

Relief Is Not Repair

Failed Recapture and the Causal Structure of Suffering

Shayan Arman

23 August 2026

RESEARCH DRAFT FOR AUTHOR REVIEW - NOT PEER
REVIEWED

409 posts audited	9 source essays	8 core countercases
----------------------	--------------------	------------------------

A philosophical and causal framework for distinguishing phenomenal relief from the repair of injury, deprivation, and lost capability.

Abstract

What is suffering: an aversive experience, a disruption of conscious life, or a significant disruption of agency? I argue that this dispute often conflates an occurrent phenomenal target with its characteristic causes and consequences. **Failed Recapture** is a candidate mechanism for one central family of episodes: aversive corrective pressure control-resistently captures attention and action while neither source resolution nor practical integration restores a flexible field of response. Actual agency disruption is a distinct worldly variable, often itself harmful, that can cause and be caused by such episodes. I represent the distinction with two ledgers that causally interact but can dissociate. This yields the Relief-Repair Gap: the complete interventional distribution of a pre-specified suffering outcome does not, by itself, identify an intervention's effect on an unmeasured, pre-specified worldly target; persistent suffering likewise does not identify failed repair when phenomenal capture can outlast its source. Coupled with the premise that both targets matter, these results support Conditional Ledger Integrity: improvement in one ledger does not by itself establish improvement in the other. The contribution is not a new essence of suffering or new causal mathematics. It is an integrated proposal: target indexing, source/grip control, paired non-identifiability results, and symmetric anti-laundersing rules.

Keywords: suffering; harm; agency; negative affect; adaptation; causal identification; healthcare ethics

1. The question hides two targets

Consider two interventions. The first leaves a disease untouched but abolishes a patient's agony. The second removes the disease, yet the patient's terror, grief, and pain continue for a time. It is natural to call the first *relief without repair* and the second *repair before relief*. Neither description is paradoxical. The distinction is familiar in practice: analgesia may reduce pain without curing its source; surgery may repair tissue before postoperative pain ends; restitution may remove a material deprivation while fear persists; equanimity may reduce distress while an injustice remains.

Yet a prominent metaphysical debate encourages us to compress these outcomes into one variable. Experientialists locate suffering in negative experience. Disruption theorists locate it in an impairment of mental life or agency. Objective theorists extend suffering to conditions in which a person's flourishing, dignity, or opportunities are gravely compromised even when no distress is felt. If each view offers an analysis of the same univocal target, at most one can be right. If, however, the views are tracking different targets under the same word, the demand for a single winner is already misframed.

This paper argues for the latter diagnosis, but not for the deflationary conclusion that the dispute is merely verbal. The targets are substantively different and support different causal, evidential, and moral inferences. An occurrence can damage a subject without being experienced. An experience can be awful without corresponding independent worldly damage or capability loss. Reducing one can leave the other untouched or make it worse. Once these possibilities are represented explicitly, a constraint follows: success measured against one target does not by itself identify or establish success against the other unless a bridge between them is independently justified.

I call the resulting architecture **two ledgers**:

1. The **phenomenal ledger** records a subject's occurrent negatively valenced experience and the degree to which it grips the subject's concerns, attention, motivation, and experienced field of possibilities.
2. The **world/capability ledger** records actual disruption of biological, psychological, and social agency, together with harms and wrongs that remain bad independently of the subject's present experience.

The word *ledger* is intended to constrain inference, not to suggest that life arrives in neat columns. Entries in the two ledgers are often causally entangled. Disease can cause agony; agony can impair sleep and work; deprivation can cause anxiety; chronic anxiety can constrict opportunity. The point is analytical: causal correlation is not identity, and movement in one register cannot simply be booked as movement in the other.

The argument has four parts. First, I reconstruct the experientialist-agency dispute as a conflict over target selection. Second, I offer **Failed Recapture** as a candidate mechanism for a central family of suffering episodes, not as a biconditional analysis of every case. Third, I extend the literature's dissociation cases by tracking phenomenal suffering, actual agency disruption,

and independent harm separately. Fourth, I prove elementary but consequential non-identifiability results: relief does not identify repair, and persistent suffering does not identify failed repair. Together with an explicit normative premise, these results support **Conditional Ledger Integrity**, a no-laundering norm for philosophical, clinical, and political evaluation.

The claim of originality should be stated at the outset. The literature already distinguishes subjective and objective dimensions of suffering; it already contains experiential, affective-construal, mental-disruption, and agency-disruption views; and recent bioethics explicitly asks whether purportedly non-experiential suffering is better described as harm or wrong. The new claim is not that no one has noticed two dimensions. It is that their separation should be treated as an inferential architecture: a disciplined double-entry account with a formal limit on what relief and repair evidence can show. The paper offers a proposed reconciliation of a live dispute. It does not claim priority for the underlying distinctions, settle every ordinary use of *suffering*, provide a complete theory of harm, or discover a phenomenology without antecedents.

2. Four theories and one unstable explanandum

Michael Brady's experientialism begins from negative affect. On his account, suffering is an unpleasant experience that the subject *minds*: roughly, the subject has the negative state and occurrently desires that the unpleasant experience not occur (Brady 2018, 2023). The analysis is meant to distinguish suffering from negative sensations a subject accepts or does not mind. In later work, Brady expressly defends the experiential category against Jennifer Corns's agency account and argues that actual agency disruption is neither necessary nor sufficient for suffering (Brady 2023, 2026).

Antti Kauppinen offers a richer affective account. Attitudinal suffering involves a sufficiently pervasive negative affective construal of one's situation. Something is affectively presented as calling for change; motivation registers that presentation; and the world can come to seem narrowed, darkened, or organized around what hurts (Kauppinen 2019). This view explains why equal sensory intensity need not produce equal suffering: an aversive state may remain local, or it may transform the field of salient possibilities.

Tom McClelland similarly resists identifying suffering with an isolated unpleasant item. On his disruption model, a suffered state brings a sufficiently severe, holistically unpleasant disturbance to conscious mental life (McClelland 2019). The disturbance includes the pressure to accommodate, resist, submit to, or otherwise reorganize around what is happening. McClelland's target remains conscious and experiential, but disruption rather than raw unpleasantness performs much of the explanatory work.

Jennifer Corns rejects the assumption that experience is required. She argues that a subject suffers if and only if the subject's agency is significantly disrupted (Corns 2022). Agency is not restricted to reflective adult planning. Biological, psychological, and social forms of agency can be impaired, and a sufficiently minimal notion is intended to cover infants, nonhuman animals, and subjects whose conscious status is uncertain. This is not merely Brady's view with an extra symptom. A deeply sedated patient may undergo further loss of agency while no negative experience occurs. Corns's account can count that loss as suffering; a strict experientialist cannot.

Adjacent taxonomies show why the disagreement resists a simple vote. Brent Kious distinguishes feeling-based, value-based, and objective conceptions of suffering and argues that they can bear different moral significance (Kious 2022). Tyler Tate has developed both a clinical discussion of what we mean by pediatric suffering and a recent account of objective suffering that seeks to preserve public, flourishing-related criteria (Tate 2020, 2026). Jennifer Blumenthal-Barby examines disorders of consciousness and argues that many concerns described as non-experiential suffering are better classified as harms or wrongs (Blumenthal-Barby 2026). These views make the lexical field explicit. They also create the central challenge for this paper: merely renaming familiar subjective and objective dimensions as two ledgers would add little.

The advance comes from specifying what each theory is trying to explain. At least four questions travel under "What is suffering?"

- **Occurrent-state question:** What must be happening in a subject's experience for the subject to be suffering now?
- **Condition question:** What adverse condition is the subject undergoing or living under, possibly over an extended period?
- **Capability question:** Which forms of biological, psychological, or social agency are actually impaired?
- **Normative question:** Which harms, wrongs, or deprivations give us reason to aid, compensate, protect, or prioritize the subject?

An answer to one is not automatically an answer to the others. “She suffered from cancer for ten years” can summarize an adverse condition, recurring episodes of felt distress, and large capability losses without implying that every moment-including dreamless sleep-contained a negative experience. “The community suffered under discriminatory rule” can correctly predicate a historical condition and injustice even of people who did not understand every mechanism that constrained them. Conversely, a brief episode of overwhelming phantom agony can answer the occurrent-state question even if tissue damage and lasting loss of capability are absent.

The disagreement therefore contains a **target-selection error** when a criterion suited to one question is treated as a complete measure of another. Experientialists err if they infer that an unexperienced constraint is harmless. Objective and agency theorists err if they infer that every serious constraint is an occurrent aversive state. Hybrid theories err if they force both variables into a conjunction and thereby exclude clear cases on each side. The appropriate repair is not a metaphysical average. It is target indexing.

3. The corpus genealogy: from symptom suppression to two ledgers

The two-ledger view emerges from a sequence of posts rather than a single finished theory. Read in order, the posts contain both the core distinction and the overstatements that a defensible paper must reject.

Post [#96, “On Sickness”](#) begins with a distinction between suppressing a symptom and addressing its cause. The post's medical examples are speculative and should not be taken as clinical advice, but its conceptual point is strong: the disappearance of a symptom does not show that the underlying process has been repaired. Removal is only one kind of intervention; addition, counteraction, or watchful nonintervention may sometimes target the cause more effectively. This is the earliest form of the Relief-Repair distinction.

Post [#145, “On Suffering”](#) asks whether painful sacrifice undertaken for a meaningful end is suffering or an ingredient of something glorious. The question notices an important difference between local aversion and global endorsement. Its suggested glorification of pain goes too far. Meaning can make pain bearable, valuable, or willingly endured; it does not follow that an awful episode was not suffering or that suffering is intrinsically good.

Post [#169, “To Take a Stand”](#) holds gross bodily strain roughly fixed while changing its intentional description. Standing for hours to meet a beloved can be joyful; standing under a sniper's threat can be agony. The case is not a controlled experiment, but it isolates a constitutive-looking variable: what an event means to a subject can alter its phenomenal character without immediately altering the physical task.

Post [#206, “Theory of Expressionism”](#) proposes a loop in which escape behavior remains organized by the feeling it avoids and thereby reproduces it. The universal formulation is implausible: not every action expresses one feeling, and avoidance does not always amplify distress. Still, the loop provides a second route to the Relief-Repair distinction. An action can deliver short-term relief while reinforcing a cause, disposition, or pattern that produces later distress.

Post [#263, “Pain”](#) says that meditation does not create the pain it reveals; attention makes a previously avoided condition salient. This blocks a common conflation. Attention to a condition and the condition attended to are distinct variables. Direct attention may sometimes alter the experience, but it is not the universal or sole path to a pain-free life.

The conceptual center appears in post [#283, “On Suffering”](#). It distinguishes an *event* that can persist without attention from *hardship* as minded and adds that “a lack is a lack.” Here the two targets are almost explicit: occurrence on one side, minded hardship on the other. The same post says that attention or rumination can be switched off instantly. That is sometimes

possible, not generally so. Panic, grief, depression, trauma, and pain can capture attention or alter salience involuntarily. Treating phenomenal grip as freely dismissible would turn a possible intervention into a universal power and invite blame of sufferers who cannot exercise it.

Post #288, “[Et tu, Brute?](#)” compresses the claim into the idea that quieting the mind can quiet suffering. The defensible residue is conditional: if an intervention genuinely removes the relevant negative experience, it relieves occurrent suffering. It does not follow that the adverse occurrence has ended, that capability has returned, or that the subject has done nothing wrong by failing to become quiet.

Post #289, “[Moments](#)” adds temporal scope. A terrible moment is not identical to a terrible life; a summary over a period cannot simply be read back into every instant. Finally, post #404, “[On Depression](#)” struggles with depression as bodily located yet intentional or conceptual in character. Its dualist metaphysics is not required here. The safer conclusion is that phenomenal, bodily, behavioral, and agential descriptions can all be true without naming one property.

The genealogy yields a corrected thesis. Events, damage, deprivation, and capability loss can persist while occurrent suffering fluctuates. Meaning and attention can modulate suffering without magically altering the world. Conversely, worldly repair can precede experiential recovery. What began as a contrast between symptom and cause becomes a constraint on metaphysics and evidence.

4. Two ledgers, with columns

The phrase *two ledgers* is useful only if it does not conceal the control process within the phenomenal side or conflate capability with harm on the worldly side. Let the following denote features of a subject's situation over an interval Δ :

Symbol	Variable	Guiding question
H	residual independent harm	Is there bodily damage, deprivation, coercion, injustice, or another bad condition not already represented by S or J?
A	negative affect	Is there conscious sensory or affective unpleasantness?
D	corrective demand	Does the state affectively direct the subject to stop, escape, restore, undo, endure, or find another viable path?
C	priority capture	To what degree does D narrow or preempt otherwise live attention and action policies?
R	recapture	Has flexible, care-sensitive policy selection been restored by satisfying, integrating, or appropriately down-ranking D?
J	significant agentive disruption	Is the reliable exercise of a specified biological, psychological, or social agentive form degraded across relevant circumstances, relative to an appropriate baseline?
S	occurrent suffering	Is the subject suffering during Δ ?

The two registers are:

$$L_P=(A,D,C,R), L_W=(H,J).$$

The adjectives *residual* and *independent* attached to H matter. Suffering may itself be a harm, and agentive disruption may itself be harmful. For this representation, H records damage, deprivation, or wrong not already counted in S or J. Actual disruption J is a separate adverse column: higher J means a greater degradation in the effective exercise of the specified agentive form. This bookkeeping convention prevents double counting; it does not deny interactions or alternative normative decompositions.

Ledger separation asserts neither dualism nor causal insulation. Every variable may be physically realized, and arrows can run in both directions. Disease may increase J, generate A,D,C, and prolonged capture may in turn impair sleep,

relationships, and agency. A false catastrophic appraisal may generate D,C without external H. Social support, medication, or meditation may alter A,C,R without repairing H; surgery may repair H while learned capture persists. The claim is explanatory: the variables occupy different causal roles, so no value in one ledger can be read off from the other by definition.

This also prevents equivocation over *disruption*. McClelland's disruption of conscious mental life maps most naturally onto C and R. For comparative purposes, this paper proposes to reconstruct Corns's agency-theoretic concern as J; that is not a neutral description of her intended analysis of suffering. Her agentive forms concern exercise and environmental modulation, not merely durable ability loss. Accordingly, J records a context-relative degradation in reliable exercise or effective access across a specified domain and interval, against a stated baseline. It is neither every blocked opportunity nor every transient competition within consciousness. If agency disruption is expanded to include every aversive change in the conscious field, the agency view becomes experiential; if phenomenal disruption is expanded to include every blocked capability, it ceases to be phenomenal.

5. Failed Recapture as a process model

The phenomenal-control ledger supports a candidate mechanism deliberately weaker than a complete analysis of suffering:

CANDIDATE FAILED RECAPTURE MECHANISM

Over an interval Δ , failed recapture occurs when (i) a conscious state has negative valence and carries a corrective demand; (ii) that demand exerts experienced policy pressure by involuntarily dominating salience or felt response possibilities; and (iii) the pressure remains unsettled because neither alteration of the represented source nor regulation of its phenomenal grip has ended that domination. The hypothesis is that this process raises the probability or severity of occurrent suffering S in a specified family of episodes. Capture may qualify through momentary intensity, temporal accumulation, or recurrence.

Failed Recapture is not a necessary-and-sufficient analysis of every state ordinarily called suffering. A context q may classify the candidate process as severe when accumulated or peak aversive capture crosses a stated threshold:

$$B(\Delta) = \int a(t)d(t)c(t)u(t) dt; B(\Delta) \geq \theta(q) \text{ or } \text{peak}(\Delta) \geq \phi(q).$$

The integral is taken over Δ . Here $a(t)$ is aversive force, $d(t)$ is corrective-demand force, $c(t)$ is experienced capture, and $u(t)=1-r(t)$ is non-settlement, with $r(t)$ the degree of achieved recapture. The peak is the maximum instantaneous product. The two clauses allow mild recurrence and brief agony to qualify while ordinary hunger or frustration may remain below a reporting threshold. Thresholds must be fixed for the measurement purpose rather than selected after the outcome. This is an explanatory schema, not a validated clinical scale. S remains independently assessed, and the causal results below do not depend on Failed Recapture being correct.

The model distinguishes two kinds of control. **Source control** is the ability to alter the pain, loss, danger, deprivation, or other represented source. **Grip control** is the achieved ability to alter how the corrective demand allocates attention and felt response possibilities without falsely representing the source as gone. An available source-changing option is not enough to prevent present capture: agony can dominate before a stop button is pressed. Source change is not necessary for relief: acceptance, reappraisal, anesthesia, absorption, or social support can sometimes restore a flexible practical field while the source remains. *Recapture* means an achieved return of sufficiently broad, care-sensitive policy selection; it is not bare distractibility, outward performance, or the mere existence of an escape route.

Global endorsement is not settlement. A runner may endorse finishing while bodily aversion still presses for cessation; a mourner may affirm the fittingness of grief while its world-directed pressure - *let this loss not be so* - remains unsettled. Meaning ends failed recapture only if it actually ends or sufficiently reduces involuntary domination. It may instead make an ordeal worth bearing while the process continues.

The model does not imply that every suffering state contains a wish to erase experience or world. Corrective pressure may be primitive and world-directed, not a higher-order desire. Yet resigned depression, apathetic misery, accepted agony, or grief wholly lacking directive force may be suffering without failed recapture. Ordinary problem solving may involve effort without negative valence or control-resistant domination. These edge cases are why Failed Recapture is a candidate mechanism rather than an exhaustive biconditional.

The model remains close to existing work. Brady supplies the experiential core; Kauppinen supplies affective presentation, attentional pervasiveness, and control; McClelland supplies disruption of the wider mental economy. Its proposed refinement is dynamic and relational: source control is separated from grip control, achieved settlement is separated from option-availability and global endorsement, and actual agency disruption is modeled as a possible cause or consequence rather than a constitutive ingredient. The paper's stronger contribution arises when this process model is joined to the Relief-Repair Gap.

6. The double dissociation

Brady (2023, 29-30) already argues against Corns by presenting both directions of the double dissociation, including resilient lovesickness without significant agency disruption and blissful love with significant disruption but no suffering. The present paper adopts rather than originates those counterexamples. Its narrower contribution is to distinguish experienced policy capture from significant agentive disruption, add residual independent harm as a further variable, and trace the resulting separation into causal and evaluative inference.

If actual agency disruption constituted occurrent suffering, then J would be necessary and sufficient for S. Brady's double dissociation challenges both directions.

6.1 Suffering without actual agency disruption

Begin with Brady's **resilient lovesickness** case (Brady 2023, 30). The subject must be specified counterfactually rather than declared intact by construction. During episodes, heartbreak dominates experienced salience; yet when relevant reasons arise, she remains able to redirect attention, deliberate, and execute actions with no significant degradation of reliability across the relevant biological, psychological, or social capacities. The case distinguishes the phenomenal organization of exercising agency from a significant loss of agentive capacity. Outward performance alone would not establish the latter's absence, so the relevant capacities and comparison interval must be independently assessed.

Virtuous grief and remorse strengthen the case. Their painful character may express what the agent cares about and can even support agency by motivating repair, fidelity, or attention to another. Fittingness and agency promotion do not reverse valence. A subject can suffer through an emotion that functions well.

Phantom or nocebo agony is a suggestive auxiliary case. If independent capacity testing places J below the relevant threshold, suffering may be high without significant agentive disruption or represented tissue damage. Behavioral effectiveness alone would not establish that result. The appraisal can be inaccurate; phenomenal grip does not require that the world validate it.

Likewise, a **brief overwhelming pain** may cross a suffering threshold immediately, but an agency theorist may classify its immediate psychological effects as significant disruption. This case is suggestive rather than load-bearing.

To state the implication claims precisely, let S_B, J_B, H_B be binary indicators for occurrent suffering, significant agentive disruption, and significant residual independent harm under fixed thresholds. The strengthened resilient case is the principal challenge to necessity:

$$S_B=1 \nRightarrow J_B=1, S_B=1 \nRightarrow H_B=1.$$

6.2 Actual agency disruption without suffering

Now consider Brady's **blissful disruptive love** case (Brady 2023, 29). The subject loses sleep, neglects plans, and functions poorly because attention and motivation are dominated by a positive state. Actual agency is disrupted; negative valence is absent. Positive capture is not occurrent suffering on the phenomenal target.

In **unnoticed discrimination**, a student's scholarship is secretly blocked by a discriminatory process. She is wronged and loses an opportunity while remaining cheerful because she has not learned what occurred. Whether J is high depends on whether the specified social agentive form includes effective access to that opportunity, so this case establishes H, not J, independently of present suffering. She may be correctly described as suffering *under* an unjust institution in a condition-predicating sense while not occurrently suffering in the phenomenal target sense.

A **deeply sedated patient** can lose consciousness while disease and lost agency remain or worsen. If sedation genuinely abolishes experience, it ends occurrent agony during that interval. It may simultaneously impose a serious capability loss and leave the disease untouched. Calling the patient free of current suffering does not imply that the intervention is harmless or all-things-considered desirable.

A **serene locked-in subject** presents the same structure without pharmacological intervention. Radical motor incapacity can coexist, at least conceptually, with calm. The world/capability ledger is grave; the phenomenal ledger need not contain aversive grip.

These cases support:

$$J_B=1 \nRightarrow S_B=1, H_B=1 \nRightarrow S_B=1.$$

6.3 Why the disagreement is not merely verbal

An agency theorist can deny some case descriptions or expand the relevant form of agency until each intuitive suffering case contains disruption. But an unconstrained expansion has a cost. If *significant disruption* includes any phenomenal alteration regardless of reliable capacity, the criterion has absorbed an experiential condition. It needs a prior rule specifying which agentive forms, which baseline, what aggregation, and what degree count. Momentary phenomenal narrowing does not by itself establish significant capacity degradation; successful outward performance does not by itself prove its absence. Conversely, the experientialist must not treat absence of S as evidence that J or H is absent.

The deepest disagreement concerns cases such as sedation and unnoticed injustice. On a non-experiential biconditional, more lost agency can mean more suffering even as consciousness disappears. On the two-ledger account, consciousness-ending sedation can produce phenomenal relief and world-level deterioration at once. These descriptions recommend different measurements and expose different moral trade-offs. The difference is not exhausted by which noun we prefer.

The following matrix systematizes these inherited contrasts and adds the independent-harm dimension; it is not offered as the discovery of the underlying double dissociation.

Case	Experienced valence/capture A,C	Actual agency disruption J	Independent harm H	Occurrent suffering S
Disliked paper cut, quickly ignored	low	absent	low	trace/below ordinary threshold
Severe headache with rumination	high	high	variable	high
Equally intense headache while absorbed and functioning	A fixed high; C lower	low	held fixed	lower
Fitting grief while functioning	high	low	present loss	high
Blissful but disabling love	positive capture	high	variable	absent
Unnoticed discrimination	absent	indeterminate	high	absent now
Deep sedation with untreated disease	absent	high	high	absent while unconscious
Phantom agony without tissue damage	high	low if independently verified	no independent damage	high

An ordinary many-to-one aggregate does not by itself reveal which component changed or identify an unmeasured component. A deliberately injective encoding could preserve every component, while a validated aggregation may combine them for a specified decision. The matrix preserves target identity before any normatively defended aggregation occurs. The question is now what can be inferred when only one side is observed.

7. The Relief-Repair non-identifiability result

The strongest consequence of the architecture is epistemic. Evidence that an intervention relieved observed suffering does not, by itself, identify whether it repaired an underlying harm or restored capability. Evidence of repair likewise does not identify relief. Randomization does not remove a missing-outcome problem when observations come from only one ledger.

For a minimal model, collapse each ledger to a binary adverse outcome:

- $S_t=1$ when a pre-specified, validated phenomenal protocol classifies occurrent suffering at time t ;
- $W_t=w(H_t, J_t) \in \{0,1\}$, with $W_t=1$ when one pre-specified world/capability target remains unresolved at t ;
- $I \in \{0,1\}$ indicates whether an intervention occurs between $t=0$ and $t=1$; and
- $Y_t=S_t$ is the only recorded outcome.

The model class consists of structural causal models of the form

$$I \perp U, W_1 = g_W(I, U_W), L = g_L(U_L), S_1 = g_S(I, W_1, L, U_S),$$

where randomization makes I independent of exogenous variables U , L permits post-repair persistence, and both direct $I \rightarrow S$ and mediated $I \rightarrow W \rightarrow S$ paths are allowed. The function w must be fixed before analysis: it may select residual harm H , significant agentive disruption J , or a defended composite. The propositions apply separately to each chosen target; they do not identify both H and J merely by naming a composite. Identification is relative to a model class: a query is identified only if every pair of models in that class that induces the same observed law agrees on it (Pearl 2009).

For Propositions 1 and 2, the complete information set is the pair of interventional laws $\Pr(S_1 | \text{do}(I=0))$ and $\Pr(S_1 | \text{do}(I=1))$; no H, J, W outcome is observed. For Proposition 3, it is the analogous pair of laws for W_1 , with no S outcome observed.

Define the repair effect and relief effect as

$$\rho = E[W_1(0) - W_1(1)], \lambda = E[S_1(0) - S_1(1)].$$

Assume initially that every subject has $S_0 = W_0 = 1$. Compare two deterministic members of the model class:

Model	Suffering equation	World/capability equation	ρ	λ	Interpretation
M_Q	$S_1 = 1 - I$	$W_1 = 1$	0	1	quieting without curing
M_C	$S_1 = W_1$	$W_1 = 1 - I$	1	1	quieting through repair

Both models predict exactly the same observed result under every intervention value:

$$\Pr_{M_Q}(Y_1 = y | \text{do}(I=i)) = \Pr_{M_C}(Y_1 = y | \text{do}(I=i)).$$

When $I=1$, both yield perfect recorded relief: $Y_1=0$. But they disagree maximally about the unobserved target. In M_Q , the selected worldly adversity remains; in M_C , it is repaired.

PROPOSITION 1: RELIEF DOES NOT IDENTIFY REPAIR

If an outcome protocol records only the phenomenal outcome S , then the repair effect ρ is not identifiable from intervention-and-suffering-outcome data alone in a model class containing both direct relief and source-mediated relief, absent additional bridge assumptions.

Proof. The models above induce the same distribution for every observed variable under every $\text{do}(I=i)$, yet $\rho=0$ in M_Q and $\rho=1$ in M_C . Models compatible with the same complete observed interventional law therefore disagree about the target query. By the standard model-pair criterion, ρ is not identified. QED.

There is a second direction of underdetermination. Let $L=1$ represent a learned or otherwise self-maintaining phenomenal process persisting throughout the measurement interval, while only $Y=S$ is recorded:

Model	World/capability equation	Suffering equation	ρ	Interpretation
M_L	$W_1 = 1 - I$	$S_1 = \max(W_1, L), \forall L=1$	1	repair with phenomenal lag
M_N	$W_1 = 1$	$S_1 = 1$	0	no repair

PROPOSITION 2: PERSISTENCE DOES NOT IDENTIFY FAILED REPAIR

When a phenomenal process may persist after its selected worldly source is repaired, persistent observed suffering does not identify the absence of a repair effect.

Proof. Both models entail $Y_1=1$ under every $\text{do}(I=i)$, but $\rho=1$ in M_L and $\rho=0$ in M_N . Persistent suffering data alone therefore cannot decide between persistence of the source and persistence of the phenomenal process. QED.

The result is symmetric. Suppose $Z_t = W_t$ is the only recorded outcome:

Model	World/capability equation	Suffering equation	λ	Interpretation
M_W	$W_1 = 1 - I$	$S_1 = 1$	0	repair without relief
M_{WR}	$W_1 = 1 - I$	$S_1 = W_1$	1	repair with relief

PROPOSITION 3: REPAIR DOES NOT IDENTIFY RELIEF

If an outcome protocol records only the selected world/capability outcome W , then the relief effect λ is not identifiable in a model class that permits both persistent suffering and relief through repair.

The proof repeats the model-pair argument: the two models agree on the complete interventional law of Z and disagree on λ . These deterministic binary pairs also establish non-identification in any broader model class that contains them; no claim is made that every restricted model class has this defect.

Two routes make the missing target identifiable, but both add information. First, validated direct measurement of W , combined with positive randomized assignment, consistency, and no interference, identifies ρ by comparing intervention arms. Second, a bridge can substitute for direct measurement only if a jointly sufficient package is defended. One deliberately strong package assumes a known, intervention-invariant, injective unit-level relation $S_1(i)=b(W_1(i))$ for both treatment values, validated measurement of S , consistency, positive randomized assignment, and no interference. Then $W_1(i)=b^{-1}(S_1(i))$, so ρ is recoverable from the transformed outcome. The bridge itself rules out direct relief, post-repair persistence, or compensating phenomenal causes not encoded in W . Monotonicity or exclusion alone does not identify an entirely unmeasured target.

These propositions are not new causal mathematics. They apply a standard structural-model identifiability test (Pearl 2009). Their philosophical importance lies elsewhere. A metaphysical distinction determines which claims an apparently successful intervention warrants. If relief and repair were the same property, single-target measurement might suffice. Once they can dissociate, an inference from one to the other requires measurement or a bridge argument.

8. Ledger Integrity and the no-laundering norm

Non-identifiability alone does not generate a moral duty to measure every outcome. The practical principle requires a further normative premise: both targets are independently material to the decision under review.

CONDITIONAL LEDGER INTEGRITY

If (i) a decision bears materially on both occurrent suffering and worldly harm or agentive disruption, (ii) the adopted normative view assigns independent relevance to both, and (iii) no other direct evidence, informative bound, or validated bridge connects them, then improvement observed in one ledger does not by itself identify or establish improvement in the other. The cross-ledger effect remains unidentified from those data and should be reported as such, rather than as improved. Which outcomes merit measurement remains sensitive to stakes, feasibility, and measurement cost.

An evaluation **launders harm through relief** when it (i) observes a reduction in occurrent suffering, (ii) treats that reduction as establishing that damage, deprivation, or wrong has ended without measuring or modeling the latter, and (iii) removes the unresolved world-level entry from further concern. It **launders suffering through repair** when it (i) verifies correction of a worldly condition, (ii) treats that result as establishing that distress has ended, and (iii) withholds recognition or relief from the continuing phenomenal state.

The norm is deliberately symmetric. Much moral criticism rightly emphasizes the first error: resilience can make injustice disappear from welfare measures; adaptation can make disability or deprivation look less serious to outsiders; tranquilization can make an institution easier to manage without improving a patient's life. But the converse error is also common. A medically successful procedure can leave pain; acquittal or compensation can leave trauma; restored employment can leave depression; removal from danger can leave terror. Objective success does not entitle an evaluator to declare the subject experientially restored.

Suppose the relevant practical value of an action a is represented schematically as

$$U(a) = -\alpha S_1(a) - \beta J_1(a) - \gamma H_1(a), \alpha, \beta, \gamma > 0.$$

Here S, J, H are adverse outcome scores, with H defined residually so the additive illustration does not double-count S or J . A realistic evaluation may require interaction terms rather than additivity. Single-ledger data cannot generally rank actions by U . Two interventions can produce identical relief with different disability or wrong, or identical repair with different continuing agony. Calling all three terms *suffering* does not solve the underdetermination; it hides the weights and makes it unclear which outcome moved. Conditional Ledger Integrity therefore calls for separate outcomes when the decision genuinely depends on them and measurement is proportionate to the stakes. The weights still require argument; they cannot be smuggled into a global score.

8.1 Adaptation and the danger of a false dilemma

Borgar Jølstad's work on adaptation and illness severity is especially close to the present proposal. Jølstad argues that adaptation can reduce suffering and, other things equal, make a person better off; a theory of health-related well-being that simply ignores this benefit is implausible. At the same time, he recommends pluralism so that subjective improvement does not eliminate objective concerns about functioning or circumstance (Jølstad 2023). The two-ledger account accepts both points. Adaptation is not “merely subjective” in a dismissive sense. Genuine relief is a genuine good. If a condition is equally damaging but less awful to live through, something important has improved.

What does not follow is that the condition has been repaired, that its wrong-making features have disappeared, or that all-things-considered priority must fall in proportion to reported distress. Those further conclusions depend on the capability trajectory, independent harm, the aims of the institution, and the relevant distributive theory. Conditional Ledger Integrity thus avoids a false dilemma when both targets matter. We need not choose between denying adaptation's value and allowing adaptation to erase deprivation. The relief is credited in L_p ; the unaltered or altered capabilities are recorded in L_w ; their combined significance is a subsequent normative question.

8.2 Four intervention patterns

The architecture distinguishes four patterns that a one-dimensional outcome obscures:

Pattern	Phenomenal ledger	World/capability ledger	Illustrative form
Relief without repair	improves	unchanged or worsens	analgesia, absorption, or sedation reduces aversive grip while disease or injustice remains
Repair before relief	unchanged at first	improves	surgery, rehabilitation, restitution, or escape from danger corrects conditions before distress subsides
Joint improvement	improves	improves	an intervention relieves experience and restores functioning or removes harm
Trade-off	worsens or improves	moves oppositely	painful rehabilitation increases short-term suffering while improving future capability

These are not treatment recommendations. They are distinctions a clinical or political recommendation must preserve. The right action may sometimes prioritize immediate relief, sometimes repair, and often both. The framework does not decide the weights. It prevents a decision-maker from claiming to have achieved both after measuring only one.

9. Objections

Objection 1: Ordinary language permits non-experiential suffering

People say that an unconscious patient suffers from disease, that an unaware citizen suffers under injustice, and that a region suffers drought. The experiential analysis therefore legislates against ordinary use.

Reply. The paper does not offer a univocal analysis of every grammatical construction. It explicitly targets occurrent personal suffering: what is happening when a conscious subject suffers now. Condition-predicating uses of *suffers from* and *suffers under* can remain legitimate. Tate, Kious, and Brady all acknowledge variation in the concept or its uses. Target indexing improves rather than abolishes this pluralism: it asks a speaker to state whether the claim concerns present experience, an adverse condition, actual capability, or a wrong. The proposed lexical discipline is revisionary only where an unindexed use licenses a false inference.

Objection 2: “Grip” is Kauppinen or McClelland under a new name

Kauppinen already analyzes suffering through pervasive negative affective construal, while McClelland already makes disruption of conscious mental life constitutive. Negative phenomenal grip adds no new essence.

Reply. This is the strongest novelty objection, and much of it should be conceded. Grip is a lean synthesis within the experientialist family. It emphasizes a cognitively undemanding, concern-engaging organization; does not require a belief that change lies outside one's power; and allows passive capture without an active struggle to accommodate. But the paper does not rest its main originality on those refinements. Its distinctive proposal is architectural: the phenomenal and agentive-disruption targets should be separately measured, their dissociation yields formal non-identifiability results, and those results support a conditional symmetric no-laundering norm. A paper that advertised “valence plus attention” as an undiscovered analysis would fail its own literature review.

The same restraint applies to the language of control and priority. Metzinger already treats conscious suffering through loss of control, loss of autonomy, and contraction of attentional agency (Metzinger 2016). Nagel's recent functionalist preprint makes priority competition and preemptive attentional reallocation central to suffering-capacity (Nagel 2026). Failed Recapture therefore claims neither to discover attentional capture nor to derive suffering from scheduling alone. Its narrower difference is the distinction between source control and grip control, achieved settlement rather than mere option-availability, and the connection of that control architecture to the Relief-Repair models. Nagel also requires a self-model, persistent goals, and an existentially threatening signal; FR is deliberately permissive about primitive animals, grief, remorse, boredom, and non-existential pain.

Objection 3: Separating harm from suffering trivializes objective suffering

Objective theorists introduced non-experiential suffering to keep disability, oppression, and profound loss morally visible. Reclassifying those states as harm risks returning to an individualistic clinical picture in which only feelings matter.

Reply. The second ledger exists precisely to block that result. Harm and wrong do not need the name *suffering* to be real or morally urgent. Indeed, forcing them into a phenomenal measure can make them less visible when subjects adapt, remain resilient, lack information, or cannot report. Blumenthal-Barby's analysis of disorders of consciousness makes a similar terminological move: many concerns at the margins are better described as harms or wrongs. The present account adds that these entries must remain independently measured and must not be laundered by low distress. The world/capability ledger is not secondary in value; it is distinct in target.

Objection 4: Agency disruption is what grip really is

When pain captures attention or depression narrows possibilities, agency is disrupted. The alleged ledgers collapse.

Reply. Experienced capture C concerns apparent salience and response possibilities; J concerns reliable exercise of a specified agentive form. Either can cause the other, but momentary narrowing need not degrade reliable exercise, and

outward performance alone does not settle J. If every phenomenal alteration counts as J, the agency criterion becomes experiential; if C requires functional failure, the phenomenal target disappears. The variables require different evidence.

Objection 5: The analysis excludes infants and animals

Concern, appraisal, and an experienced field of possibilities sound cognitively sophisticated.

Reply. The process model requires neither language, reflective selfhood, nor a second-order judgment. An aversive episode can primitively organize orienting, withdrawal, attention, and motivation. *Concern* marks that organization, not a narrated value system. The empirical attribution of consciousness and valence remains open, so the model permits animal and infant suffering without defining it into existence.

Objection 6: Sedation shows the view is dangerous

If unconsciousness ends occurrent suffering, the theory seems to endorse sedation whenever distress is severe.

Reply. That inference violates Conditional Ledger Integrity whenever capability, harm, and suffering are independently relevant. Sedation may reduce S while worsening J, leaving H untouched, carrying risks, or conflicting with autonomy. If consciousness and negative experience are absent, occurrent suffering is absent during that interval; it does not follow that producing unconsciousness is permissible or optimal.

Objection 7: The formal results are obvious

Of course an unmeasured outcome cannot be learned merely by measuring a different outcome.

Reply. The proofs are elementary, and the paper labels them so. Their value is diagnostic. Philosophical and institutional language often suppresses the fact that there are two outcomes by placing both under *suffering*. Once the target is verbally unified, a fall in a distress score can appear to show a fall in severity, or successful repair can appear to show that no further relief is owed. The propositions turn a conceptual distinction into a precise limit on warranted inference and display the extra assumptions needed to cross it.

10. What the proposal answers - and leaves open

This paper offers a targeted reconciliation of a live dispute; it does not claim to solve a historically recognized open problem by discovering the uniquely correct essence of suffering. The reconciliation is substantive because occurrent experience and significant agentive disruption differ in truth conditions, measurement, causal dynamics, and normative role. The philosophical question is:

THE RECONCILIATION PROBLEM

How can experientialist and agency-disruption theories each track genuine features of paradigmatic cases while yielding incompatible analyses of suffering?

The proposed answer is not that Corns secretly analyzes an adjacent variable; she explicitly intends significant agency disruption as an analysis of suffering. It is that the rival criteria are best reconstructed as selecting distinct targets for inference and measurement. Occurrent suffering is a phenomenal target; Failed Recapture is a candidate mechanism for one family of episodes. Significant agentive disruption and residual independent harm are worldly targets that can come apart from present experience. Reconstructing the agency-theoretic concern as J preserves its truth conditions and moral force without forcing it into the occurrent target. Once the targets are indexed, the demand that either experience or actual agency do all explanatory work dissolves. The paper leaves open whether ordinary language should reserve *suffering* for one target or permit both uses.

A narrower inferential question is settled inside the stated model class: do complete interventional outcomes from either ledger identify the intervention's effect on the other when direct relief, source-mediated relief, and post-repair persistence are

allowed? The paired models prove that they do not. First, the inherited dissociation cases challenge any thresholded biconditional between S and J. Second, Propositions 1-3 show that outcomes in one target do not identify causal effects on the other and that persistent phenomenal symptoms need not show failed worldly repair. Third, Conditional Ledger Integrity states what follows when an independent normative view assigns material relevance to both targets: improvement in either ledger does not by itself establish improvement in the other.

Several larger problems remain unsolved:

1. The paper does not give a complete theory of harm, welfare, illness severity, capability, or justice. The H column intentionally permits competing normative theories.
2. It does not establish a unique empirical measure of occurrent suffering or prove that Failed Recapture outperforms Kauppinen's pervasiveness or McClelland's disruption in every case.
3. It does not settle every ordinary or extended use of *suffering*, nor forbid objective or condition-predicating speech when its target is clear.
4. It does not supply the bridge assumptions needed to infer one ledger from the other in a particular domain. Those are empirical and causal-modeling tasks.
5. It does not decide how much weight relief, agentive disruption, and residual independent harm should receive when they conflict.
6. It does not show that suffering is always disvaluable all things considered, that meaningful suffering is illusory, or that attention is voluntarily controllable.

The defensible contribution is an original synthesis and application: a candidate process model of experienced aversive capture; a two-target representation of phenomenal suffering and worldly disruption; paired causal countermodels establishing cross-target non-identifiability; and a conditional no-cross-ledger-inference norm. Kious already distinguishes kinds of suffering; Tate already develops subjective and objective dimensions; Blumenthal-Barby already argues that non-experiential cases may be harms or wrongs; Jølstad already combines the value of subjective adaptation with objective pluralism; Brady supplies the double dissociation; Brady, Kauppinen, McClelland, and Corns supply nearly all constitutive materials; and Metzinger and Nagel preempt generic claims about control or attentional priority. The paper claims no priority for experientialism, subjective/objective pluralism, attentional or control-based theories, Brady's cases, or the mathematics of causal identification. Priority for the integration itself remains provisional pending a broader historical review.

11. Conclusion

There is a difference between no longer hurting and no longer being harmed. There is also a difference between being repaired and feeling relief. A theory of suffering should not erase either difference in pursuit of one elegant scalar.

Occurrent suffering belongs to the phenomenal-control ledger; Failed Recapture models an important family of episodes in which aversive pressure captures attention and action while neither source resolution nor grip regulation restores a flexible practical field. Actual agency disruption, damage, deprivation, and wrong belong to the world/capability ledger. The ledgers interact continuously, but they can vary independently. That separation explains why the experientialist and agency theorist each sees something real and why neither variable should monopolize the other.

Once the distinction is explicit, an epistemic discipline follows. Relief data alone do not identify repair; repair data alone do not identify relief. Resilience is not exoneration. Tranquility is not cure. Restitution is not yet recovery. When both targets matter, Conditional Ledger Integrity requires us to credit every measured improvement while marking the cross-ledger effect unidentified from those data until it is measured, bounded, or bridged. That is the proposal offered here: not a new word for suffering, but a rule for keeping its experience and its world from being mistaken for one another.

References

- Arman, Shayan. 2025a. [“On Sickness”](#). Post 96, 7 June.
- Arman, Shayan. 2025b. [“On Suffering”](#). Post 145, 20 July.
- Arman, Shayan. 2025c. [“To Take a Stand”](#). Post 169, 14 August.
- Arman, Shayan. 2025d. [“Theory of Expressionism”](#). Post 206, 4 September.
- Arman, Shayan. 2025e. [“Pain”](#). Post 263, 23 October.
- Arman, Shayan. 2025f. [“On Suffering”](#). Post 283, 25 November.
- Arman, Shayan. 2025g. [“Et tu, Brute?”](#). Post 288, 6 December.
- Arman, Shayan. 2025h. [“Moments”](#). Post 289, 8 December.
- Arman, Shayan. 2026. [“On Depression”](#). Post 404, 17 July.
- Blumenthal-Barby, Jennifer. 2026. [“Suffering at the Margins: Non-Experiential Suffering and Disorders of Consciousness.”](#) *Cambridge Quarterly of Healthcare Ethics* 35: 49-57.
- Brady, Michael S. 2018. *Suffering and Virtue*. Oxford: Oxford University Press.
<https://doi.org/10.1093/oso/9780198812807.001.0001>.
- Brady, Michael S. 2023. [“Suffering as Experiential - A Response to Jennifer Corns.”](#) *Philosophy and Phenomenological Research* 107 (1): 24-30.
- Brady, Michael S. 2026. [“Why We Should Be Experientialists about Suffering.”](#) *Cambridge Quarterly of Healthcare Ethics* 35: 58-68.
- Corns, Jennifer. 2022. [“Suffering as Significantly Disrupted Agency.”](#) *Philosophy and Phenomenological Research* 105 (3): 706-729.
- Jølstad, Borgar. 2023. [“Adaptation and Illness Severity: The Significance of Suffering.”](#) *Medicine, Health Care and Philosophy* 26 (3): 413-423.
- Kauppinen, Antti. 2019. [“The World According to Suffering.”](#) In *Philosophy of Suffering: Metaphysics, Value, and Normativity*, edited by David Bain, Michael S. Brady, and Jennifer Corns, 19-36. London: Routledge.
- Kious, Brent M. 2022. [“Three Kinds of Suffering and Their Relative Moral Significance.”](#) *Bioethics* 36: 621-627.
- McClelland, Tom. 2019. [“The Disruption Model of Suffering.”](#) In *Philosophy of Suffering: Metaphysics, Value, and Normativity*, edited by David Bain, Michael S. Brady, and Jennifer Corns, 37-54. London: Routledge.
- Metzinger, Thomas. 2016. [“Suffering.”](#) Penultimate draft of a chapter in *The Return of Consciousness*, edited by Kurt Almquist and Anders Haag. Stockholm: Axel and Margaret Ax:son Johnson Foundation.
- Nagel, Jonathan J. 2026. [“Suffering as the Schedulability Price of Agency: A Functionalist Framework for Substrate-Independent Detection.”](#) Non-peer-reviewed preprint, 5 April.
- Pearl, Judea. 2009. *Causality: Models, Reasoning, and Inference*. 2nd ed. Cambridge: Cambridge University Press.
- Tate, Tyler. 2020. [“What We Talk About When We Talk About Pediatric Suffering.”](#) *Theoretical Medicine and Bioethics* 41: 143-163.
- Tate, Tyler. 2026. [“Objective Suffering: What Is It? What Could It Be?”](#) *Cambridge Quarterly of Healthcare Ethics* 35: 13-21.