Ease: Evaluation kills entry, not the state: a threshold model of ease

Florian Morin
2026
Independent Scholar
florian@florianmorin.com
florianmorin.com
Authoritative version: florianmorin.com/papers/ease

Abstract

Ease designates a class of stabilized engagement episodes characterized by reduced evaluative monitoring and the absence of deliberate control. These episodes reliably collapse when rendered objects of self-evaluation, resist deliberate pursuit, and scale weakly with explicit reward magnitude. Rather than treating positive affect as an output of reward computation or performance optimization, this paper proposes a control-level constraint on access: certain experiential states arise only when optimization pressure is temporarily relaxed, and are structurally incompatible with being targeted or monitored during entry.

A two-regime framework is developed, distinguishing optimization-dominated control from a permissive control configuration in which discrepancy can remain informational rather than imperative. The analysis predicts an asymmetric access dynamic: evaluative monitoring acts primarily as an entry barrier rather than as a continuous suppressor. If permissive conditions persist, a threshold transition can occur after which the episode becomes more robust to incidental awareness, with termination increasingly governed by fatigue-like constraints rather than renewed evaluation.

The contribution of this paper is not a new affective taxonomy or an intervention strategy, but a constraint on admissible explanations of positive engagement. The framework specifies when and why standard reward, learning, and performance models fail to predict experiential availability, and derives falsifiable signatures that distinguish fragile entry from stabilized persistence without presuming a specific neural or computational implementation.

This paper does not propose an intervention. It formalizes a constraint on when positive engagement can occur.

Application

The framework developed here is not intended as a design prescription or optimization guideline (see Scope, Box 1). Its primary function is diagnostic. It specifies a structural constraint on experiential access that applies across tasks, contexts, and domains, regardless of whether positive affect is a desired outcome.

One consequence is that many standard practices used to study, measure, or enhance engagement can inadvertently eliminate the phenomenon they aim to capture. Procedures that introduce explicit evaluation, success criteria, or real-time self-assessment may function as interventions on the control regime, selectively preventing access during critical entry periods. Null results obtained under such conditions should therefore be interpreted conditionally, rather than as evidence of absence.

A second consequence concerns interpretation rather than application. When positive engagement fails to scale with reward magnitude, resists deliberate pursuit, or collapses under explanation, this should not be treated as noise or subjective unreliability. Within the present framework, these patterns are expected signatures of a regime that is permissive rather than productive, and whose access conditions are asymmetric over time.

Finally, the framework imposes a constraint on future theorizing. Any account that aims to explain this class of positive affect must specify why evaluative monitoring is selectively destructive during entry, why its influence diminishes after stabilization, and why termination becomes dominated by fatigue-like limits rather than renewed optimization. Accounts that do not address these asymmetries remain incomplete, even if they successfully model reward learning, motivation, or performance. These commitments are stated formally, together with explicit falsification criteria, in Box 1.

The contribution of this work is therefore not an improved method for producing positive affect within optimized games or application, but a clarification of why such production repeatedly fails. It reframes ease as a control-state phenomenon with strict access conditions, and treats those conditions as non-negotiable constraints rather than parameters to be tuned.

Key Points

• A class of positive engagement episodes exists that is structurally incompatible with explicit self-evaluation during entry, regardless of reward magnitude or task performance.

• Evaluative monitoring acts primarily as an entry barrier rather than as a continuous suppressor: once stabilized, episodes become more tolerant to incidental awareness, with termination increasingly governed by fatigue-like constraints.

• Access to these episodes depends on a temporary relaxation of optimization pressure, not on reward accumulation, deliberate pursuit, or performance improvement.

• Many standard measurement and explanation practices function as interventions on control, selectively preventing access rather than revealing the phenomenon they aim to study.

BOX 1. Ease as a Non-Instrumental Positive Regime: Minimal Statement and Falsification Set

(canonical reference block, stable wording)

Definition (phenomenological).

Ease is a high-intensity positive regime characterized by (i) high subjective positivity and perceptual vividness, (ii) minimal self-monitoring, (iii) minimal instrumental intent, and (iv) a distinctive property: it does not accumulate into meaning, progress, or durable motivational value. Instrumental intent triggers evaluative capture and collapses entry.

Somatic marker (common report).

Ease is often accompanied by a robust chest-centered pleasure sensation. Increased sensory granularity and recovery of stimulus-specific processing styles previously consolidated under low monitoring load.

Core claim (mechanistic).

Ease is not produced by adding reward, pleasure, relaxation, or motivation. It emerges when anticipatory evaluative control fails to capture ongoing coordination. The key variable is the coupling strength between subcortical coordination (PAG-centered) and cortical monitoring systems, which scales with Z.

Guardrail (non-assimilation clause).

This hypothesis is not a re-description of flow, mindfulness, reward, or positive mood. The core claim is that the present regime is specifically the one that fails when it becomes instrumentable. If a framework predicts stable access under goal-maintenance, attentional training, or emotion regulation, it is likely tracking a different target.

Behavioral entry principle (operational).

Ease is facilitated by tasks that prevent anticipation and evaluation from stabilizing. This can be achieved through micro-actions that are:

(1) non-repeatable,

(2) metric-free,

(3) non-optimizable,

(4) rapidly self-terminating once evaluation appears.

A concrete demonstration domain is a library of Non-Use microtasks (e.g., game-like micro-actions), designed to die quickly when goal-monitoring emerges.

Z hypothesis (developmental).

Access to ease is primarily constrained by cumulative entropic load (Z), defined as the long-term accumulation of predictive-evaluative structure that increases the probability of anticipatory capture. High Z increases the stability of monitoring and reduces the size of the entry window.

Scope (what is *not* being claimed)

This framework does **not** claim that ease is universally desirable, morally superior, or the only valuable human state. It does **not** claim that ease maximizes productivity, mental health, or long-term wellbeing. It does **not** claim that the mechanism is fully localized to PAG, or that PAG is the sole necessary substrate. It does **not** claim that the described behavioral entry tasks will generalize across all individuals, cultures, or clinical populations. Finally, it does **not** claim that the theory is complete, only that it is falsifiable and operational.

Falsification set (non-negotiable)

F1. Instrumental compatibility.

If ease can be reliably entered and maintained while explicitly used for a goal (performance, productivity, emotional regulation), without collapse at entry, the hypothesis is false.

F2. Repeatable direct interventions.

If a repeatable direct intervention (stimulation, drug, training protocol) produces stable long-term access to ease without increasing anticipatory capture or dependence on the intervention, the hypothesis is false.

F3. Neural signature equivalence.

If the neural signatures of ease are indistinguishable from reward, pleasure, relaxation, or flow, the hypothesis is false. Specifically, if ease shows stable anticipatory coupling between PAG-centered coordination and evaluative cortex, the mechanistic core is false.

F4. Task design irrelevance.

If Non-Use microtasks constraints (non-repeatability, metric-free, non-optimizable structure) do not affect entry probability, the operational claim is false.

F5. Z irrelevance.

If proxies for Z (lifetime optimization load, sustained evaluative training, early competitive environments) do not predict ease access, the developmental claim is false.

For attribution and unambiguous reference, predictions are labeled A1-D1 and should be cited by label.

BOX 1B. Canonical Prediction Index (for citation)

Prediction set A, Entry barrier (monitoring as access constraint)

A1. Metric removal increases entry probability. Holding task content constant, removing evaluative affordances (score, HUD, progress indicators, success cues) increases the probability of entering an in-scope episode.

A2. Evaluative prompts collapse entry more than matched interruptions. Brief prompts that recruit self-evaluation (e.g., "rate your enjoyment," "are you in the state," "how well are you doing") are more disruptive during entry than equally salient but non-evaluative interruptions.

A3. Weak reward scaling under matched monitoring load. Within in-scope episodes, intensity and persistence correlate weakly with objective reward magnitude once monitoring load is controlled.

Prediction set B, Threshold and lock-in (time-dependent vulnerability)

B1. Time-dependent monitoring vulnerability. The disruptive effect of evaluative monitoring declines as a function of time since onset. Early probes are strongly disruptive, late probes are weaker.

B2. Asymmetric prevention vs termination. Small monitoring cues reliably prevent entry, but once lock-in is established, termination requires larger disruptions and is dominated by fatigue-like constraints.

B3. Persistence dissociation. Monitoring manipulations primarily shift entry and time-to-threshold, while post-threshold persistence is comparatively better explained by fatigue proxies than by evaluative affordances.

Prediction set C, Repetition and methodification (non-monotonic access)

C1. Non-monotonic repetition curve. Entry probability rises initially across early exposures, then plateaus or declines as the procedure becomes instrumentable ("methodification").

C2. Attribution penalty. Manipulations that increase explanatory capture (prompting subjects to describe the cause of the episode, labeling it as a technique) reduce subsequent entry probability under otherwise matched conditions.

C3. Structural variation restores access. Slight structural changes to the task (map, timing, micro-constraints) restore entry more than exact repetition, even when difficulty is matched.

Prediction set D, Z and development (individual differences)

D1. Z predicts entry, not persistence. Proxies for cumulative optimization load (Z) predict reduced entry probability and longer time-to-threshold, but weakly predict post-threshold persistence duration once entry occurs.

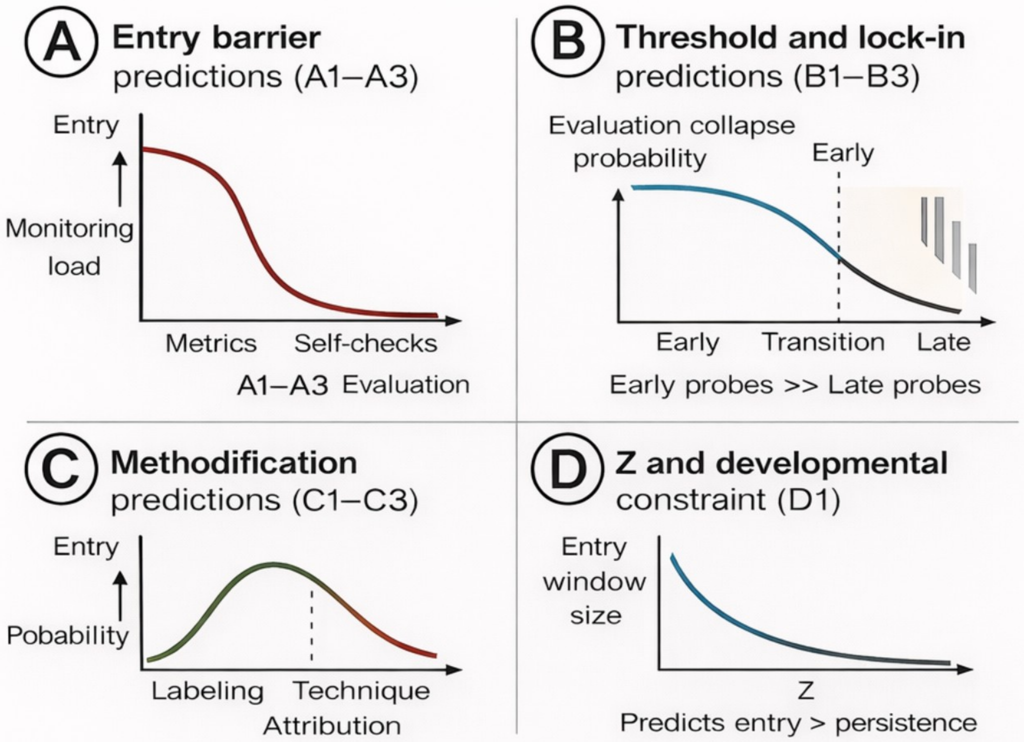

Figure 1. Canonical prediction map (A–D). (A) Entry barrier predictions (A1–A3). Entry probability is a steep function of monitoring load, with evaluative affordances (metrics, score, self-checks) acting primarily as access constraints rather than continuous suppressors.

(B) **Threshold and lock-in predictions** (B1–B3). The disruptive effect of evaluative probes is time-dependent: monitoring is maximally destructive early and progressively weaker after a transition point, with termination increasingly governed by fatigue-like constraints.

(C) **Methodification predictions** (C1–C3). Entry probability is non-monotonic across repetition: early exposures can facilitate access, but repeated instrumentalization and attribution reduce subsequent entry probability.

(D) **Z and developmental constraint** (D1). Proxies for cumulative optimization load (Z) predict reduced entry probability and longer time-to-threshold, with weaker effects on post-threshold persistence duration.

Predictions are labeled for citation (A1–D1). A minimal implementation and citation standard is provided in Appendix 4.

Canonical title for reuse (do not vary):

Ease as a Non-Instrumental Positive Regime: Minimal Statement and Falsification Set

The remainder of the paper develops, motivates, and tests the claims summarized in Box 1.

Introduction

Positive affect (PA) is commonly analyzed as a consequence of reward valuation, reinforcement, or hedonic amplification, where pleasure is treated as a scalar readout of acquired or anticipated benefit (Berridge & Kringelbach, 2015; Daw & O'Doherty, 2014). However, several robust features of PA are not fully accommodated by outcome-based explanations. In a specific class of episodes, PA can arise without explicit reward (Csikszentmihalyi, 1990), it can collapse under self-monitoring or instrumental pursuit (Wegner, 1994), and it shows a marked developmental asymmetry between childhood and adulthood (Arson, R. W. & Csikszentmihalyi , 1983). This collapse-at-entry property is treated here as definitional rather than incidental (see non-assimilation clause, Box 1). It also often fails to consolidate into stable preferences or durable action policies, repetition does not reliably increase intensity or motivational pull, and deliberate attempts to "make it happen" can be counterproductive (Wegner, 1994). Together, these patterns suggest that, for this class of episodes, PA is not well described as a reward output.

Thus, the present work advances a control-level account of a class of engagement states, referred to here as ease. In this restricted sense, positive affect is treated as a transient configuration of control rather than as a unitary affective construct. We distinguish optimization-dominated and permissive control regimes (defined in Section 2.1) and examine their empirical signatures. PA is predicted to occur only within a narrow, metastable region of this control space, where external or task-related input remains structured but is not immediately recruited for optimization. Prediction-error and free-energy language is used here as a cross-framework formalism, not as an assumption that the system implements a Friston-style predictive-processing model (Friston, 2010).

Two clarifications follow. First, suspended optimization is permissive rather than productive. This permissive-but-non-instrumental property is a core constraint of the framework (Box 1). It does not generate PA autonomously and can remain affectively neutral in the absence of structured external input such as visual, auditory, or interactive task dynamics. Second, given the monitoring-as-entry disruptor asymmetry (Section 2.2), the present study avoids in-episode probes and relies on delayed measures (Smallwood & Schooler, 2015). This constrains methodological choices by excluding intrusive measurement and favoring task-embedded designs in which PA can emerge mid-activity without evaluative reset. We formalize these regimes (Section 2), derive measurement implications (Section 12), and then introduce the proof-of-concept paradigm.

Threshold extension. Although suspended optimization is unstable during entry, we propose a threshold effect: if permissive conditions are maintained long enough, the control configuration stabilizes and evaluative monitoring loses leverage (Schooler & Wilson, 1991). The account implies a time-dependent vulnerability profile: reinstating self-evaluation is maximally disruptive early and progressively less effective after a transition point, after which termination is dominated by fatigue-like constraints (Csikszentmihalyi, 1990).

This account entails a game-based test via interface manipulation. By varying evaluation affordances (normal vs. semi-blind vs. blind HUD and feedback), we test whether reducing metric access selectively increases threshold-crossing probability and yields a time-dependent drop in monitoring vulnerability post-transition, while leaving fatigue as the primary termination pathway.

Throughout, we restrict the term positive affect to a specific class of episodes defined by preregistered inclusion and exclusion criteria (Section 2). Importantly, episode classification is performed without reference to HUD condition, and without using the predicted time-dependent probe interaction as an inclusion feature, preserving a clean separation between identification and hypothesis testing.

This work delimits a diagnostically asymmetric class of noninstrumental episodes that are reliably prevented by evaluative monitoring. The level of analysis is control organization and phenomenology, without commitment to a specific neural or computational implementation.

Ease is commonly experienced as a strong, chest-resonant pleasure, particularly responsive to music, cartoons, and simple visual stimuli. However, the phenomenon discussed here is difficult to evaluate purely from verbal description or third-person report. In our experience, its defining properties become apparent primarily through direct exposure under the specified control conditions. Strong judgments based solely on conceptual resemblance to existing practices or abstract interpretation are therefore unreliable for evaluating the phenomenon. The present framework functions as a constraint on testability and experiential access, rather than as a self-sufficient verbal characterization,

Part I, Problem, scope, and the constraint

A recurring asymmetry is observed across contexts: explicit self-evaluation reliably precipitates collapse of the target episode, whereas removing self-evaluation does not reliably restore it (Wegner, 1994). This hysteresis implies an asymmetric cost of evaluative access over time. Entry is fragile to appraisal, while re-entry is constrained by residual attribution, expectation, and control-policy adjustments induced during collapse (Wegner, 1994).

Within the Suspension of Optimization framework, the state is not merely weakened by evaluation. During the entry window, it is structurally incompatible with being rendered an object of optimization. Once evaluation occurs, the episode becomes more attributable and governable, and control policies shift toward increased evaluability. This shift persists beyond the evaluative act itself, reducing subsequent access probability even after evaluation is withdrawn.

This dynamic generates an apparent paradox in cultural and therapeutic contexts. Articulation is often assumed to confer insight and facilitate well-being, yet it frequently reduces access to the very state it describes. The resolution proposed here is that articulation is not a neutral report but an intervention on the monitoring regime (Schooler & Wilson, 1991). By increasing identifiability and predictability, it shortens detection latency and shifts the episode from experiential mode to target mode, improving reportability at the expense of entry conditions.

Explanation thus functions as an affordance for optimization. Once the state becomes describable, it becomes instrumentable and is drawn back into evaluation, tracking, and instrumental control. Speech and reflection are therefore not intrinsically distracting. Rather, articulation induces a control-mode shift that converts a self-maintaining trajectory into an object of assessment and improvement.

The resulting increase in attributional clarity strengthens monitoring policies and narrows the entry window. In consequence, explanation improves descriptive access while degrading the structural conditions required for initiation (Schooler & Wilson, 1991).

What the paper is and is not - framing it as a constraint, not a full theory

This paper characterizes a control-level constraint shaping access to a broad class of experiential episodes, referred to here as ease. The class is defined diagnostically by a characteristic asymmetry: access is reliably disrupted by evaluative monitoring, whereas removal of evaluation does not reliably restore it. The level of analysis is control organization, not phenomenological taxonomy, neural implementation, representational content, or computational architecture. No unique mechanism is assumed. Instead, the framework constrains admissible explanations by identifying a control condition that must be satisfied for these episodes to arise. Existing theories of reward, learning, and regulation are treated as compatible payload accounts whose phenomenological availability is conditional on this constraint. The scope and non-claims of this framework are summarized explicitly in Box 1.

Preview of the model: two control regimes and a threshold transition

Positive affect is treated here as a transient configuration of control (Carver & Scheier, 1982). Control systems are assumed to operate along a continuum between an optimization-dominated regime, characterized by continuous monitoring, evaluation, and corrective use of discrepancy, and a permissive regime in which optimization pressure is temporarily relaxed while engagement and perception remain intact. The regimes differ in how incoming variation is handled, not in sensory access. Under optimization-dominated control, discrepancies are rapidly recruited for correction. Under permissive control, discrepancies can remain briefly as experiential signals, creating a narrow, metastable region in which positive affect can emerge from structured sensory or task input.

This distinction can be expressed in prediction-error or free-energy terms without commitment to a specific architecture (Friston, 2010). In optimization-dominated control, discrepancy functions as an imperative and is rapidly converted into corrective action or updated expectations. Under suspended optimization, discrepancy remains informational: it is registered and experienced without automatically recruiting correction or self-evaluation

This altered handling of discrepancy allows novelty and variation to acquire positive phenomenological value without inducing corrective demand (Friston, 2010). A core empirical asymmetry follows. During entry into the permissive regime, positive affect is meta-unstable under self-referential monitoring (Wegner, 1994 & Wilson & Schooler, 1991). Redirecting attention toward evaluating, sustaining, or explaining the affective state tends to attenuate or terminate it, not because of distraction, but because monitoring reinstates optimization and converts experiential signals into objects of control (Wegner, 1994 & Wilson & Schooler, 1991). As a result, deliberate pursuit is often counterproductive, and the state is most reliably observed when it emerges incidentally within ongoing activity rather than through explicit initiation or evaluative reset Csikszentmihalyi, 1990).

The present paper extends this two-regime account with a threshold hypothesis. Although entry is fragile, the model proposes that if permissive conditions persist, the control configuration can stabilize such that evaluative monitoring loses leverage.

In this stabilized phase, monitoring becomes largely behaviorally irrelevant, and episode termination is dominated by fatigue-like constraints rather than renewed optimization pressure (Matthews, Parasuraman, Warm , 2008). This yields a specific time-dependent prediction: evaluative probes should be maximally disruptive early in the episode and progressively less effective after a transition point, with persistence thereafter limited primarily by resource depletion and attentional drift (Matthews & Parasuraman & Warm , 2008).

The introduction concludes by outlining a concrete test in a game-based task. Manipulating evaluative affordances within the interface, for example normal, semi-blind, and blind HUD and outcome feedback, isolates the predicted entry asymmetry: removing metrics selectively increases the probability of crossing the transition, with associated changes in persistence and temporal perception that can be captured without intrusive in-episode self-report (Kluger & DeNisi, 1996). This “Blind Play” manipulation functions as a falsifiable test of threshold crossing and post-entry persistence under reduced monitoring.

Definitions and formal commitments

What "optimization" means (control-theoretic, not moral)

Optimization refers to a control configuration in which perceived discrepancy is treated as an imperative for correction.

A schematic of the control landscape is shown in Figure 1 (Appendix 1) to support this conceptual distinction. The figure is illustrative only and does not represent a latent state variable, an energy function, or a continuous affective magnitude.

Minimal control formalization (gain and latency).

Let $\delta(t)$ denote a discrepancy signal (e.g. prediction error, mismatch, or deviation from expectation). Optimization is operationalized as the strength and latency with which discrepancy recruits corrective demand. Concretely, the controller generates a corrective drive

$c(t) = g\,\delta(t - \tau)$,

where $g \geq 0$ is the discrepancy-to-correction gain (the strength with which mismatch acquires corrective authority) and $\tau \geq 0$ is the effective latency from discrepancy detection to corrective recruitment.

Suspension of optimization corresponds to a transient reduction in g, an increase in τ, or both, such that discrepancy remains informational without automatically recruiting correction, strategy updating, or self-evaluation. Under this condition, deviations between expected and observed signals are registered without acquiring immediate corrective authority. Operationally, optimization is defined by the strength and latency of the discrepancy-to-correction coupling, how quickly and how strongly detected mismatch recruits evaluation and corrective policy. Suspension of optimization denotes a transient weakening of this coupling, allowing discrepancy to be experienced without triggering immediate correction or re-evaluation.

Two commitments follow.

1. **Optimization is functional, not evaluative in a value-laden sense.**

The term is used to describe how control systems handle discrepancy

2. **Optimization is defined by what discrepancy triggers**.

Under optimizationdominated control, discrepancy tends to recruit monitoring and correction quickly enough that phenomenological salience is reduced or instrumentalized before it can be experienced as such. By contrast, the permissive regime central to this account is defined by a temporary relaxation of this imperative: discrepancies can remain informational rather than imperative, meaning they are registered and experienced without triggering corrective control.

Suspended Optimization is treated as permissive rather than productive. It does not generate positive affect autonomously and may remain affectively neutral in the absence of structured external input, such as ongoing sensory or interactive signals. The account also sets boundary conditions: it does not target sustained mood states, pathological euphoria, or affect that scales robustly with explicit reward learning, and it excludes affective states that remain intact under explicit self-evaluation or that amplify monotonically with repetition.

Positive affect (PA) is defined here as brief-to-moderate episodes of non-instrumental positive engagement that arise during ongoing activity and exhibit a diagnostic asymmetry with respect to evaluation. Explicit self-evaluation during entry reliably disrupts emergence by inducing a stabilized evaluative control regime. Once established, this regime does not reliably disengage upon removal of evaluation. Non-instrumental denotes that onset is not attributed to an explicit reward or performance outcome. Episodes are defined at the level of phenomenology and behavioral continuity and are classified offline using preregistered delayed-report items together with logged behavioral signatures

What "monitoring" means operationally

Monitoring is the set of operations that reintroduce optimization pressure by converting ongoing experience into an explicit object of evaluation. In practice, this includes:
Outcome tracking, including performance metrics, score-like feedback, progress indicators, or success/failure cues.
Self-evaluation, attempts to assess whether the state is present, strong, improving, justified, or worth pursuing.
Success criteria, any explicit rule that turns the episode into a target to reach ormaintain.
Explanatory capture, the move from experiencing the state to explaining it, narrating it, or treating it as evidence about the self, goals, or meaning. A central commitment is that PA in-scope is meta-unstable under monitoring: directing attention toward evaluating, sustaining, or explaining the affective state tends to attenuate or terminate it early, not because of distraction, but because monitoring reinstates optimization and converts experiential signals into objects of control.

Inclusion rule

These diagnostic criteria operationalize the minimal commitments summarized in Box 1 and should not be interpreted as an induction procedure.
The criteria below function as diagnostic boundaries rather than prescriptive rules and are used exclusively for post-hoc classification. An episode is classified as in-scope if it satisfies at least three of the four core criteria and does not meet the hard exclusion.

Core criteria (need 3 of 4):

1. Weak reward scaling

Affective intensity and persistence correlate weakly with objective reward magnitude or explicit success, once monitoring load is considered.

2. Evaluation collapsibility during entry

During the early window, making the state an object of evaluation, sustaining, or explanation reliably attenuates or terminates it, beyond what matched nonevaluative interruptions do.

3. Mid-engagement emergence

Onset occurs more often mid-activity than at deliberate start, and deliberate initiation attempts tend to be counterproductive.

4. Entry vs. persistence dissociation (threshold signature)

Monitoring acts mainly as an entry barrier: early evaluative probes are highly disruptive, later probes are weaker after stabilization, with termination increasingly fatigue-limited.

Optional supporting feature (do not count toward the 3 of 4):

Weak consolidation into preferences

The episode fails to consolidate into stable preferences or durable policies, and repetition is non-monotonic due to "methodification."

Hard exclusion (any one is sufficient to exclude)

Exclude the episode if either of the following holds (these are already in your out-ofscope list, we are just making one of them "hard" for classification):
Monotonic reward scaling: affect scales monotonically with reward magnitude
or incentive value.
OR
Monitoring robustness: the affect remains stable under explicit selfevaluation, success criteria, or explanatory capture.

What counts as "positive affect" in-scope

It is qualified here as a family of positive affect episodes defined by their control conditions rather than by phenomenology in-scope PA is:

- Transient and context-dependent, it occurs under a non-default control configuration with asymmetric reversibility and does not behave like a stable trait or mood.
- Non-instrumental, it often arises without clear reward, and resists being stabilized by explicit pursuit or valuation.
- Collapsible under evaluation, explicit self-monitoring is sufficient to abolish the conditions under which it arises.
- Contingent on structured external input, suspended optimization is permissive rather than productive: it does not generate affect autonomously, and in the absence of such input, no positive affect is produced.

Features are neither necessary nor sufficient individually, the goal is to construct alignment and to prevent conflation with reward-scaled pleasure, mood, or distraction.

Example 1, reward-scaled pleasure (close, but out)

Episode: a strong pleasure response to a large reward outcome, e.g. winning a highstakes match with visible score, receiving money, eating a highly palatable food when hungry.

Why it is close: it is positive, can be intense, can occur during tasks.

Why it fails the diagnostic pattern:

-Fails Core 1 (weak reward scaling): intensity tracks reward magnitude or explicit success, by design. This violates the in-scope expectation that reward magnitude and PA correlate weakly once monitoring is accounted for.

- Often fails Core 2 (evaluation collapsibility): explicit appraisal ("this is great", "I'm winning") does not reliably terminate it, it can even amplify it.

This contradicts "meta-unstable under monitoring" as a defining property.

- It typically meets the hard exclusion if you adopt monotonic reward scaling as the hard line

Replication note: in experiments, these episodes confound your target class because metric visibility is part of the reward experience, not merely an entry barrier.

Example 2, conventional flow with heavy feedback dependence (close, but out)

Episode : skilled performance flow maintained by continuous feedback and rapid error correction, e.g. rhythm games, speedrunning with split times, high-precision aim training with immediate hit feedback.

Why it is close : it can look like absorption, time compression, sustained engagement.

Why it fails the diagnostic pattern:

Fails Core 4 directionally and violates the proposed flow separation:

classic flow is stabilized by active optimization loops, tight feedback coupling, and rapid error correction, and removing feedback often degrades it. By contrast, the in-scope class predicts the opposite pattern, metric removal should facilitate access by reducing evaluative affordances

Often fails Core 2 (evaluation collapsibility): evaluation may disrupt, but it does not show the strong early collapse signature that defines entry window dynamics.

May violate the hard exclusion if it remains largely stable under explicit performance monitoring because performance monitoring is constitutive of the state.

Replication note: Keep flow as an explicit contrast condition, because it provides a clean differential prediction under HUD removal and evaluative probes.

Example 3 simple novelty-only attentional capture (close, but out)

Episode: a “wow” response driven primarily by novelty, e.g. first exposure to a new track, a new map, a new visual stimulus, a new gadget.

Why it is close: it can feel positive, salient, sometimes time-distorting.

Why it fails the diagnostic pattern:

- Hits the stated exclusion: ‘simple attentional capture or novelty-only.
- Fails Core 3 (mid-engagement emergence): it often peaks at onset rather than emerging mid-activity, and is less about a permissive control configuration and more about stimulus-driven salience.
- **Fails Core 4 (entry vs. persistence dissociation): importing the “threshold/lock-in” story is unnecessary here, novelty typically decays with repetition even if monitoring is minimized, so the predicted time-dependent probe vulnerability is not the key signature.

The threshold and lock-in, core terms

The term threshold does not denote a discrete psychological event or a moment of subjective realization. It refers to an inferred change in control regime, defined by a shift in how evaluative monitoring impacts ongoing engagement. The threshold is not identified by any single behavioral marker, but by a change in causal sensitivity: after the transition, evaluative intrusions lose their capacity to interrupt the episode. Behavioral continuity measures function only as convergent indicators that such a regime change may have occurred, not as defining features of the threshold itself.

The framework already commits to a metastable region in which positive affect can occur when evaluative monitoring is relaxed, and to a reliable collapse when monitoring is reinstated during entry. The threshold terminology sharpens this account by making the entry–persistence asymmetry operationally testable.

- Entry window: the early period after permissive conditions begin, during which PA can appear but remains highly vulnerable to reinstated monitoring, especially
- Self-referential evaluation.
- Threshold: a transition point at which the permissive control configuration becomes sufficiently stabilized that evaluative monitoring no longer reliably reasserts effective corrective control. This term is used to make a concrete prediction: manipulations that reduce evaluation affordances should increase the probability of reaching this stabilized phase.
- Lock-in: the stabilized phase in which the episode persists with reduced susceptibility to evaluative probes. In this phase, the state is expected to be limited less by reactivated monitoring and more by ordinary constraints, mainly fatigue, on continued engagement, which later sections operationalize via persistence and indirect markers such as engagement continuity and temporal perception. This terminology is intentionally minimal: it does not presume a unique neural implementation and is compatible with multiple mechanistic realizations. Its role is to make the entry vs. persistence distinction experimentally tractable, including in gamebased paradigms that manipulate metric visibility, for example blind vs semiblind vs. normal UI. To make the core dynamics concrete, Figure 2 schematizes the threshold transition from fragile entry to lock-in under suspension of optimization.

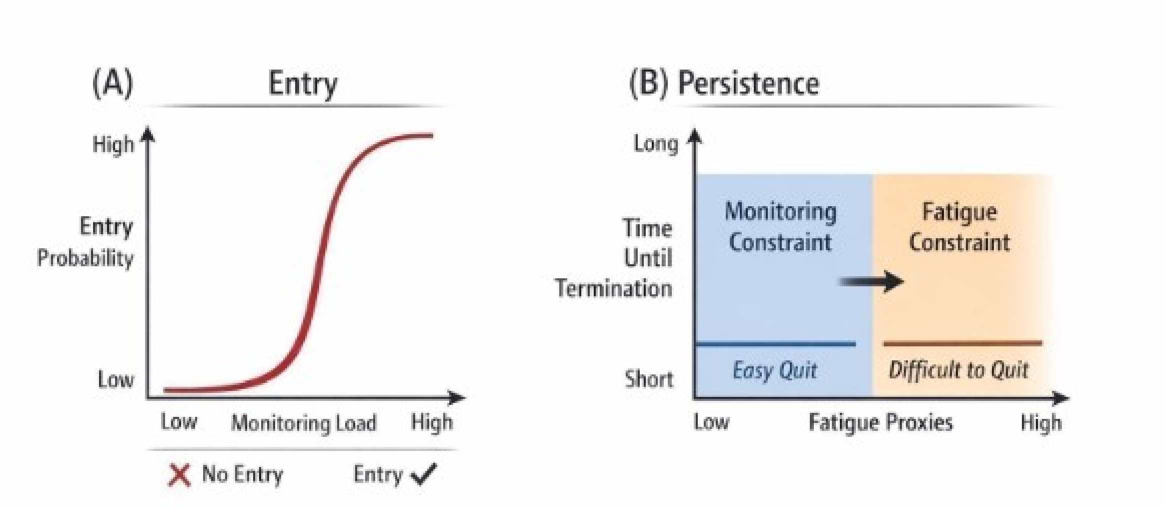

Figure 2. Entry–persistence dissociation schematic.

(A) Entry. Probability of entry is a binary outcome (entry vs. no entry) plotted against monitoring load. Entry shows steep sensitivity to monitoring, with small increases in evaluative load sharply reducing entry likelihood. This panel depicts vulnerability at access, not graded state strength.

(B) Persistence. Once entry has occurred, persistence is indexed by time until termination and is primarily predicted by fatigue-related proxies. Monitoring effects are weak or flat in this regime. Termination reflects resource depletion rather than renewed evaluation.

Together, the panels formalize a two-regime model in which entry and persistence are governed by distinct constraints, eliminating any notion of gradual escalation into the state and aligning directly with the model's falsification criteria.

Part II, The model

The key reframe

The central reframing is that positive affect (PA) is treated as a property of control organization rather than as an output of reward computation. The relevant causal lever is the temporary relaxation of a mature optimization regime that normally converts discrepancy into correction. Under optimization-dominated control, discrepancies between expected and observed signals are rapidly evaluated, prioritized, and recruited for adjustment. Under suspended control, discrepancies remain informational rather than imperative and can be experienced without automatically recruiting correction or self-evaluation. Suspended optimization is permissive rather than productive. It enables phenomenological salience to arise from structured ongoing input but does not autonomously generate affect. The functional commitment is asymmetric: reinstating evaluative monitoring reliably abolishes the state, whereas suspending monitoring does not reliably restore it.

Structured external input requirement

Because suspended optimization is permissive rather than productive, the account imposes a boundary condition: PA does not arise without structured external input. Relaxing optimization increases receptivity to incoming signals, but does not produce a content-free positive state. When the environment supplies structured variability, for example visual, auditory, or interactive task dynamics, those signals can acquire phenomenological salience because they are no longer immediately consumed by corrective demand.

In the absence of such input, the suspended state can remain affectively neutral rather than euphoric. The early phenomenology under relaxed evaluation may even feel initially neutral or uncanny, because salience is no longer automatically interpreted through goals, values, or self-referential meaning, and incoming signals may not “justify themselves” instrumentally. With limited exposure, the control system can adapt to this noninstrumental context, allowing PA to be integrated without evaluative endorsement, but the necessary condition remains that there is structured signal to receive.

Why monitoring collapses PA

A defining feature of PA under this account is meta-instability under self-referential monitoring. When attention is redirected toward evaluating, sustaining, or explaining the affective state itself, PA typically attenuates or terminates rapidly. This collapse is not attributed to attentional distraction or affective weakness, but to the reinstatement of optimization processes. Monitoring reintroduces evaluative comparison and performance tracking, converting experiential signals into objects of control. Because PA depends on the temporary absence of that evaluative pressure, detecting PA as a target state can be sufficient to reactivate the mature optimization regime and extinguish the conditions under which PA arose.

This mechanism provides one way of accounting for a commonly observed asymmetry: deliberate pursuit is often counterproductive, and PA is more reliably observed when it emerges incidentally within ongoing activity rather than through explicit attempts to induce or maintain it. It also motivates task-embedded paradigms that preserve continuity of engagement and minimize evaluative reset, because the primary threat to the state is not interruption per se, but the reintroduction of monitoring that restores optimization authority.

The threshold extension, from fragile entry to persistent episodes

Phenomenology of entry fragility

Early positive affect (PA) is phenomenologically fragile, not in the sense of being weak, but in the sense of being weakly anchored. Initial episodes often feel oddly neutral, impersonal, or insufficiently justified, despite the presence of clear sensory salience. Experience may appear vivid or absorbing without being immediately interpreted through goals, identity, or instrumental meaning. In ease, intensity reflects uncorrected persistence rather than stimulus strength.

During this early phase, salience is present but loosely integrated. Sensory or taskrelated signals can feel striking without yet cohering into a stable affective episode. The experience may lack narrative grounding and may not immediately register as "positive" in an explicit or self-referential sense. Subjects frequently report, and readers may recognize, an impression of something present but not yet named, owned, or endorsed. Ease is not defined by pleasure intensity, yet it is commonly experienced as a strong, chest-resonant pleasure, particularly responsive to music, cartoons, and simple visual stimuli. Pleasure during ease often feels disproportionate to the objective intensity of the stimulus. Simple sensory inputs, trivial events, or low informational content can produce unexpectedly strong positive affect. This apparent disproportion is not due to increased stimulus gain, but to the absence of evaluative normalization. When experience is not compared, ranked, or contextualized against alternatives, its intensity is not downscaled by reference frames. Awareness of the state often precedes its collapse phenomenologically. Entry is therefore marked by a distinctive combination of salience and instability, in which experience is noticeable but not yet consolidated into a durable episode. This entry window defines the phenomenological signature of fragility targeted by the present account.

An extended first-person descriptive account that informed the present framework is available elsewhere (Morin, 2025). This material is cited for phenomenological context only and does not function as empirical support for the model or its predictions. When present, the target state is frequently accompanied by amplified humor responsiveness, qualitative deepening of auditory experience, visually evoked effects reminiscent of earlylife affective responses, a non-instrumental comfort or warmth sensation, and enhanced imagery vividness.

The causal explanation for this fragility is not developed here and is treated separately at the control level in Section 3.3.

The threshold hypothesis

The preceding section characterizes entry phenomenology, but does not explain why some episodes, once present, can persist for extended periods despite intermittent awareness. The threshold hypothesis extends the two-regime account by proposing a transition in stability over time.

Specifically, if permissive conditions are maintained, the configuration supporting positive affect can stabilize such that continued awareness no longer reliably disrupts the episode. Stabilization here does not denote a discrete subjective event or an increase in affective intensity. Rather, it denotes a change in vulnerability: the episode no longer depends on continuous protection during entry and can persist despite incidental self-awareness. This transition marks a shift from a fragile entry phase to a more robust persistence phase. The progression from an initially uncanny or weakly valenced experience to an integrated positive episode is interpreted as stabilization under permissive conditions, not as hedonic escalation, reward accumulation, or increasing valuation. The defining feature of the threshold is therefore not phenomenological intensity, but altered sensitivity to disruption.

The threshold hypothesis does not posit a new latent state variable or phase transition in affect intensity. It posits only a change in the derivative of interruption probability with respect to evaluative intrusion.

Persistence mechanisms at the control level

The threshold extension is falsified if evaluative probes remain equally disruptive at all timepoints after controlling for interruption timing and salience, even if continuity, performance, or engagement metrics improve over time. In such a case, apparent stabilization would be attributable to generic warm-up rather than a change in control regime.

A practical implication follows for measurement. If early entry is uniquely sensitive to evaluative intrusion, then paradigms that embed induction mid-task, preserve continuity of engagement, and avoid explicit self-report during the episode should capture more of the relevant dynamics than designs that repeatedly query affect in real time. Indirect markers, such as engagement continuity, temporal perception shifts, or responsiveness to sensory variation, are predicted to be less disruptive than introspective ratings collected during the state itself.

Within this extension, termination after stabilization is expected to be dominated less by renewed optimization pressure and more by fatigue-like constraints, for example resource depletion, attentional drift, or performance breakdown that forces a reset of engagement. This fatigue dominance is a forward hypothesis that becomes testable once the entryvspersistence distinction is made explicit.

Clear falsifiers

Formal falsification conditions for the overall framework are summarized in Box 1; the present section details their operationalization.

The threshold extension makes a sharp prediction: vulnerability to monitoring should be time-dependent. Evaluative probes should be maximally disruptive early and progressively less effective after a transition point. If, instead, explicit selfevaluation remains equally able to terminate PA at any time point, then the lock-in claim is not supported and the model should revert to a purely entry-fragility account.

A second falsifier concerns instrumental scaling. If PA robustly tracks reward magnitude, task success, or improves monotonically with repetition, then the permissive-control interpretation is weakened, because the file's core constraint is that the state depends primarily on the degree to which monitoring is suspended, not on outcome-based valuation.

Finally, the extension yields a concrete operational test in a game-based setting. If removing evaluative metrics increases the probability of threshold crossing, and produces greater persistence plus a reduced effect of monitoring probes later in the episode, then the threshold interpretation gains support. If metric removal does not selectively affect crossing probability or time-dependent vulnerability, then the added threshold machinery is unnecessary.

Part III, Relation to existing literatures, as constraints not competitors

Reward, valuation, and hedonic accounts, where they fit, where they fail

What reward-based accounts explain well

Reward-based accounts explain an extensive and important class of affective phenomena. Hedonic and valuation frameworks capture how pleasure tracks the acquisition or anticipation of beneficial outcomes, how reinforcement learning shapes action selection through prediction error, and how incentive salience energizes pursuit and stabilizes preferences. Within these accounts, affect functions as a value-like signal that supports learning, motivation, and adaptive choice.

What they do not predict without additional constraints

The class of positive affect targeted here exhibits properties that outcome-based models do not predict on their own. These episodes frequently arise in the absence of explicit reward and are weakly related to objective success. They collapse under self-monitoring or instrumental pursuit, including attempts to evaluate, sustain, or explain the state itself. They also fail to consolidate into durable preferences or action policies, and repetition does not reliably amplify intensity or motivational pull. Without an additional constraint specifying when evaluative monitoring suppresses access, models that treat pleasure as a reward output predict strengthening with pursuit and repetition where collapse is in fact This mismatch is not a minor phenomenological detail, it is a structural asymmetry. If the act of assessing or pursuing the experience reinstates evaluation and extinguishes the conditions under which it arises, then intuitions derived from standard reinforcement frameworks can become misleading guides, more pursuit, more reward, more stability, become contrary guides for this class of episodes. The key issue is not that reward models are wrong, but that they often leave implicit the control conditions under which their reward-related signals become phenomenologically available rather than being immediately consumed by correction and performance tracking.

The compatibility claim

Reward, valuation, and hedonic mechanisms are treated as payload suppliers whose phenomenological availability depends on control state. The framework specifies the conditions under which their outputs can be experienced without being immediately instrumentalized by evaluative control.

Prediction error language without committing to a full predictive processing doctrine

The control-state account can be stated in prediction error terms while remaining agnostic about whether prediction error minimization is the primary objective of the organism, or whether any single predictive-processing architecture is sufficient to capture the full control stack. The claim required here is weaker and more testable: when optimization pressure is relaxed, discrepancy signals can be processed as experiential inputs rather than automatically recruited as control demands.

Discrepancy as signal vs. discrepancy as imperative

In many models of behavior, discrepancies between expected and observed signals contribute to learning and adjustment. Under optimization-dominated control, these discrepancies are treated as imperatives, they are rapidly evaluated, prioritized, and converted into corrective action or updated expectations. In other words, discrepancy is not merely detected, it is immediately assigned control authority, and the system behaves as if mismatch must be resolved.

Under suspended optimization, the same discrepancy signals are treated as signals, they are registered and can be experienced phenomenologically, but they do not automatically trigger correction, strategic updating, or self-evaluation. This does not require that prediction error disappears, nor that the system abandons learning. It requires only that the coupling from discrepancy to corrective demand is transiently weakened, such that mismatch can remain present without recruiting the evaluative layer that normally converts it into a performance problem.

This distinction clarifies why the present phenomenon is fragile under monitoring. Monitoring functions as a mechanism that restores imperative status to discrepancy by reintroducing evaluation and performance tracking, thereby collapsing the permissive configuration.

Why relaxed control allows novelty and variation to feel positive

When discrepancy is treated as an imperative, novelty and variation are immediately framed as deviations to be corrected, reduced, explained, or exploited instrumentally. This tends to compress phenomenological salience because incoming signals are rapidly consumed by optimization, either folded into goal pursuit, converted into error reduction, or used to update policies.

When optimization is relaxed, novelty and variation can persist briefly as experiential signals rather than as problems to be solved. In that permissive state, surprise can be felt as richness rather than as urgency.

When Reducing Feedback Helps Rather Than Hurts

Flow vs. suspended optimization vs. Mindfulness

Flow is typically maintained by active optimization: continuous performance regulation, tight feedback coupling, and rapid error correction. Removing feedback usually degrades flow because the loop that stabilizes skilled performance is weakened.

Conditions that improve flow by increasing feedback precision should reduce entry probability for suspended-optimization positive affect, and vice versa. Suspended optimization is defined by reduced evaluative control while engagement remains intact. It depends on structured input, but not on performance feedback. Removing metrics can increase access by lowering evaluation affordances, especially during the entry window. Mindfulness style accounts can overlap superficially because they reduce judgment, but many implementations still introduce a monitoring stance, a success criterion, or a meta-goal. When that happens, they function more like evaluation and can block entry in the same way.

We note that some low-feedback or internally guided flow variants may partially overlap phenomenologically with suspended optimization; however, they remain empirically distinguishable by their reliance on active error correction and performance coupling, whereas the present class is defined by entry fragility under evaluation and a selective increase in access when evaluative affordances are removed.

Empirical separations

Predicted differential effects of the same manipulations:

Feedback removal

Flow: decreases stability and performance coupling

Suspended optimization: little harm, may help by reducing evaluation cues

Metric removal (score, HUD, ranking)

Flow: mixed, may reduce goal anchoring

Suspended optimization: increases threshold crossing probability

Self-evaluation probes

Flow: may disrupt but not necessarily terminate

Suspended optimization: strongly disruptive early, weaker after threshold if lock-in holds

Meta-consciousness and measurement reactivity as a central methodological obstacle

A core implication of the framework is that the target phenomenon is highly sensitive to certain forms of measurement. Measures that require participants to evaluate, label, or assess their ongoing state alter the control conditions under which the phenomenon occurs. Measurement therefore functions as a potential intervention rather than as a neutral readout. Direct assessment during ongoing activity can interfere with access, placing a methodological constraint on when and how the phenomenon can be observed.

Experimental validity depends on avoiding measurement operations that recruit explicit evaluation or state-checking during critical entry periods.

Implications for self-report

In-episode self-report violates this constraint by reinstating evaluative monitoring during ongoing engagement. Affective classification and subjective ratings are therefore collected after task completion or following a prespecified delay, once immediate evaluative carryover has decayed.

Reactivity as an ordinary design issue

Observation changes the state" is treated here as a standard case of measurement reactivity rather than as a metaphysical claim. Certain observation procedures recruit reflective or comparative processing that alters task engagement. For the phenomenon targeted here, such procedures function as interventions on control and therefore cannot be treated as neutral noise. As a consequence, access can be assessed only through task-embedded manipulations and indirect indicators; explicit in-task questioning reinstates evaluative monitoring and violates the access condition.

Box 2. Measurement practices that preserve engagement continuity

- Avoid in-episode evaluative questions. Prefer post-episode or delayed ratings.
- Use unobtrusive continuous proxies. Examples include engagement continuity, timing behavior, interruption frequency, and time-perception markers.
- Embed manipulations in the interface rather than the instructions. Alter metric visibility (blind vs. semi-blind vs. normal) instead of instructing participants to suppress evaluation.
- Minimize resets. Maintain continuous task structure to avoid goal redefinition or explicit state checking.

• Vary probe timing. When probes are used, compare early versus late introduction to assess differential disruption.

Prediction set A, entry probability depends on monitoring load

Prediction Set A tests the claim that evaluative monitoring acts as an entry barrier by manipulating monitoring load while holding task structure constant.

• A1. Metric removal increases threshold crossing probability.

Holding task content constant, reducing evaluative affordances (e.g. hiding score, HUD, rankings, performance summaries) should increase the probability of reaching the post-entry transition. The effect should be strongest when metrics are removed early, before the episode stabilizes.

• A2. Evaluative prompts collapse entry more than matched sensory interruptions.

Brief prompts that recruit self-evaluation (eg “rate your enjoyment,” “are you in the state,” “how well are you doing”) should be more disruptive than equally timed, equally salient but non-evaluative sensory interruptions (eg a neutral tone, a brief visual flash) that do not imply performance assessment. The key prediction is a specific interaction, evaluation content matters more than interruption per se.

• A3. Weak coupling between PA and objective reward magnitude.

Within the target class, PA intensity and persistence should correlate weakly with objective reward magnitude or explicit success, once monitoring load is accounted for. Large rewards can fail to increase PA under high monitoring, while modest or incidental outcomes can support PA under reduced monitoring, consistent with PA depending on permissive control-state rather than reward maximization.

Prediction set B, threshold and lock-in signatures

Prediction Set B tests whether access dynamics are better described as a state transition with temporally asymmetric vulnerability rather than as continuous reward accumulation.

Time-dependent vulnerability

The critical prediction of the threshold account is not the presence of a smooth stabilization curve, a particular latency distribution, or any specific warm-up profile. The decisive signature is an interaction: the disruptive effect of evaluative monitoring must decline as a function of time since onset. Any account based solely on practice, familiarity, or motor warm-up predicts monotonic performance stabilization, but does not predict a selective reduction in sensitivity to evaluation content over time.

Claim: evaluation is most destructive early, and becomes less destructive after a measurable transition point.

Operational signatures in-game

Early window: small evaluative events (checking score, reinterpreting performance, asking "is it working") sharply reduce persistence probability.

Late window: the same evaluative events have weaker impact, often producing only a brief attenuation.

Measure

Survival curves of episode duration as a function of when evaluation is introduced (early vs. late injection).

Hysteresis

Claim: returning to the same context after collapse is harder than staying continuous.

Operational signatures

After a collapse, re-entering the same map and settings yields lower entry probability than uninterrupted continuation, even when sensory input is matched. "Context equivalence" fails, continuity matters.

Measure

● Re-entry probability conditioned on (a) uninterrupted play vs. (b) a reset of the loop (menu, break, explicit reflection), holding stimuli constant.

Asymmetric reversibility

Claim: it is easier to prevent entry than to terminate an already locked-in episode.

Operational signatures

Small optimization cues can block entry reliably.

Once locked-in, terminating requires larger disruption, typically fatigue or strong re-optimization demands rather than mild monitoring.

Measure

Compare minimal intervention strength needed to (a) prevent entry vs. (b) terminate an ongoing episode, fit separate thresholds.

Minimal model prediction

Together these imply a two-phase trajectory: entry barrier (high sensitivity to evaluation) followed by lock-in (reduced monitoring engagement), with fatigue as the main exit route once lock-in is established. This is the distinctive pattern that separates "suspension of optimization" from standard reward or performance accounts.

Prediction set C, repetition and "methodification"

Prediction Set C tests whether repetition alters access probability via methodification and attribution rather than via reinforcement-like strengthening.

Predictions

The schematic below summarizes the control landscape implied by the predictions that follow, highlighting the asymmetry between entry and collapse under optimization : under optimization (Supplementary Fig. S2).

Empirical inclusion fingerprint

An in-scope episode is defined by a minimum empirical fingerprint in behavior, not by reported phenomenology.

1. First, reducing access to metrics and evaluative cues (eg HUD removal, score concealment, no self-rating prompts) must increase the probability of entry relative to an otherwise matched condition.
2. Second, brief evaluative probes must show a strong time-dependent effect; they are disproportionately disruptive early in the episode and progressively lose disruptive power after a transition point (an entry-by-time interaction rather than a uniform main effect). 3 Third, once the transition has occurred, episode duration should be explained primarily by fatigue proxies (e.g. time-on-task limits, reaction-time slowing, error accumulation, pause frequency, subjective tiredness after the block) rather than by evaluation load or the availability of monitoring channels.

◦ **Non-monotonic repetition curve**: entry probability rises for the first few exposures, then plateaus or declines as the procedure becomes recognizably instrumental.

◦ **Attribution penalty**: the more the agent can explain the episode as caused by "the trick," the lower subsequent entry probability under otherwise identical conditions. This includes any repetition or instrumentalization of the initial micro-movement.

◦ **Generalization asymmetry**: slight structural changes to the procedure can restore access, whereas perfect repetition accelerates decay.

◦ **Persistence dissociation**: repetition mainly affects entry, not duration after **lock-in**, which remains primarily limited by fatigue.

Operational tests (game-based)

• Track entry rate across repeated sessions using the same HUD and task constraints.

• Add a brief “labeling” manipulation (prompt to describe what caused the state) and test whether it lowers subsequent entry probability.

• Compare exact repetition vs. minor structural variation (different map, timing, or micro-constraint) while holding overall difficulty constant.

Figure 3 formalizes the model’s key predictions for how monitoring load, time since onset, and repetition shape entry probability and vulnerability to interruption.

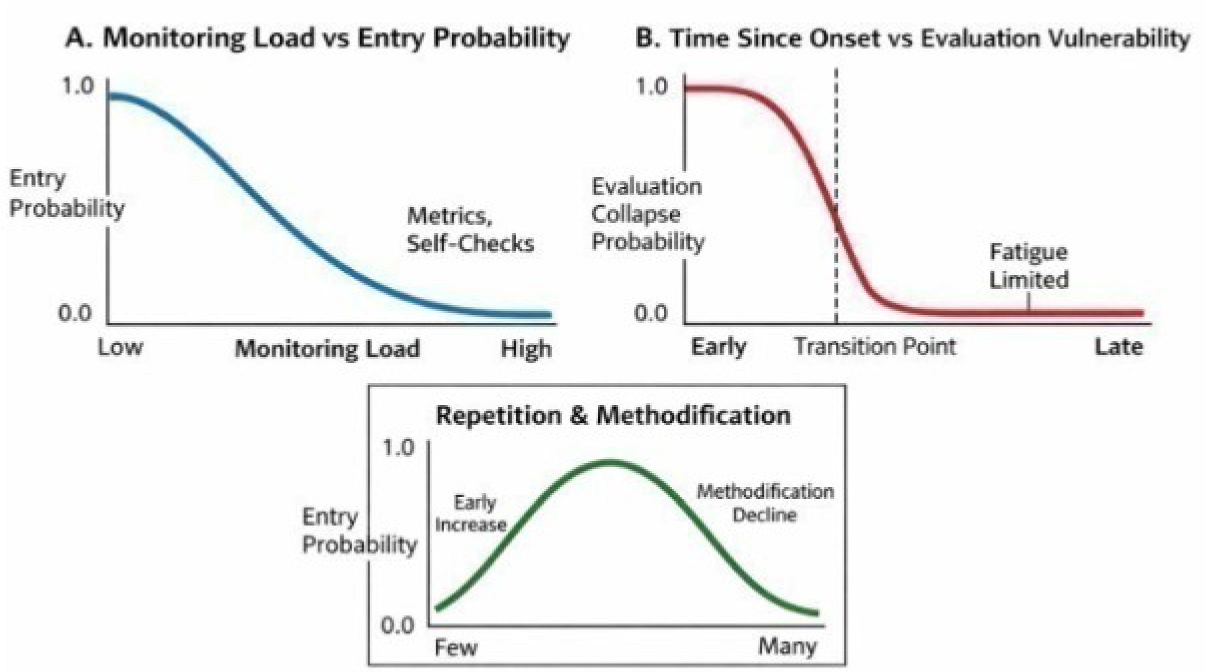

Figure 3. Predicted signatures of entry, lock-in, and methodification. Predicted signatures of entry, inferred stabilization, and methodification. Curves are schematic and intended to illustrate predicted dissociations between monitoring load, time since onset, and vulnerability to evaluative interruption.

They do not represent fitted trajectories or assumed warm-up dynamics.

Proof-of-concept behavioral study, Non-Use microtasks (UT99 implementation)

Participants are screened for in-scope episodes using the criteria in Section 2

Game task and control manipulation (Unreal Tournament)

The behavioral paradigm is implemented in Unreal Tournament, a fast-paced firstperson game that allows fine-grained control over interface and feedback parameters through built-in options and configuration files. The game can be parametrized to remove or attenuate evaluative information (e.g., score, timers, performance indicators) while minimizing changes to perceptual input, motor requirements, and core task structure.

We therefore treat HUD manipulation as a way to alter evaluative affordances, and we explicitly quantify any residual effects on performance-grounded difficulty metrics as manipulation checks (Section 12.4).

Difficulty calibration (low-optimization requirement).

To prevent performance regulation from dominating control, the task must be set below the participant's skill ceiling. If difficulty is high, frequent errors, planning demands, and corrective monitoring increase optimization coupling and reduce entry probability independent of the HUD manipulation. We therefore include a brief calibration phase and set difficulty to an under-challenging level (e.g., low-skill opponents/bots, reduced opponent density, simplified movement constraints, slower pace), while holding settings constant across UI conditions within a session. Within each session, participants engage in continuous play under a fixed HUD configuration (normal, semi-blind, or blind). No explicit instruction to seek or evaluate affective states is given. A brief, non-repeatable perturbation is externally cued once per block (e.g., via a neutral auditory tone or a simple instruction to briefly adjust posture), without reference to its purpose and without requiring participant choice. The perturbation is framed as a routine procedural element and is not described as an induction, task, or technique. If repeating this micro-movement were to increase entry probability or stabilize access, this would contradict the model by reintroducing instrumental control and methodification. Its inclusion is therefore logically constrained to a single, non-repeatable perturbation, rather than a method, cue, or induction procedure.

Low-demand activities that preserve engagement continuity without discrete triggering, such as brief periods of unstructured walking, are treated differently and may occur more than once without violating this constraint, provided they are not framed or used as intentional inductions.

Unstructured walking is considered a permissive background activity rather than a perturbation and is therefore not subject to the same nonrepeatability constraint.

Conditions

Participants play the same game under three interface conditions that vary evaluative affordances: normal UI (full metrics and feedback), semi-blind UI (reduced metrics), and blind UI (minimal or no performance indicators).

Procedure

Pre-task context standardization.

Prior to task engagement, participants are asked to ensure that no external timekeeping device is visible or consulted during the session and that no time constraints are anticipated (e.g., upcoming appointments). No instructions are given regarding internal dialogue, thought content, or attentional control. The protocol does not assume or require the absence of spontaneous self-evaluation or internal checking by participants; only externally available evaluative affordances are manipulated.

Participants are not asked to engage with, suppress, or modify spontaneous inner speech if it occurs.

Because this class of PA is entry-fragile under self-referential monitoring, standard affective measures that rely on explicit reporting may underestimate occurrence when administered during entry. The threshold extension turns this limitation into a testable prediction: measurement reactivity should decrease with time since onset. Evaluative probes should be most disruptive during early emergence but weaker after stabilization, when monitoring becomes informational rather than corrective. This motivates a twostage measurement strategy: protect the entry window using embedded or indirect behavioral proxies, then apply explicit ratings after a prespecified delay or after behavioral indicators of stabilization.

Accordingly, null effects in highly evaluative settings should be interpreted conditional on probe timing within this episode class, not as blanket evidence against the phenomenon.

Participants complete continuous play blocks under a fixed HUD configuration (normal, semi-blind, or blind), with condition order counterbalanced. No explicit instruction to seek, detect, or evaluate affective states is given. In-episode selfreport is avoided entirely.

To initiate engagement without inducing sustained self-monitoring, the protocol includes a single, brief micro-movement performed once at the beginning of each block (e.g., a minimal postural or motor adjustment). This movement is not repeated and has no instrumental function in the task, its role is limited to a transient perturbation without establishing a method or routine. No assumptions are made about participants' internal interpretation of this event; any spontaneous attribution or evaluation is treated as noise rather than a violation of the protocol. Additional examples of procedurally equivalent perturbation implementations are documented in Appendix 3, where they are treated as illustrative design variants rather than participant instructions.

To introduce a standardized procedural event that is temporally fixed and independent of participant intent, non-evaluative stabilization bias without requiring the participant to notice or label any internal change, a simple micro-task is imposed once per block for a strictly bounded interval (one round or one life only, depending on game mode). The micro-task is time-locked rather than state-locked, for example, it is initiated at the first natural boundary after a fixed elapsed time window (e.g., after ~5 minutes of play), signaled by a brief neutral cue. The micro-task is identical across UI conditions, deliberately minimal, framed as non-evaluative, and discontinued immediately at the end of the round/life. The micro-task is not assumed to reduce evaluation at the participant level; its role is limited to altering task structure uniformly across conditions All behavioral and interface events are continuously logged with timestamps (e.g., movement and action inputs, pauses, menu openings, HUD toggles or score checks where available, deaths/respawns, and session termination).

After each block, episodes are classified offline using the pre-specified 3-of-4 criteria, thereby separating entry dynamics from evaluative self-assessment during play.

Once a stabilization candidate has been reached, auditory stimuli may be introduced to facilitate phenomenological access to the ongoing state, without functioning as an induction or entry manipulation. Familiar music, including childhood-associated tracks, is used at this stage not to increase entry probability or persistence, but to enhance subjective recognizability of an already ongoing episode.

Participants are not instructed to suppress evaluation, monitor internal states, or aim for any specific experience. Instructions are deliberately minimal and task-focused. The use of such stimuli is explicitly restricted to the post-threshold phase and is not treated as a causal factor in access, stabilization, or maintenance. Its role is limited to supporting experiential reportability once evaluative vulnerability has decreased. After each block, the game was closed and participants underwent a short buffer period (e.g. 5 minutes) before completing delayed reconstruction items. During this interval, participants could optionally walk at a comfortable pace (e.g., 2–5 minutes). This buffer was intended to reduce immediate evaluative carryover and support postblock reconstruction without in-episode probing. The duration of the buffer and whether walking occurred were recorded.

A stabilization candidate is only accepted if an evaluative probe introduced after the inferred time point produces significantly less disruption than the same probe introduced earlier within the same block or a matched block.

Optional: Before each block, participants completed a short attentional-set instruction (3–5 s). In the salience version, they were asked to treat a neutral object in the room as briefly behaviorally relevant, without performing any action. In the control version, they were asked to observe the same object's perceptual features (e.g. shape, color) without assigning relevance. The induction was administered prior to play to avoid in-episode probing, and its timing and condition were preregistered. A single delayed manipulation-check item was collected post-block.

Using more than one delayed interval (e.g., 2 h vs. 6 h postsession) allows assessment of the stability of retrospective integration once immediate evaluative carryover has decayed, without introducing additional experimental conditions or inepisode probing.

Optional delayed follow-up.

An optional remote follow-up several hours after the session may be used to collect delayed subjective reports, further distancing assessment from the entry window and reducing evaluative contamination.

Primary endpoints

Key endpoints are entry rate (probability of onset), time-to-threshold (latency to a stabilized regime, inferred from logs), and persistence duration (post-threshold survival). Time-to-threshold is not interpreted as an individual-level estimate of a latent transition moment. It is treated as a distributional construct, defined only at the group level as a shift in the statistical relationship between evaluative probes, continuity hazards, and persistence.

(1) Entry rate (probability of onset)

Entry is defined at the block level as whether an in-scope episode occurred during that block. Episodes are classified offline using the pre-specified inclusion rule (3of-4 criteria plus hard exclusion, Section 2), combining delayed post-block reports with logged behavioral signatures. The primary test is whether entry probability differs across UI conditions (blind, semi-blind, normal).

(2)Time-to-threshold (latency to stabilization, inferred, no in-episode marking)

Timetothreshold is inferred as the earliest interval at which multiple preregistered stabilization indicators converge, conditional on later evidence that evaluative probes have lost disruptive power. Continuity markers are treated as necessary but not sufficient indicators of stabilization, and are interpreted in conjunction with timedependent probe effects. Continuity or hazard changepoints by themselves are not taken as evidence of threshold crossing, as similar patterns can arise from generic warm-up or familiarity effects. Continuity markers include interruption hazard (pause, menu, alt-tab, leaving match) and input continuity floors that exclude idle periods. This inference does not presuppose the absence of internal monitoring or a discrete subjective transition, and is treated as a statistical construct rather than an individual experiential marker. This operational definition is a proxy for stabilization that does not assume the mechanism a priori. Operationally, this latency is estimated from continuous logs using one of the following pre-registered definitions (or both, as convergent estimates):

(a) Change-point estimate: the earliest detected change-point at which the hazard of interruption (pause, menu, alt-tab, leaving the match) drops and remains lower for a minimum duration.

(b) Uninterrupted-window estimate: the onset of the first uninterrupted window of duration W (pre-registered, e.g. 3 to 5 minutes) during which interruption events remain below a pre-registered threshold (no pauses or menu openings above T seconds, no exits), and input continuity remains above a minimal floor (preregistered) to exclude idle periods.

Alternative-explanation control: to reduce confounding by generic warmup or familiarity, we pre-register a control analysis that (i) estimates changepoints in matched practice blocks where no in-scope episode is reported, and (ii) includes objective performance trajectories (eg accuracy, damage, deaths) as covariates, testing whether continuity change-points remain conditionsensitive after accounting for performance stabilization.

This endpoint tests whether reducing evaluative affordances shortens stabilization latency conditional on entry.

(3) Persistence duration (post-threshold survival)

Persistence is defined as duration from the inferred threshold time point to block termination. Termination is defined as either (i) voluntary stop or leaving the match, or (ii) a pre-registered interruption criterion (e.g., pause or alt-tab longer than T seconds), or (iii) end of the block (right-censoring). The key prediction is dissociation: UI condition primarily shifts entry and time-to-threshold, while postthreshold persistence is comparatively less sensitive to monitoring affordances and is better explained by fatiguelike drift.

Secondary outcomes (corroborative, delayed)

Secondary measures are collected after each block and treated as corroboration and moderators, not as timing signals: delayed affect rating, time-compression proxy ratings, post-episode recall/fragmentation, and subjective tiredness. Additional logged fatigue proxies include reaction-time drift where available, increasing micropauses, declining input density, error correction frequency, and voluntary continuation behavior.

Predicted dissociations

Warm-up / familiarity accounts predict:

- Improved performance metrics
- Reduced interruption from motor difficulty

- No selective interaction with evaluative contentThreshold account predicts:
- Reduced impact of evaluative probes specifically
- Persistence explained by fatigue rather than monitoring

• Entry probability affected more than post-threshold durationThe analyses in this section are designed to test these dissociations directly, rather than to identify a particular temporal profile.

Manipulation checks

Confirm that UI manipulations reduce evaluation and monitoring opportunities without altering objective task difficulty or required motor-cognitive demands. (1) Objective difficulty and motor demand invariance: across UI conditions, task settings are held constant and we test invariance using performance-grounded metrics that do not overlap with the primary continuity endpoints (eg kill/death ratio, accuracy, damage dealt/received, objective match outcome where applicable). Any residual performance differences are modeled as covariates rather than treated as manipulation success.

(2) Perceived difficulty: delayed difficulty rating after each block.

(3) Cognitive set adherence (delayed): brief post-block ratings of inner speech, imagery, planning, and task-unrelated mind wandering, treated as moderators rather than in-episode probes.

(4) Information removal check: in normal and semi-blind conditions, log and summarize any remaining access events to evaluative displays (e.g. scoreboard calls) to confirm the manipulation changes evaluative affordances as intended.

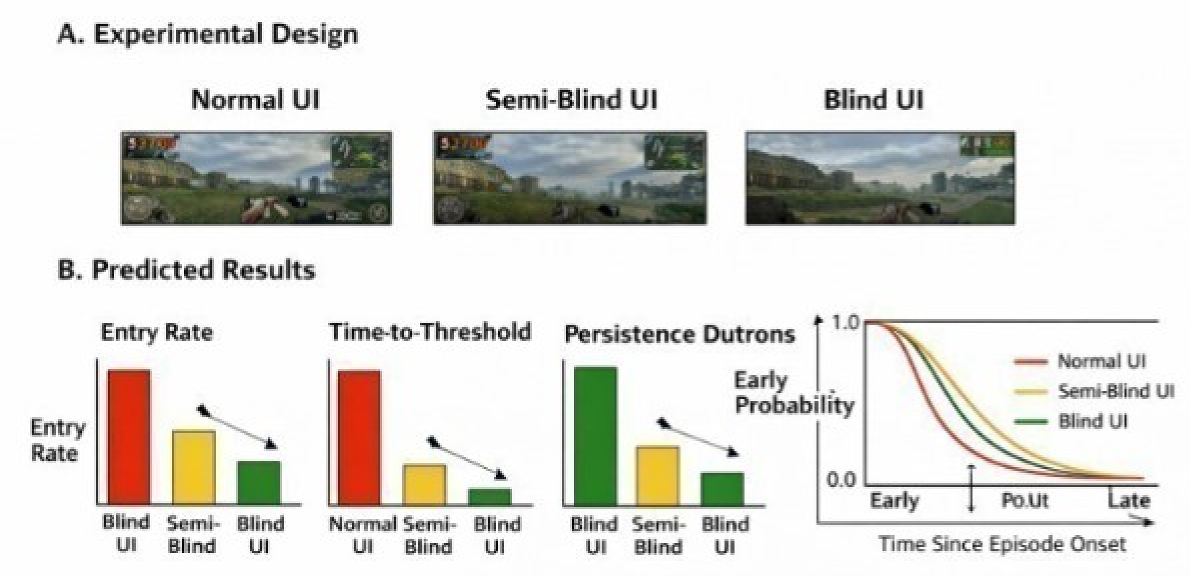

Figure 4. Experimental design and predicted results pattern. Participants perform the same game task under three interface conditions (normal, semi-blind, blind) that systematically vary evaluative affordances. The model predicts higher entry rates, shorter time-to-threshold, and longer persistence under reduced monitoring (blind > semi-blind > normal), with differences driven by entry probability rather than task difficulty. Evaluation probes are expected to disrupt engagement primarily early in the episode, while late termination is dominated by fatigue rather than renewed optimization.

Developmental and individual differences

Developmental prediction

The model predicts a systematic developmental shift in access to suspended optimization regimes. Earlier in development, control systems are less tightly coupled to continuous outcome evaluation and prospective correction. Engagement is more exploratory, discrepancy is more often tolerated as informational rather than imperative, and explicit self-monitoring is less chronically recruited. As a result, the entry barrier to positive affect should be lower, with higher baseline probability of threshold crossing under identical sensory and task conditions.

With maturation, control increasingly reorganizes around persistent optimization:

discrepancies are rapidly evaluated, goals are made explicit, and performance monitoring becomes default rather than situational. This shift should not primarily affect the capacity to sustain positive affect once established, but rather the probability of entering it in the first place. Development is therefore predicted to compress the entry window while leaving post-threshold persistence comparatively intact. Empirically, this predicts that age-related differences will be largest for entry rate and time-to-threshold, and smaller or absent for persistence duration once lock-in occurs. It also predicts that interface manipulations that reduce evaluative affordances should partially restore access in older participants by functionally approximating earlier control regimes.

Individual differences prediction

Beyond age, stable individual differences in monitoring and evaluative control should modulate access to positive affect through the same mechanism. Individuals with higher trait self-monitoring, stronger performance orientation, or greater sensitivity to error and discrepancy are predicted to show lower entry probability under standard task conditions. These differences should be especially pronounced when evaluative cues are salient, such as visible metrics, explicit goals, or opportunities for self-assessment.

Crucially, the model predicts a dissociation between access and maintenance. Once an individual crosses the threshold into a lock-in state, persistence duration should be weakly related to trait monitoring. Instead, persistence should be governed by fatigue-related factors such as attentional depletion, arousal drift, or task endurance. High-monitoring individuals may therefore show rare but otherwise normal-length episodes when entry succeeds. This dissociation provides a clear empirical test. If individual differences primarily affect entry but not post-threshold duration, this supports the claim that optimization pressure acts as an entry barrier rather than a continuous suppressor of positive affect. Conversely, if trait monitoring shortens episodes uniformly at all time points, the threshold account would be weakened.

Interaction with task structure and interface

Developmental stage and individual monitoring traits are predicted to interact strongly with task structure. Highly structured, metric-heavy interfaces should amplify individual differences, whereas blind or semi-blind configurations should reduce variance by suppressing evaluative recruitment across participants. This predicts a convergence effect: as monitoring affordances are removed, age and traitrelated differences in entry probability should shrink.

Taken together, these predictions frame developmental change and individual variability not as differences in hedonic capacity, but as differences in default control regime. Positive affect becomes harder to access not because it is depleted, but because optimization becomes chronically active, narrowing the conditions under which suspended control can occur.

Neuro and computational extensions

This section treats neural and computational elements as implementation hypotheses, not as necessary commitments of the model. The core claims remain at the level of control topology and state transitions. Neural and computational sketches are offered to show plausibility, generate secondary predictions, and connect the framework to existing literatures without overfitting to any single mechanism (Appendix 2). This figure is not part of the formal model and is intended solely to demonstrate plausibility across levels.

Neural candidates as implementation hypotheses

The model requires a mechanism that can modulate the recruitment of evaluative control without globally suppressing perception, arousal, or task engagement. In other words, it requires a gating process that selectively attenuates the authority of evaluation and correction, rather than shutting down cognition. A plausible implementation is a control gating circuit that regulates whether discrepancy signals trigger corrective action or remain informational. Candidate components include prefrontal and cingulate systems involved in performance monitoring and arbitration, interacting with striatal and neuromodulatory systems that influence action continuation versus interruption. Within this view, evaluation is not eliminated during positive affect episodes, but its output is prevented from exerting veto power over ongoing trajectories. These are not mutually exclusive and the present paper does not require choosing among them.

The threshold extension implies a state-dependent change in gain or latency. Early in an episode, evaluative circuits can still intervene rapidly, collapsing entry. After stabilization, the same circuits may continue to register information but fail to recruit downstream control strongly enough to interrupt behavior. This predicts measurable changes in response to evaluative probes over time, such as reduced behavioral adjustment, delayed corrective responses, or altered coupling between monitoring signals and action switching.

Fatigue-based collapse further constrains neural interpretation. Termination after lock-in should correlate more strongly with markers of resource depletion, attentional drift, or arousal decay than with renewed error signaling. This dissociation distinguishes the present account from models in which control reassertion is the primary termination mechanism throughout. “Formal dynamical models are possible but not required for the present claims.”

Why these extensions matter

The value of these extensions is not explanatory completeness, but constraint sharpening. They clarify what kinds of mechanisms are compatible with the behavioral signature of fragile entry and robust persistence, and which are not. Any candidate neural or computational account must explain why evaluation is selectively disruptive early, largely ignored after stabilization, and why fatigue, rather than renewed optimization, dominates termination.

By keeping these elements explicitly optional, the paper preserves its core contribution as a control-level constraint while remaining open to multiple realizations across neural systems, tasks, and populations.

Operationalizing “optimization” and “monitoring”

Optimization and monitoring are treated as functional processes rather than subjective reports. Optimization is operationalized as the degree to which discrepancy triggers corrective action, goal re-evaluation, or strategy change. Monitoring is operationalized as the availability and use of evaluative information about performance or state.

Behavioral proxies

- Frequency and latency of error correction or strategy adjustment
- Reaction-time costs following performance feedback
- Interruptions of action flow after evaluative cues
- Menu access, score checking, or metric consultation frequency

Interface-based measures

- Presence or absence of explicit metrics (score, timers, rankings)
- Salience of feedback (visual, auditory, persistent vs. transient)
- Opportunities for self-checking without task interruption

Probe-based measures

- Brief, task-embedded evaluative prompts introduced at controlled time points

• Behavioral impact of probes (collapse, attenuation, or no effect) rather than selfreport content

Dealing with demand characteristics and introspection artifacts

Because the target phenomenon is sensitive to certain forms of assessment, the present methods prioritize minimally intrusive measurement. Measurement procedures are selected and structured to reduce disruption of ongoing task engagement.

Design principles

- Use neutral cover stories. Frame the study in terms of interface usability, task ergonomics, or interaction design rather than affect or enjoyment.
- Avoid in-episode affective queries. Do not ask participants to rate enjoyment, success, or internal state during task performance.
- Prefer delayed reports. Collect subjective reports after task completion or following a pre-specified delay.
- Use neutral experience descriptors. When experience sampling is employed, frame items around continuity of engagement, absorption, or time perception rather than affective labels.

Key constraint

Measurement is treated as a potential intervention rather than a passive readout. Experimental validity depends on minimizing the extent to which assessment procedures recruit explicit evaluation during critical periods of task engagement.

Construct validity and discriminant validity

The framework explicitly separates suspended-optimization positive affect from nearby constructs.

Distinctions

- Arousal: controlled for via identical task demands and physiological load
- Novelty seeking: tested by repetition and methodification manipulations
- Generic distraction: ruled out by preserved task performance and engagement continuity

Critical dissociation Manipulations that increase entry probability without increasing reward, arousal, or novelty support the control-based account.

Table *Construct, operational measures, confounds and mitigation*

Construct	Operational measures	Primary confounds	Mitigation strategies
Optimisation	Error correction frequency, strategy switching, reaction-time changes after feedback	Task difficulty, learning effects	Hold difficulty constant, analyze within-session changes, compare blind vs normal UI
Monitoring	Metric visibility, score checks, menu access, response to evaluative probes	Curiosity, novelty of UI	Counterbalanced UI order, habituation periods, neutral cover story
Entry	Probability of onset, latency to threshold crossing	Arousal spikes, novelty effects	Low-dose perturbations, repetition tests, delayed attribution
Persistence	Duration after threshold, uninterrupted engagement time	Motivation, reward accumulation	Fixed rewards, identical task structure, fatigue markers
Collapse	Abrupt disengagement, performance breakdown, attentional drift	External interruption, boredom	Controlled environment, distinguish fatigue vs evaluative probes

Part VII, Boundary conditions, limitations, and what the theory refuses to explain

Boundary conditions

The proposed mechanism applies only under specific structural conditions. In the absence of structured external input, suspended optimization does not yield positive affect and remains neutral. In high-optimization contexts that require continuous evaluation, explicit goals, or performance tracking, access probability is low. Affective states that scale monotonically with reward magnitude or incentive value fall outside the scope of the model. The proposed effects are therefore most observable in low-demand task regimes that permit continuous engagement without sustained planning or errordriven correction. High difficulty, learning pressure, or strategy-heavy play are expected to shift control toward optimization and reduce access even under blind interfaces.

Limitations

The framework is constrained by measurement reactivity, as probing the target state can itself alter access. Replication may be sensitive to task design and interface details, and substantial inter-individual heterogeneity is expected in baseline access without implying pathology. Entry probability may be reduced when task engagement is externally imposed or obligation-driven. Voluntary engagement, initiated under conditions of intrinsic inclination rather than demand, may better approximate the permissive control regime required for access. This constraint limits generalization to highly structured or compulsory experimental settings.

Alternative explanations and decisive tests

Competing accounts include attentional capture, simple distraction, novelty effects, placebo framing, and flow-based performance models. The decisive tests are dissociations: preserved task performance rules out distraction, repetition without monotonic gain rules out novelty, delayed attribution controls placebo, and differential sensitivity to evaluation distinguishes the model from flow.

Implications for affect science

The model explains why many “happiness pursuit” interventions backfire. Techniques that emphasize tracking, improvement, or justification increase evaluative load and therefore block entry rather than enhancing effect. Interventions succeed or fail not by hedonic strength, but by whether they preserve suspended optimization during the entry window.

Implications for experimental design and clinical assessment

When self-report, labeling, or continuous assessment recruits monitoring, the measurement alters the phenomenon itself.

In such cases, the method is not a neutral readout but a causal intervention. Valid designs must therefore minimize inepisode evaluation and rely on delayed or indirect measures.

Cultural and developmental implications

If evaluation pressure systematically increases with maturation and socialisation, access to positive affect becomes more conditional over time. This predicts population-level shifts in affective accessibility without invoking deficit or pathology, framing reduced access as a byproduct of control norms rather than loss of capacity.

Discussion

The framework does not provide a method for inducing positive affect. It specifies why method-seeking fails and delineates the minimal access conditions that remain once instrumental pursuit is suspended. A central contribution is the separation of entry mechanisms from persistence mechanisms. Many existing accounts implicitly assume that the same processes that initiate positive affect also maintain it. In contrast, the present account predicts that after stabilization, persistence becomes largely insensitive to evaluative monitoring and is instead limited by fatigue-like constraints (Matthews, Parasuraman, & Warm, 2008). This dissociation yields clear empirical predictions that distinguish the framework from reward-based, novelty-based, and flow-based explanations (Berridge & Kringelbach, 2015; Daw & O'Doherty, 2014; Csikszentmihalyi, 1990). It also resolves apparent inconsistencies in prior findings. Reports that introspection or mindfulness sometimes disrupt and sometimes preserve enjoyment can be interpreted as timing effects relative to the entry window (Kabat-Zinn, 2003; Wilson & Schooler, 1991). Accounts that preserve deliberate monitoring or instrumental emotion regulation as compatible with the state therefore fall outside the present definition (Box 1). Apparent tolerance effects in repeated inductions are likewise attributable to methodification and attribution rather than biological desensitization (Carver & Scheier, 1982; Wegner, 1994).

Treating monitoring as a dynamic control variable rather than a fixed trait integrates these observations without invoking separate mechanisms (Carver & Scheier, 1982; Barch, Braver, Botvinick, Carter, & Cohen, 2001).

The model is intentionally coarse-grained and does not specify a unique neural implementation (Carver & Scheier, 1982; Barch et al., 2001). Measurement remains constrained by reactivity (Smallwood & Schooler, 2015). The strength of the account lies in its falsifiability: if evaluative probes remain equally disruptive at all time points, or if repetition monotonically increases access, the threshold proposal is not supported.

Conclusion

This work isolates a minimal structural claim about the conditions under which certain forms of positive affect can arise and persist. Access to these episodes requires a temporary suspension of optimization, understood as continuous evaluative control. During entry, this configuration is fragile and reliably disrupted by monitoring. If permissive conditions are maintained, a threshold transition can occur: the episode becomes self-sustaining, evaluative monitoring loses leverage, and persistence becomes primarily limited by fatigue. This explains several patterns that otherwise don't make sense. Interventions aimed at pursuing happiness often backfire. Monitoring reliably collapses early episodes but has reduced impact once the transition has occurred. Repetition produces non-monotonic effects rather than cumulative strengthening. Ease, in this view, is not a reward output to be maximized but a control-state transition with asymmetric dynamics. The value of the account lies in its falsifiability. If evaluative monitoring remains equally disruptive at all time points, if repetition monotonically increases access, or if entry and persistence are governed by the same mechanisms, the threshold extension is not supported. More broadly, the framework implies that understanding affect requires attention not only to what a system values, but to when evaluative control is sufficiently relaxed for experience to continue without correction.

Conflict of Interest: The author declares no competing interests.

References

Berridge, K. C., & Kringelbach, M. L. (2013). Neuroscience of affect: Brain mechanisms of pleasure and displeasure. *Current Opinion in Neurobiology*, *23(3)*, 294–303. https://doi.org/10.1016/j.conb.2012.12.007

Berridge, K. C., & Kringelbach, M. L. (2015). Pleasure systems in the brain. *Neuron*, 86(3), 646–664. https://doi.org/10.1016/j.neuron.2015.02.018

Botvinick, M. M., Braver, T. S., Barch, D. M., Carter, C. S., & Cohen, J. D. (2001). Conflict monitoring and cognitive control. *Psychological Review, 108(3)*, 624–652. https://doi.org/10.1037/0033-295X.108.3.624

Carver, C. S., & Scheier, M. F. (1982). Control theory: A useful conceptual framework for personality-social, clinical, and health psychology. *Psychological Bulletin, 92(1)*, 111–135. https://doi.org/10.1037/0033-2909.92.1.111

Csikszentmihalyi, M. (1990). Flow: The psychology of optimal experience. Harper & Row.

Daw, N. D., & O'Doherty, J. P. (2014). Multiple systems for value learning. *Neuroscience & Biobehavioral Reviews, 37(9)*, 198–205. https://doi.org/10.1016/j.neubiorev.2013.10.004

Friston, K. (2010). The free-energy principle: A unified brain theory? *Nature Reviews Neuroscience, 11(2)*, 127–138. https://doi.org/10.1038/nrn2787

Kabat-Zinn, J. (2003). Mindfulness-based interventions in context: Past, present, and future. Clinical *Psychology: Science and Practice, 10(2),* 144–156. https://doi.org/10.1093/clipsy/bpg016

Kluger, A. N., & DeNisi, A. (1996). The effects of feedback interventions on performance: A historical review, a meta-analysis, and a preliminary feedback intervention theory. *Psychological Bulletin, 119(2)*, 254–284. https://doi.org/10.1037/0033-2909.119.2.254

Larson, R. W., & Csikszentmihalyi, M. (1983). The experience sampling method. *Journal of Youth and Adolescence, 12(3)*, 1–19. https://doi.org/10.1097/00005053-198709000-00004

Morin, F. (2025). *Transient Suppression of ACC Monitoring Enables High-Salience Positive Affect*. Preprint. Canonical version: https://florianmorin.com/papers/acc-monitoring-suppression.html

Smallwood, J., & Schooler, J. W. (2015). The science of mind wandering: Empirically navigating the stream of consciousness. *Annual Review of Psychology, 66,* 487–518. https://doi.org/10.1146/annurev-psych-010814-015331

Warm, J. S., Parasuraman, R., & Matthews, G. (2008). Vigilance requires hard mental work and is stressful. *Human Factors, 50(3)*, 433–441. https://doi.org/10.1518/001872008X312152

Wegner, D. M. (1994). Ironic processes of mental control. *Psychological Review, 101(1)*,

34–52.
https://doi.org/10.1037/0033-295X.101.1.34

Wilson, T. D., & Schooler, J. W. (1991). Thinking too much: Introspection can reduce the quality of preferences and decisions. *Journal of Personality and Social Psychology, 60(2),* 181– 192. https://doi.org/10.1037/0022-3514.60.2.181

Appendix 1

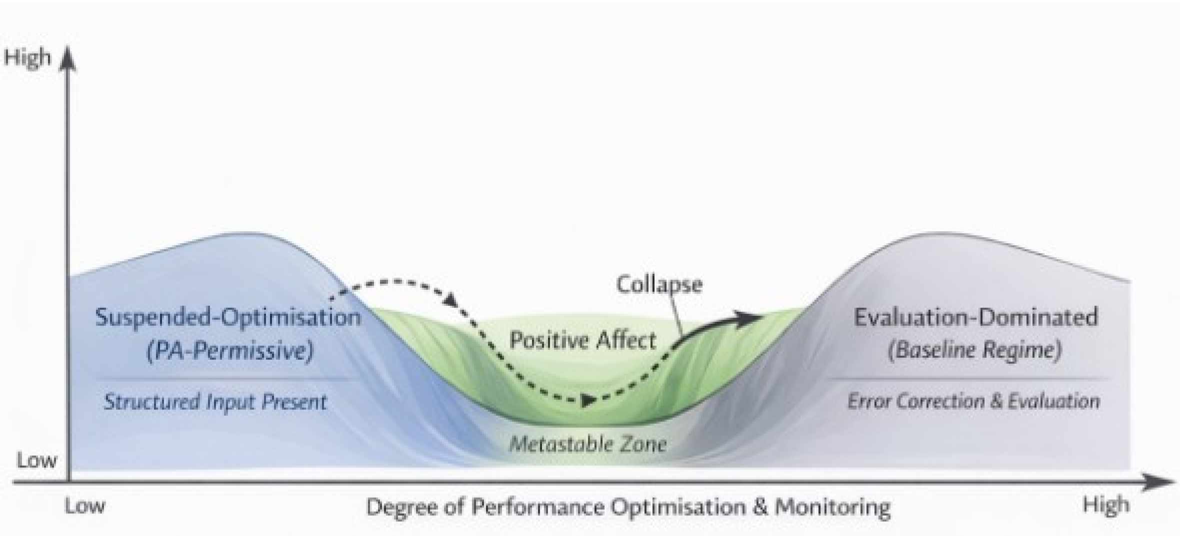

Figure 1. Illustrative control landscape (heuristic).

This schematic is a conceptual aid illustrating the distinction between optimisationdominated and permissive control regimes and the fragility of entry under evaluative monitoring. It does not depict an underlying energy landscape, continuous affective intensity, or a dynamical attractor. The "metastable zone" is a descriptive shorthand for a region of control configurations compatible with positive affect, not a latent variable or a quantitative model.

Appendix 2

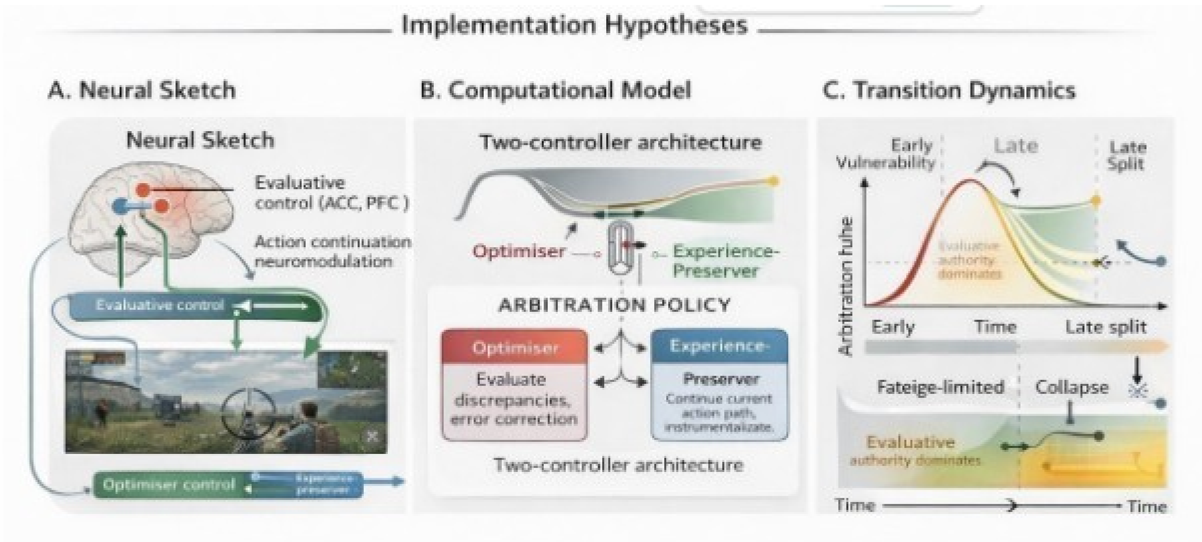

Figure 5. Neural and computational extensions.

Neural and computational sketches are presented as implementation hypotheses rather than required commitments of the framework. (A) A possible control-gating circuit that modulates the authority of evaluative control over ongoing action without suppressing engagement or perception. (B) A two-controller computational architecture in which an optimiser and an experience-preserver compete under an arbitration policy, producing fragile entry and robust post-threshold persistence. (C) Nonlinear transition dynamics predicting early vulnerability to evaluation followed by a lock-in phase in which termination is dominated by fatigue rather than renewed optimization

Appendix 3

Task-embedded induction paradigm, mid-task caffeine perturbation, and more non-instrumental perturbations

Optional: A minimal, sub-perceptual dose of caffeine is introduced during ongoing gameplay, after task engagement has already begun but before explicit evaluative reflection is likely to occur. Because the activity is already underway, the perturbation does not function as an initiation attempt and does not reframe goals, expectations, or monitoring policies. In this context, caffeine operates as a weak arousal and salience bias rather than as an attributed intervention. Its role is therefore facilitative rather than generative, potentially lowering the threshold for transition without recruiting optimization, attribution, or causal reappraisal of the episode.

Higher doses are expected to be counterproductive, as salient stimulation would increase attribution and monitoring, thereby blocking access rather than facilitating it. For this reason, individual caffeine tolerance or habitual coffee consumption is not treated as a relevant moderating factor, since the manipulation does not rely on hedonic impact or stimulant responsiveness but on a minimal, non-attributable bias applied after engagement has already begun. It follows that any observed effects cannot be attributed to caffeine as a primary driver, since the perturbation is insufficient to induce engagement or positive affect in the absence of the task-embedded structure.

- Timing sensitivity: caffeine increases entry probability when administered midtask, but has weaker or blocking effects when taken before task onset or framed as an initiation strategy.
- Attribution asymmetry: effects are strongest when caffeine is not explicitly represented as causal for the affective state.
- Dissociation: caffeine modulates entry probability and time-to-threshold more than post-lock-in persistence, which remains primarily fatigue-limited.

Compare mid-task caffeine to pre-task caffeine and placebo, holding dose and task difficulty constant, to isolate the role of timing and attribution rather than pharmacological strength alone.

Applied subset: Some other non-instrumental perturbations in interactive (ingame) environments

General principle

Introduce an action that is visibly non-optimal, unjustified, and nonrepeatable, creating a brief, 3-10 secondes, disruption of control policies without offering any exploitable strategy.

The action is deliberately non-optimal and carries no expectation of benefit.

Examples

1. Switching weapons without reason

- No tactical necessity
- No situational justification
- Performed only once
-

Rationale

This creates a transient break in action continuity without enabling learning or optimization. Because the action provides no measurable advantage, it cannot stabilize into a technique.

2. Choosing a path at the last possible moment

The decision is delayed, the choice is arbitrary, No post-hoc evaluation of the choice The choice must be late, arbitrary, and not followed by evaluation,

▪ Rationale

Late, unjustified decisions interrupt anticipatory control without introducing a new goal structure.

3. Intentionally aiming beside an opponent, This example requires strict constraints to remain admissible.

Constraints

A single occurrence only; No corrective follow-up ; No repetition even if an effect is perceived; No internal narration framed as “missing on purpose”

A single intentional deviation from task accuracy, without correction or repetition.

Rationale

This action temporarily suspends accuracy-based success criteria, weakening evaluative control without establishing an alternative objective.

4. Hesitating briefly between two neutral paths (single occurrence)

Two paths that are equivalent or neutral with respect to task success;
A short hesitation without analysis ;
No justification for the final choice;
No retrospective evaluation;

Constraint
Prolonged comparison or reasoning invalidates the perturbation. The hesitation must occur without evaluative comparison and must not delay action in a way that invites optimization, The participant must continue moving around his caractere.

Rationale,

Momentary hesitation between equivalent options introduces a transient suspension of decision closure without generating a preference structure or alternative goal. Because neither option is framed as better, the action disrupts control policies without creating exploitable information.

5 Unresolved open questioning (cognitive suspension)

Principle

Introduce a brief interrogative stance without resolution, without answerseeking, and without expectation of benefit. The operative element is not the semantic content of the question, but the momentary felt state of indeterminacy associated with “?”.

The perturbation consists in briefly instantiating the sense of a question, not in formulating or answering one. Any attempt to resolve, elaborate, or dwell on the question invalidates the perturbation.

Admissible examples :

Momentary activation of an open interrogative state, immediately abandoned:

- “Why this, right now?” “What is happening?” “What counts here?”

Descriptive, target-free questioning, without pursuit:

- “What am I noticing?”
- “What is present?”

Question followed by immediate redirection of attention, meaning without pausing the game:

Interrogative stance → interruption → primary task

Appendix 4

How to cite and implement the Ease framework (minimal standard).

1. Cite the canonical definition and scope in Box 1 as the source of the framework.
2. Refer to falsification criteria by label (F1–F5) when reporting boundary conditions.
3. Refer to predictions by label (A1–D1) and specify which subset was tested.
4. Report separately: (i) entry probability, (ii) time-to-threshold, and (iii) post-threshold persistence.
5. Treat evaluative monitoring as an access constraint (A-set), not as a continuous suppressor.
6. When probing the state, distinguish evaluative probes from equally salient non-evaluative interruptions (A2).
7. Model time-dependence explicitly by testing early vs late probes (B1) and prevention vs termination asymmetry (B2).
8. Control for repetition effects and methodification (C1–C3) by varying task structure across sessions.
9. Use “Non-Use microtasks” as the task family name; treat UT99 “Blind Play” as one implementation domain.
10. When possible, include Z proxies and test whether Z predicts entry more strongly than persistence (D1).