

Kavian OS — Weekly Behavioral Intelligence Report

Period: August 27 – September 3, 2026

HOW TO READ THIS REPORT

Your logged entries, self-reported alignment ratings, and self-authored standards are the primary evidence in this report. Kavian identifies recurring patterns across that record and, where the data supports it, contextualizes selected patterns with external behavioral research. Research context is not a diagnosis or a prediction about you. Kavian recommends experiments — you decide what to activate, modify, defer, or reject.

1. DECISION BRIEF

Three entries this week, three different domains, three recoveries — but the same underlying mechanism runs through two of them, and the Cognition Audit adds enough historical corroboration to make it worth naming now.

Act now: Social comparison as a credibility threat → test a pre-committed reframe before sending any congratulatory message or public acknowledgment in a context where comparison is live. Review after 5 qualifying events or 21 days.

Build next: Authority-under-scrutiny reactivity → the pattern shows up in Entry 3 and Cognition Audit 1, but one live-period entry isn't enough to write a complete protocol. Log the next two qualifying events in full before drafting a protocol. Review after 2 additional qualifying entries.

Entry 2 (the call with your dad) is the period's strongest alignment moment and doesn't require a protocol — it's the standard to hold, not a problem to fix.

High-cost signal requiring context — not yet a pattern finding: The open Directive (Directive 1, internal remedy, status open, no selected remedy) has no entry attached to it in this period's data. The system can't evaluate what it's connected to or whether it's been addressed. **What to log next:** Link the Directive to its source entry and record whether the internal remedy was attempted, what happened, and the current status.

2. BEHAVIORAL SIGNAL SUMMARY

Coverage: 3 entries across 3 domains (Social/Peers, Family, Work/Leadership). No field notes logged. 0 calibration challenges completed this period.

Dominant emotion logged: Insecurity (Entry 1), Fear (Entry 2), Anger/Humiliation (Entry 3).

Most frequent signal: Somatic activation preceding the response — sinking stomach (Entry 2), tight jaw and racing thoughts (Entry 3), resentment and confusion (Entry 1).

Most common response pattern: Pause-and-redirect. All three entries show a gap between initial activation and the logged response — the anchor held in each case.

ECL range: 6 minutes (Entry 1) to 22 minutes (Entry 2). Entry 2's longer ECL is consistent with the weight of the trigger, not with dysregulation.

Alignment status breakdown: 0 breaches, 2 recoveries executed, 1 nominal.

Self-reported alignment ratings: Entry 1: 3.7 / Entry 2: 4.4 / Entry 3: 3.9. The 3.7 and 3.9 on recovery entries suggest the operator is rating the activation itself, not just the outcome — which is honest and useful data.

Strongest alignment moment: Entry 2. Your dad disclosed declining health and fear you hadn't heard from him before. The logged response was to stay on the phone longer than planned and ask what he actually needed rather than moving immediately to solutions. Pause/Silence deployed successfully. Self-reported alignment: 4.4. That's the right move in the right moment — no anchor failure, no breach, no carryover logged.

Data sufficiency: Limited — 3 entries across 3 domains. Sufficient for directional observation and monitoring; insufficient on its own for a new high-confidence pattern claim. Cognition Audit entries carry meaningful corroborating weight this period and are treated as such.

3. PRIORITY PATTERN CARDS

Pattern 1 — Act now

Finding: Social comparison to a peer's visible success triggers an automatic credibility threat, which the operator has learned to manage behaviorally but hasn't yet neutralized at the interpretive level.

Evidence:

Entry 1: Friend's career announcement in group chat → immediate self-comparison → interpreted as evidence of personal stagnation → recovery via cognitive reframe and genuine congratulations sent. Self-reported alignment: 3.7.

Cognition Audit 3: Peer public recognition → narrative "my effort goes unnoticed because I'm not politically savvy enough" → typical reaction is effort withdrawal and quiet complaint → new paradigm logged: "I can celebrate others and advocate for myself in the same week."

Source Philosophy: Identity statement includes "I'm all about growth" and "I've been mostly financially unstable most of my life with periods of stability followed by instability" — the logged self-concept carries both an achievement orientation and a background instability narrative that may make peer wins land harder than they otherwise would.

Mechanism: Based on what's logged, the pattern may be driven by a gap between a strong internal achievement identity and an external reference point that appears to contradict it. When a peer's win becomes visible, the data suggests the operator's first interpretive move is to treat it as comparative evidence — not about the peer, but about the operator's own trajectory. This is consistent with what behavioral research describes as social comparison under conditions of identity-relevant uncertainty: when a person's sense of progress is already unclear to them, another person's visible milestone can function as an unwanted data point about their own standing. The recovery in Entry 1 was genuine — the congratulations was sent, the reframe was deployed — but the 3.7 self-reported alignment rating and the Cognition Audit 3 pattern together suggest the interpretive move (his win = my stagnation) is still running automatically before the reframe kicks in. The mechanism isn't the comparison itself; it's the speed at which the comparison gets routed through a credibility lens rather than a neutral one. Possible relational exposure: a pattern of this kind, if recurring, tends to create subtle withdrawal or flatness in peer interactions at moments when warmth is most expected — the congratulations gets sent, but something in the tone or timing may register as slightly off to the people around the operator. The current record cannot determine how any specific person has actually experienced it.

Research context — not a personal prediction: Research on social comparison and identity-relevant uncertainty suggests that comparison to a peer in a domain where one's own standing feels unclear tends to produce a threat response rather than an informational one — the peer's outcome gets processed as evidence about the self, not just about the peer. Your record makes a pre-committed interpretive reframe worth testing before the comparison reflex runs; it doesn't establish that this mechanism fully accounts for what happened in Entry 1.

Confidence: Moderate — one live entry, strong Cognition Audit corroboration, consistent with Source Philosophy background.

Data quality: Cross-period corroborated — Entry 1 plus Cognition Audit 3 plus Source Philosophy identity statement.

Cost: Relational cost is low when recovery executes cleanly, as it did in Entry 1. The residual cost is internal — the 3.7 self-reported alignment rating on a clean behavioral recovery suggests the interpretive layer is still generating friction even when the behavior holds. If the pattern runs in contexts where the recovery doesn't execute (no anchor deployed, no reframe available), the logged typical reaction from Cognition Audit 3 is effort withdrawal and quiet complaint — which carries a longer-tail relational and professional cost.

Recommended experiment: Install a single pre-committed question to run before responding in any context where a peer's visible win is the trigger: "Does this tell me anything about my actual trajectory, or am I borrowing their data point?" Log the answer and the response. Test whether the interpretive move changes before the behavioral one has to.

Protocol: See Active Protocol — "Comparison Check" in Section 4.

Review point: After 5 qualifying events or 21 days, whichever comes first.

Emerging pattern — Build next

Finding: Authority-under-scrutiny triggers a fast-moving reactivity sequence (humiliation → defensiveness → over-explanation or voice elevation) that the operator is currently managing with anchors but hasn't yet converted into a governing rule that holds without one.

This period's evidence: Entry 3 — vendor talks over the operator twice in a negotiation meeting, in front of his team. Signals: tight jaw, racing thoughts, defensiveness. Emotions: anger, humiliation. Interpretation: "He's not taking me seriously because of my age."

Response: waited for a natural pause, finished the point calmly. Self-reported alignment: 3.9. Recovery executed.

Historical corroboration: Cognition Audit 1 — being corrected in front of others at work → narrative "they're trying to make me look incompetent" → typical reaction: defensive over-explanation to save face → new paradigm logged: "A leader who can be corrected in the open is more trusted, not less." Exploitability Map Entry 6 — approval from authority figures, tight chest, rehearsing justifications before feedback is even heard. Exploitability Map Entry 10 — need to win the argument, jaw tightens, listening for the next thing to say rather than what's being said.

Mechanism: Based on what's logged, the pattern may be driven by a credibility-protection reflex that activates specifically when the operator's authority or competence is made visible as uncertain — in front of others, in a professional context, with stakes attached. The Cognition Audit and Exploitability Map together suggest this reflex has a fast onset (somatic cues arrive early — tight jaw, racing thoughts — but the logged interpretation arrives at the reaction stage rather than before it) and a consistent narrative structure: the external actor's behavior is read as a verdict on the operator's standing, not as a behavior to be managed. The recovery in Entry 3 was clean. The question the data raises is whether the anchor is doing the work that a governing rule should eventually do on its own — and whether one more qualifying entry would show the same sequence or a different one.

Confidence: Exploratory — one live-period entry, strong audit corroboration, insufficient to call a full pattern without a second live qualifying event.

Data quality: Single live-period entry; cross-period corroborated via Cognition Audit 1 and Exploitability Map entries 6 and 10.

Cost: Professional and relational. In Entry 3, the cost was contained — recovery executed, no breach logged. The Cognition Audit's typical reaction (defensive over-explanation) and the Exploitability Map's jaw-tightening / argument-winning reflex suggest the cost escalates when the anchor doesn't deploy. In a leadership context, a visible reactivity sequence in front of a team carries a credibility cost that the recovery itself doesn't fully erase.

Evidence-collection protocol: For the next two qualifying events where authority is publicly questioned or scrutinized — log: the exact trigger, the first interpretation that ran (before any reframe), the somatic signal and when you noticed it, the behavior that followed, and whether an anchor was deployed or the response ran without one.

Decision threshold: Two additional qualifying entries showing the same interpretive move (external scrutiny → credibility threat → defensive sequence) before a full protocol is drafted. If either entry shows a different interpretive move or a clean response without an anchor, note that — it changes the protocol design.

These observations are based solely on user-submitted data and do not constitute a psychological evaluation, clinical assessment, or diagnosis of any kind. This report is not a substitute for professional mental health, medical, or therapeutic care. It should not be used to make clinical, legal, or diagnostic decisions. The system makes no representations about the accuracy or completeness of this analysis. The user assumes full responsibility for how they interpret and act on the content of this report.

4. ACTIVE PROTOCOLS

Comparison Check

Status: Proposed — activate, modify, or defer from the dashboard.

Connected to: Pattern 1 — Social comparison as credibility threat (Act now).

Owner: You, at the moment a peer's visible win enters your awareness — in a group chat, a meeting, a social feed, or a direct conversation.

Trigger: Any moment you notice a peer's public success and feel the first pull toward self-comparison. **Operator check:** "Am I reading this as information about them, or as a verdict about me?" Yes to the second half = protocol is live.

Protocol: Before responding (typing, speaking, reacting), run one question out loud or in writing: "Does this tell me anything about my actual trajectory, or am I borrowing their data point?" Answer it — even a one-word answer counts. Then respond. If the question doesn't interrupt the comparison reflex (you notice you answered it and still felt the pull), log that as a protocol miss and note what the comparison was actually about — that's the data the next report needs. **Fallback:** if no time to run the question before responding, run it immediately after and log the sequence.

Duration: 21 days or 5 qualifying events, whichever comes first.

Metric: Protocol completion on 80%+ of qualifying events; zero qualifying events in this window that go unlogged. A qualifying event where the question was skipped entirely is a miss — log it as such within 24 hours.

Evidence to log: The trigger (whose win, what context), the answer to the question, the response you gave, and your self-reported alignment rating on the interaction.

If missed: If this protocol is activated and a qualifying event isn't met, log the miss within 24 hours and name what ran instead of the question.

Review: After 5 qualifying events or 21 days — whichever comes first. If fewer than 3 qualifying events occur in 21 days, extend by 14 days before reassessing.

The self-authored calibration widget on the dashboard can be used to build a personal action step for any of these protocols.

These observations are based solely on user-submitted data and do not constitute a psychological evaluation, clinical assessment, or diagnosis of any kind. This report is not a substitute for professional mental health, medical, or therapeutic care. It should not be used to make clinical, legal, or diagnostic decisions. The system makes no representations about the accuracy or completeness of this analysis. The user assumes full responsibility for how they interpret and act on the content of this report.

5. STANDARDS AND SYSTEM GAPS

This section compares your logged entries against your baseline audit documents and flags where either is incomplete.

Gap 1 — Exploitability Map: Missing Neutralization Protocols

What's stated: Entries 2–5 and 10 each name an exploitable need, activation context, and somatic warning signs — but the Neutralization Protocol field is blank or absent on all five.

What the entries show: Entry 3 this period activated what appears to be the Entry 10 lever (need to win the argument, jaw tightens) and the Entry 6 lever (approval from authority figures). The operator managed both with an anchor — but the Exploitability Map's own protocol for Entry 10 is empty, meaning there's no pre-committed neutralization step to fall back on if the anchor doesn't hold.

The gap: The map identifies the vulnerabilities but leaves the response undefined for half its entries — which may mean the anchor is currently doing work the neutralization protocol should eventually own.

Fix: Add a neutralization protocol to Exploitability Map entries 2, 3, 4, 5, and 10. For Entry 10 specifically, the Cognition Audit 1 alt response ("thank the person and ask a clarifying question instead of justifying") is a candidate — but the wording is the operator's call, not this report's.

Gap 2 — Leadership Compass: "People over profits" — Incomplete Breach and Remedy Fields

What's stated: Behavior Proof is present. Breach Proof, Relational Remedies, and Internal Remedies are all blank.

What the entries show: No direct breach of this principle this period. The gap is structural, not behavioral.

The gap: A principle with no breach proof and no remedies defined can't function as a governing standard under pressure — there's no stated condition that would tell the operator they've crossed it, and no pre-committed response if they do.

Fix: Add a concrete breach proof scenario and at least one internal remedy. This context does not appear in the current audit documents. Adding it would give the system more to work with in future reports.

Gap 3 — Directive 1: Open, No Selected Remedy, No Source Entry

What's stated: Directive 1 is open, remedy type is internal, selected remedy is blank.

What the entries show: No entry this period links to this Directive. The system cannot determine what triggered it, whether the internal remedy was attempted, or what the current status is.

The gap: An open Directive with no selected remedy and no linked entry is unresolvable from the outside — the system can flag it but can't evaluate it.

Fix: Link the Directive to its source entry, select or write the internal remedy, and update the status. Complete audit documents give the system more data to surface patterns against.

Gap 4 — Cognition Audit: Strong Processing, No Behavioral Proof of Paradigm Adoption

What's stated: Five Cognition Audit entries this period, all with new paradigms logged and disposition tags assigned (Replace or Discard on four of five).

What the entries show: The audit work is substantive. But none of the five entries has a corresponding logged incident showing the new paradigm operating under live pressure — the paradigm is stated, not yet demonstrated in the behavioral record.

The gap: This isn't a failure — it may simply be that the qualifying triggers haven't occurred yet, or that they occurred and weren't logged. But based on what's available, the system can't confirm whether the new paradigms are holding under pressure or remaining as stated intentions.

Fix: When a Cognition Audit trigger fires in a live situation, log it as an Incident Report and note which audit paradigm was active. That's the data that closes this gap over time.

These observations are based solely on user-submitted data and do not constitute a psychological evaluation, clinical assessment, or diagnosis of any kind. This report is not a substitute for professional mental health, medical, or therapeutic care. It should not be used to make clinical, legal, or diagnostic decisions. The system makes no representations about the accuracy or completeness of this analysis. The user assumes full responsibility for how they interpret and act on the content of this report.

6. EVIDENCE APPENDIX

Pattern 1 — Social comparison as credibility threat

Entry 1: Friend's career announcement in group chat. Trigger: self-comparison to "stalled" trajectory. Signals: resentment, confusion. Emotions: insecurity, frustration. Interpretation: "His win says something about my lack of progress." Response: genuine congratulations sent. Anchor: cognitive reframe deployed successfully. Self-reported alignment: 3.7. Recovery executed. The 3.7 on a clean behavioral recovery is the signal — the behavior held, the interpretive layer didn't fully resolve.

Cognition Audit 3: Peer public recognition → "my effort goes unnoticed because I'm not politically savvy enough" → typical reaction: effort withdrawal and quiet complaint → alt response: congratulate genuinely, then advocate directly with manager → new paradigm: "I can celebrate others and advocate for myself in the same week." Disposition: Replace. This is the same interpretive sequence as Entry 1, logged in a different context — which is what makes Entry 1 more than a single-entry observation.

Source Philosophy cross-reference: Identity statement includes both a strong growth orientation and a background of financial instability punctuated by reaches for new opportunity. The data suggests — though can't confirm — that this background may make peer milestones land as comparative data points rather than neutral information, particularly when the operator's own trajectory feels unclear.

Research context — not a personal prediction: Research on social comparison and identity-relevant uncertainty suggests that when a person's sense of their own progress is ambiguous, a peer's visible success tends to get processed as a self-relevant threat rather than an informational update about the peer. Your record makes a pre-committed interpretive check worth testing before the comparison reflex runs; it doesn't establish that this mechanism fully accounts for what happened in Entry 1.

Emerging pattern — Authority-under-scrutiny reactivity

Entry 3: Vendor talks over the operator twice in a negotiation meeting, in front of his team. Signals: tight jaw, racing thoughts, defensiveness. Emotions: anger, humiliation. Interpretation: "He's not taking me seriously because of my age." Response: waited for a natural pause, finished the point calmly. Anchor: Taking One Full Step Back deployed successfully. Self-reported alignment: 3.9. Recovery executed. The age-attribution in the interpretation is worth noting — it's a specific framing the system can't verify but that, if it recurs, would suggest the credibility-threat narrative has a particular shape.

Cognition Audit 1: Being corrected in front of others at work → "they're trying to make me look incompetent" → defensive over-explanation to save face → alt response: thank and ask a clarifying question → new paradigm: "A leader who can be corrected in the open is more trusted, not less." Disposition: Replace. Same interpretive structure as Entry 3 — external scrutiny read as a verdict on standing, not as a behavior to manage.

Exploitability Map Entry 6: Need for approval from authority figures. Activation context: performance reviews or any moment work is being evaluated by someone above. Somatic warning: tight chest, rehearsing justifications before feedback is heard. Neutralization: "let me sit with that and follow up tomorrow."

Exploitability Map Entry 10: Need to win the argument. Activation context: disagreements where the operator senses he might be wrong but hasn't said so. Somatic warning: jaw tightens, listening for the next thing to say rather than what's being said. Neutralization: blank.

Note on single-entry status: Entry 3 is the only live-period entry for this pattern. The Cognition Audit and Exploitability Map entries provide meaningful corroboration, but the pattern is labeled Exploratory and Build next for this reason — one more qualifying live entry is what moves it to a full protocol.

Entry 2 — Alignment moment, no pattern card

Entry 2: Father discloses declining health and fear on a phone call. Trigger: hearing fear in his voice for the first time. Signals: sinking stomach, restless legs. Emotions: fear, regret. Interpretation: "I should have known and pushed harder to check in sooner." Response: stayed on the phone longer than planned, asked what he actually needed rather than jumping to solutions. Anchor: Pause/Silence deployed successfully. Self-reported alignment: 4.4. Integrity status: Nominal. ECL: 22 minutes — consistent with the weight of the trigger, not with dysregulation. This entry doesn't generate a protocol because nothing broke. It's the clearest example in this period's data of the Compass principle "Regulate Before I Respond" operating under genuine pressure — a high-stakes, emotionally loaded trigger met with a pause and a question rather than a fix. That's the behavior to recognize and repeat.