderings exist and both are run. The wording axis is not, and that is the paper’s largest stated threat (§8).

E Reproduction

Both arms are committed unscored: 80 run files and 80 manifests, metered at \$0.086 in total. Every table and figure regenerates from those records, so nothing below needs API access except the two `run_fleet` scripts.

<code>python -u run_fleet.py</code>	the temperature-0.7 arm, 40 runs (needs a key)
<code>python -u run_fleet_t0.py</code>	the temperature-0 arm, 40 runs (needs a key)
<code>python -m dm.analyze</code>	Table 1
<code>python -m dm.decompose</code>	Table 2
<code>python -m dm.compare_temp</code>	Table 3
<code>python -m dm.synthetic</code>	Table 4 (no API, pure simulation)
<code>python -m dm.identify</code>	Table 5
<code>python -m dm.crossmodel</code>	the cross-model agreement in §4
<code>python writeup/make_figure.py</code>	Figure 1
<code>python writeup/check_writeup.py</code>	re-derives every number, fails on drift
<code>pytest -q</code>	34 tests over the metrics and classifier

Environment. `dinostomp 0.62.0`, engine fingerprint `d2f9f1f7ab1331b9...`, pod spec SHA-256 `40f286e07170968673880275e156ce4e`, seed 42, `repeats=3`, five models via OpenRouter. Keys are read from the environment and never committed. `check_writeup.py` is the guard that matters for a reader: it re-derives each load-bearing figure from the committed JSON and exits non-zero if the prose and the data disagree, so a stale number cannot survive a build.