A Regime Theory of Joy

The Ease Regime as a Permissive Control Configuration

Florian Morin
Independent Researcher in Affective Science
florianmorin.com
ORCID: 0009-0004-7301-8524

May 2, 2026

Abstract

Joy is not well captured by accounts that treat it as a direct product of reward, learning, or goal-directed activity. Instead, it depends on access conditions and is gated by evaluative monitoring, becoming available only when such monitoring falls below a critical threshold.

This transition defines a distinct control regime, termed the ease regime. Access is governed by a state variable, Z, indexing the level of evaluative monitoring and optimization. A key manifestation of high Z is a persistent pressure to select and update the next action, constraining affective intensity through continuous correction and performance tracking.

When Z falls below threshold, this pressure is reduced. Discrepancies remain informational rather than being converted into corrective signals, allowing high-intensity positive affect to emerge. Entry is abrupt, consistent with a regime shift rather than gradual learning.

A central blocking mechanism, $Zshift$, reflects self-reinforcing monitoring triggered by attempts to produce or verify the state. The model predicts a measurement paradox: methods that recruit self-monitoring tend to suppress the phenomenon they aim to observe.

This framework provides a falsifiable alternative to reward-based accounts by formalizing joy as a threshold-dependent regime of system dynamics, governed by access conditions rather than optimization.

Key points

- High-intensity positive affect depends on access to a control regime defined by reduced evaluative coupling, rather than reward or optimization processes.
- Entry is threshold-dependent and is disrupted by monitoring, evaluation, and goal-directed control.
- Prediction error contributes directly to experiential intensity only within this regime.
- The regime enables, but does not require, high-fidelity reactivation of previously encoded experiences.
- Memory reactivation is a consequence of regime access, not its primary mechanism.
- High-intensity positive affect requires both reduced evaluative coupling and sufficient coherence of discrepancy signals to support their propagation.

1 Introduction

Contemporary models of positive affect predominantly conceptualize joy as the outcome of reward processing, goal attainment, or affective valuation (Berridge & Kringelbach, 2015; Schultz, 2016). Within these frameworks, positive affect is typically treated as a continuous and directly producible output of identifiable variables such as reinforcement, prediction error, or appraisal (Sutton & Barto, 2018; Rutledge et al., 2014; Kahneman, 1999; Diener, 2000).

However, this view struggles to account for reports of high-intensity positive states that appear abruptly, resist voluntary induction, and collapse under attempts at control or stabilization. Initial structured observational work suggests that such states may depend on a distinct configuration of cognitive-affective dynamics rather than on the accumulation of reward signals (Morin, 2026).

Notably, these states cannot be reliably produced through effort or optimization and are often disrupted by monitoring or control processes, consistent with evidence that introspection can interfere with ongoing cognition and performance (Schooler, 2002; Wilson & Schooler, 1991; Beilock, 2010).

Standard affect measurement introduces evaluative demands, requiring individuals to monitor and report their internal states. This engagement of control processes may prevent the emergence of the phenomenon itself, yielding a measurement paradox: the more precisely a high-intensity state is targeted, the less likely it is to occur. This concern is consistent with broader accounts of reactivity and observer effects in psychological measurement (Campbell, 1957; Kelley, 1967).

In response to these limitations, the present paper proposes a regime-based account of joy. Rather than treating positive affect as a scalar output of reward systems, joy is conceptualized as emerging from a specific configuration of cognitive control in which discrepancy signals remain informational and are not immediately converted into corrective demands. This configuration defines a permissive regime characterized by reduced evaluative stabilization and altered dynamics of prediction error processing (Friston, 2010; Clark, 2013).

Within this framework, entry into joy is governed by threshold dynamics rather than gradual accumulation. The transition occurs when evaluative control fails to stabilize discrepancies, allowing prediction error to function not as a corrective signal but as a driver of experiential intensity. This model generates distinct empirical predictions, including abrupt onset, all-or-none phenomenology, sensitivity to contextual evaluation load, and incompatibility with sustained goal-directed optimization.

By reframing joy as a regime shift rather than a reward output, this account aims to resolve limitations in existing models and to motivate experimental approaches that do not interfere with the conditions under which high-intensity positive states emerge.

2 Reframing "Ease" as a Computational Regime

The term "ease" is used here as a provisional label for a specific configuration of cognitive-affective dynamics. This account does not introduce a new primitive category of affect, but identifies a configuration of control parameters that gives rise to a distinct class of experiential states. This distinctiveness is emergent rather than fundamental. A regime is defined as a stable configuration of relationships between key functional components, including discrepancy detection, evaluative processes, and action selection. Regimes differ not by the presence or absence of these components, but by how they are coupled. The present account focuses on a configuration in which discrepancies are not immediately converted into corrective or goal-directed processes, reducing the pressure to select and update the next action. Under standard conditions, discrepancies between expected and observed states trigger rapid evaluation and correction, supporting efficient goal pursuit but imposing a continuous evaluative structure on experience. In contrast, under reduced evaluative coupling, discrepancies persist as informational signals rather than initiating control processes.

This shift has three main consequences. First, discrepancies are no longer rapidly resolved, altering the temporal dynamics of processing. Second, prediction error contributes directly to experiential intensity rather than functioning solely as a transient control signal. Third, the dominance of goal-directed structure is reduced, and the pressure to determine what to do next is weakened. Importantly, this regime does not reflect a deficit or pathological failure of control. The underlying architecture remains intact, and evaluative processes can be re-engaged at any moment. Rather, the regime reflects a change in control policy, specifically a reduction in the gain assigned to evaluative processes. Accordingly, "ease" does not denote a separate affective category, but a configuration within a broader space of possible control regimes. Access to the permissive regime depends on reduced evaluative coupling, while the emergence of high-intensity positive affect additionally requires that discrepancy signals remain sufficiently coherent to propagate within the system.

The present proposal isolates one such configuration without claiming phenomenological exclusivity. Its contribution lies in situating the phenomenon within established computational frameworks, including predictive processing, control theory, and reinforcement learning, while allowing it to be described in terms of parameter changes and coupling dynamics rather than as an exceptional category. The regime is not defined by specific content, including memory. Its defining property is reduced evaluative coupling. Under these conditions, two classes of phenomena may emerge. First, high-intensity positive affect becomes accessible as a regime property. Second, previously encoded experiential configurations may be reactivated with high fidelity. Such reactivation is contingent on regime access and should be interpreted as a consequence of the underlying control dynamics rather than as a primary mechanism. Z is a composite, macro-level index of evaluative load, while g denotes the local coupling parameter through which this load is expressed in the transformation of discrepancy into corrective signals, and behaviorally, in the pressure to select and update subsequent actions.

3 A Regime Space Account of Brain Function

The present account treats the “ease” regime as a specific region within a broader space of control configurations. A regime is defined by the coupling between monitoring, evaluation, and action selection. This space is parameterized by global evaluative load (Z) and coupling strength (g), which together determine whether discrepancies are corrected or allowed to persist.

Most regimes operate under high evaluative load, supporting continuous monitoring and correction. By contrast, the ease regime emerges when Z falls below a critical threshold ($Z < Z$), weakening coupling (low g) while preserving sensitivity to incoming signals. In this configuration, prediction error contributes directly to experiential intensity.

Transitions are threshold-dependent and discontinuous. This framework shifts the focus from reward or optimization to control architecture as the primary determinant of subjective experience.

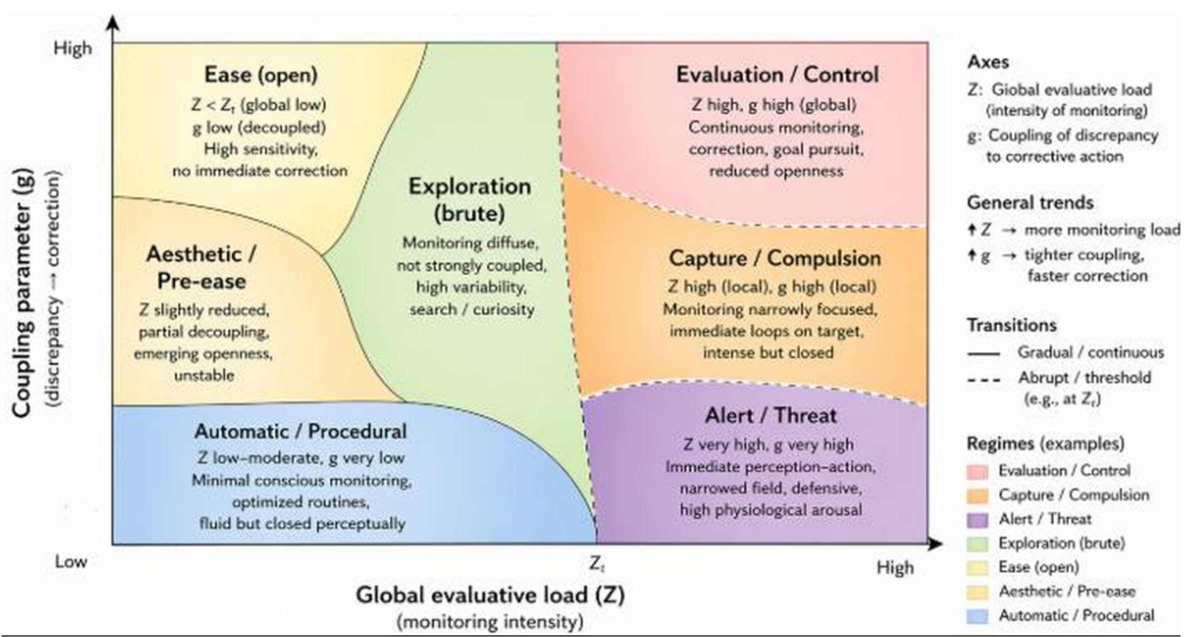

Figure 1: Brain function is organized along two dimensions: global evaluative load (Z) and coupling strength (g). Distinct regions correspond to stable regimes of monitoring, evaluation, and action selection. High Z and g yield evaluation/control, while low Z and g define the ease regime. Intermediate regions include exploration, aesthetic/pre-ease, and procedural states. The threshold Z marks a transition below which access to ease becomes possible.

The regime-space account also explains why genuinely joyful regimes may be comparatively rare. Most stable configurations of brain function are organized around control, prediction, correction, defense, efficiency, or narrative coherence. These configurations may be functional, calm, meaningful, or even rewarding, but they do not necessarily produce joy. Joy requires a more constrained combination: global evaluative load must be sufficiently low, discrepancy-to-correction coupling must be weakened, and sensitivity to incoming signals must remain high. This makes joy a narrow region of regime space rather than a default affective outcome. Most regimes are not affectively optimal. They are operationally useful. Joy appears only in a narrow region of regime space where evaluative load is low, corrective coupling is weakened, and sensitivity remains high.

The regimes illustrated here should not be interpreted as an exhaustive taxonomy, but as coarse attractors within a continuous regime space. The underlying space likely supports a large number of hybrid, transient, or unstable configurations that are not captured by discrete labels. The present framework is therefore generative rather than enumerative.

Box 1. The adult baseline problem

The present framework challenges a central assumption in well-being science: that adult reports capture the relevant range of positive experience. If high-intensity joy becomes largely inaccessible and poorly remembered after development, then adult well-being research may be calibrated on a truncated experiential baseline.

On this view, joy is not merely stronger happiness. It is a distinct experiential regime whose absence has been normalized. Developmental accounts of childhood openness and flexible inference (e.g., Gopnik, 2009; Harris, 2000) are therefore relevant because they point to a mode of cognition in which such access may be less constrained.

If access is re-established under low-evaluative conditions, for example in contexts similar to the M-ZRT, the contrast is not gradual but categorical. It forces a reassessment of what counts as intensity, value, and well-being itself.

3.1 Z definition

We distinguish three levels of description.

- Monitoring refers to a control-layer process that continuously tracks discrepancies between expected and observed states.

- Evaluation and optimization are behavioral consequences of sustained monitoring, reflecting the systematic conversion of discrepancies into corrective, goal-directed adjustments.

- Z is defined as a macro-level index of this evaluative load, integrating multiple sources of monitoring across timescales. This load is implemented locally through a coupling parameter g, which governs the transformation of discrepancy signals into corrective drives.

Accordingly, the relationship between these components can be summarized as follows:

monitoring $\rightarrow$ evaluative load (Z) $\rightarrow$ coupling strength (g) $\rightarrow$ conversion of discrepancy into correction, expressed behaviorally as pressure to select and update the next action.

3.2 Brain dynamics as regime-like organization

Recent developments in neuroscience increasingly characterize the brain not as a collection of fixed modules, but as a dynamical system capable of occupying multiple large-scale regimes. Neural activity is now commonly described in terms of network configurations, state transitions, and context-dependent coupling, rather than stable, localized functions (Friston, 2010; Clark, 2013; Shine et al., 2019). Within this view, cognitive and affective phenomena depend on the global organization of interacting processes, including prediction, control, and salience, whose influence varies across regimes. This shift toward a dynamical, regime-based perspective provides a natural framework for interpreting abrupt transitions, state-dependent processing, and variability in experiential access, and is consistent with the present proposal that certain forms of positive affect emerge from changes in control configuration rather than from incremental modulation of specific signals.

In particular, changes in large-scale coupling may be expressed behaviorally as variations in the pressure to select and update subsequent actions, providing a tractable link between system-level dynamics and observable patterns of cognition and behavior.

4 Relation to Existing Frameworks

The regime proposed here overlaps superficially with several well-established constructs in the study of attention, affect, and consciousness. A comparative analysis helps clarify these distinctions. All phenomena described in this framework are derived from a single control relationship: evaluative monitoring increases evaluative load (Z), which in turn modulates coupling strength (g), thereby determining whether discrepancy signals are rapidly corrected or allowed to persist.

4.1 Meditative absorption (e.g., jhāna states)

The regime described here does not require sustained attentional control. Entry is abrupt and cannot be achieved through deliberate stabilization.

4.2 Flow states

The present regime is not tied to performance. An extended phenomenological and functional comparison

Dimension	**Flow**	**Ease**
Core condition	Sustained evaluative control with goal alignment	Reduction of evaluative load enabling permissive control
Control logic	Continuous monitoring and optimization	Absence of monitoring and optimization without replacement
Entry condition	Gradual, built through skill–demand matching	Threshold-based, requires reduction of evaluative load below critical level
Transition profile	Continuous variation in intensity	Discontinuous, all-or-none shift
Relation to goals	Requires goal pursuit and structured tasks	Goal pursuit reinstates evaluative control and prevents access
Response to optimization	Stabilized and improved through optimization	Any attempt to optimize constitutes a structural failure condition
Temporal structure	Maintained through continuous engagement	Emerges abruptly, persists without active maintenance
Stability	Stable under sustained control	Metastable, persists without control but collapses upon reintroduction of monitoring
Failure mode	Breakdown under overload or disengagement	Collapse via reactivation of evaluative monitoring
Measurement compatibility	Compatible with self-report and performance tracking	Measurement acts as an intervention that reinstates evaluative control
Strategy extraction	Effective strategies can be identified and reused	No stable strategy can be extracted or reused
Sensory profile	Enhancement of task-relevant features	Global amplification of sensory and affective salience
Stimulus role	Dependent on task relevance	Stimuli modulate intensity only after regime shift, not sufficient for induction
Emotional tone	Engagement, absorption, satisfaction	High-intensity positive affect, often somatic and diffuse
Access mechanism	Scales with effort, skill, and feedback	Not scalable, not reproducible through effort or repetition
State modulation	Influenced by task difficulty and engagement	Constrained by evaluative load, not directly inducible

Table 1: Control-dependent versus control-incompatible regimes of positive affect. Flow is stabilized by evaluative control, whereas the reduced-control regime depends on its absence and is disrupted by its reintroduction.

is provided in Appendix A. This comparison highlights that ease and flow rely on distinct control configurations, with ease emerging under reduced evaluative monitoring and exhibiting a discontinuous, high-intensity profile.

5 Reward-based and Reinforcement Learning Accounts

Standard accounts of positive affect link pleasure to reward prediction error, reinforcement signals, or dopaminergic valuation processes. Within these models, affect is typically treated as a function of outcomes, expectations, and their discrepancies.

In the present framework, prediction error is not eliminated but persists under reduced evaluative coupling, contributing directly to experiential intensity rather than being rapidly converted into corrective signals. The key difference lies not in its presence, but in its mode of integration within the control architecture.

This distinction can be decomposed into two roles. Prediction error provides the energetic component of the experience, while reduced evaluative coupling defines the access condition. Accordingly, prediction error is necessary but not sufficient for the emergence of high-intensity positive affect, as its impact depends on the underlying control regime. Consistent with this view, dopaminergic signals are widely interpreted as encoding prediction error rather than subjective pleasure per se (Schultz, 2016, 2021), reinforcing the separation between computational learning signals and phenomenological experience.

Prediction error should therefore be understood as a carrier of experiential intensity within the permissive regime, rather than as its defining mechanism, which is determined by the control architecture mediated by evaluative load (Z) and coupling (g).

These considerations suggest that prediction error alone cannot account for high-intensity affective states, but instead undergoes a control-dependent transformation within the proposed framework.

6 From Prediction Error to Affective Experience in the Permissive Regime

These effects are modulatory rather than generative, as they operate only once the control regime has shifted. In more recent formulations, increased model precision and control have been associated with reduced exploratory dynamics, suggesting that heightened evaluative constraint may suppress modes of experience characterized by open-ended or high-variability signal propagation (Clark, 2019).

Within the permissive regime, certain classes of sensory input appear to modulate the intensity of positive affect without being sufficient to induce the state. In particular, visual features associated with high salience, such as brightness, contrast, or properties resembling rarity cues (e.g., glow-like or "treasure-like" signals), are frequently reported to amplify subjective intensity. These features may act by increasing the persistence or propagation of discrepancy signals under reduced evaluative coupling, without immediately recruiting corrective or action-selection processes.

A second, partially distinct mode of positive affect is also observed. This mode is typically associated with slow-tempo, low-frequency auditory input, often involving simple or minimally structured melodic patterns. It is described as less localized and less sharply defined than the primary mode, and may involve diffuse, low-contrast perceptual qualities.

Importantly, these inputs do not reliably produce the regime in isolation. Their effects appear contingent on prior access to the regime and should therefore be interpreted as modulatory. This distinction suggests that different input classes may interact with the regime through partially dissociable pathways, potentially corresponding to different modes of discrepancy propagation and different degrees of engagement with action-selection pressure.

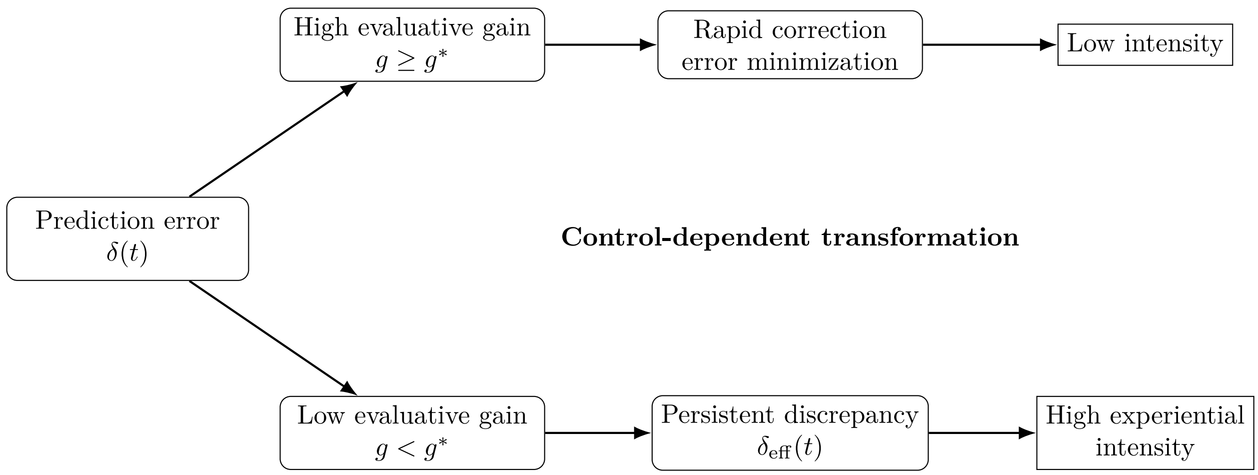

Figure 2: Control-dependent transformation of prediction error. The same discrepancy signal (δ) leads to different outcomes depending on evaluative gain (g). Under high gain, discrepancies are rapidly corrected and have limited experiential impact. Under low gain, discrepancies persist and accumulate, contributing directly to subjective intensity. This dissociation explains why prediction error alone is insufficient to generate high-intensity positive affect.

Within the present framework, the permissive regime does not originate from memory processes, but may enable the reactivation of previously encoded positive experiential configurations. For example, exposure to a previously encountered object (e.g., a childhood-related stimulus) may be accompanied by a rapid and coherent re-emergence of its associated affective tone. This reactivation appears to involve the joint reinstatement of sensory, affective, and bodily components, rather than the retrieval of abstract or narrative memory content.

Importantly, this process may occur automatically and resist suppression, consistent with a system-level change in access conditions rather than a controlled retrieval process. From a mechanistic perspective, this phenomenon is consistent with processes resembling hippocampal pattern completion, in which partial cues trigger the reactivation of previously encoded configurations.

Notably, the permissive regime does not generate new affective content, but restores access to configurations that remain encoded yet are typically inaccessible under standard control dynamics. Accordingly, reactivation inherits the same dependency on Z and g as all other regime-level phenomena.

7 All-or-None Entry and Regime Shift Dynamics

Entry into the permissive regime is characterized by abrupt, unambiguous transitions rather than gradual changes. Probabilistic access and abrupt entry refer to distinct aspects of the same process and should not be conflated. The probability of entering the regime depends on the system's proximity to a critical threshold, shaped by fluctuations in evaluative load (Z). By contrast, once this threshold is crossed, the transition itself is discontinuous and experienced as abrupt. Accordingly, access is probabilistic, while entry is all-or-none. Reports frequently describe a clear onset, with little to no intermediate states.

Formally, this transition corresponds to the crossing of a critical boundary in evaluative load (Z), such that a reduction below threshold weakens the coupling between discrepancy and correction ($g < g^*$). Behaviorally, this corresponds to a reduction in the pressure to select and update the next action. Importantly, this all-or-none characterization applies to regime entry rather than to the internal dynamics of the state. Once the regime is entered, experiential intensity may vary continuously as a function of ongoing signal dynamics.

The phenomenological profile associated with this transition is not used as direct evidence for the underlying mechanism, but as a set of constraints that competing explanations must satisfy. Alternative interpretations, such as immersive distraction or gradual mood elevation, generate distinct and testable predictions.

The transition can be described as a regime shift, marked by a sudden and sustained increase in positive affect. In some cases, the resulting intensity is reported as unusually high, occasionally approaching levels that are difficult to tolerate. Importantly, the shift is multidimensional. Alongside elevated positive affect, reports frequently include increased sensitivity to reward, enhanced emotional variability, vivid imagery, spontaneous laughter, and a general amplification of experiential salience.

The state is typically described as self-sufficient and non-teleological, emerging independently of explicit goals or outcomes. Joy may present as a chest-centered pleasure, but is more generally characterized by a global increase in experiential intensity. A notable feature of this transition is its immediate recognizability. In some cases, the initial phase may include atypical or difficult-to-classify affective qualities, reflecting the abrupt reconfiguration of control dynamics associated with reduced evaluative load (Z) and weakened coupling (g).

Taken together, these observations suggest that access to high-intensity states is not gradual, but depends on whether the system enters a distinct operational regime.

8 Joy as a Threshold-Dependent Regime Property

In the present framework, joy is not treated as a direct product of inputs, actions, or optimization processes. Instead, it is conceptualized as a regime-dependent property of the system, contingent on the level of evaluative control. This control is captured by a variable Z, which reflects the degree to which the system engages in continuous comparison, evaluation, and correction. Z indexes the strength of evaluative coupling that transforms discrepancies into goal-directed adjustments, behaviorally expressed as pressure to select and update the next action. The system is governed by a critical threshold, $Z_t hreshold$. When Z remains above this threshold, the system operates in an evaluation-dominated regime. When Z falls below this threshold, a qualitative shift occurs, and the system becomes permissive to a distinct regime in which high-intensity positive affect becomes accessible.

Formally:

- If $Z > Z_t hreshold$, the system remains in an evaluation-dominated regime
- If $Z < Z_t hreshold$, the system transitions into a permissive regime

Accordingly, the threshold applies to access conditions, not to the magnitude of the resulting affect. This transition reflects a discontinuous change in system dynamics rather than a gradual modulation. In the evaluation-dominated regime (Z high), processing is characterized by continuous monitoring, goal-directed behavior, performance tracking, and rapid error correction and a persistent pressure to determine what to do next. Experience is structured by narrative integration and outcome relevance, and positive affect tends to be limited, unstable, and dependent on external contingencies. By contrast, in the permissive regime (Z low), evaluative processes fail to stabilize discrepancies. This results in a relative suspension of optimization, reduced coupling between perception and action selection, and a weakening of next-action pressure. Within this configuration, high-intensity positive affect becomes accessible as a property of the system's dynamics rather than as the result of reward accumulation.

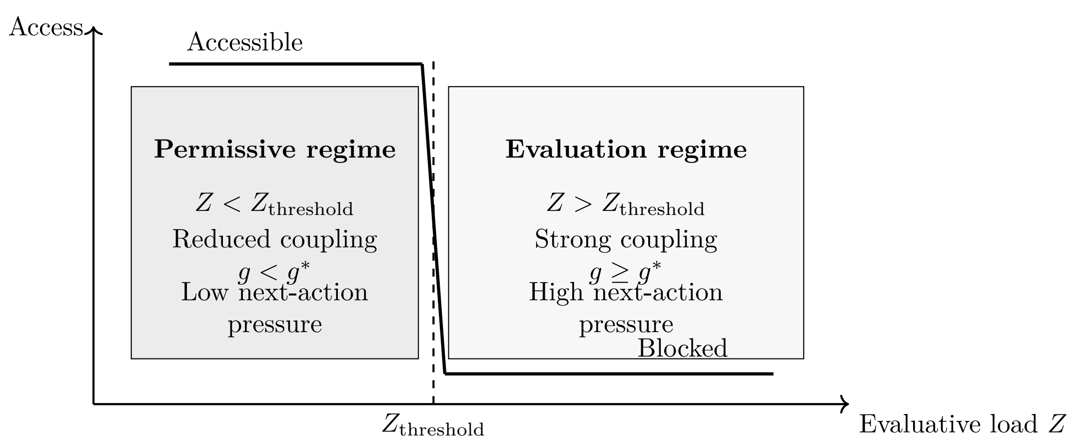

Figure 3: Threshold regime model. Access to the permissive regime depends on evaluative load (Z), with a discontinuous transition at $Z_{\text{threshold}}$. The regimes differ in coupling strength (g) and in the degree of next-action pressure.

9 Anti-Instrumental Access and Control Sensitivity

Within the proposed framework, attempts to intentionally use, optimize, or reproduce the state are predicted to increase evaluative load, thereby reducing the probability of regime access. More precisely, evaluation increases evaluative load (Z), which in turn restores the coupling between discrepancy and correction ($g \geq g^*$), and behaviorally reintroduces pressure to select and update the next action, thereby preventing or terminating regime access. Goal-directed control and monitoring processes increase Z, pushing the system above the critical threshold.

Accordingly, access to the permissive regime is disrupted by conditions that introduce explicit or implicit optimization demands, including goal setting, performance-oriented repetition, and real-time evaluation. This property distinguishes the present account from frameworks in which positive affect can be enhanced through training or reward maximization. In contrast, increasing control in this sense reduces access by increasing evaluative load (Z), strengthening coupling (g), and reinstating action-selection pressure.

Interventions should therefore target the structure of control rather than affect directly, for example by reducing monitoring or limiting narrative capture. This constraint applies not only to explicit goals but also to continuous micro-level corrections applied to ongoing experience.

Within the permissive regime, external stimuli modulate the intensity of the state without generating it. Inputs that introduce structured but manageable discrepancies, such as small and frequent prediction errors without strong demands for resolution, are predicted to amplify experiential intensity. Under reduced evaluative load (Z), weakened coupling (g) allows these discrepancy signals to persist without being immediately converted into corrective or action-selection processes.

A key implication is that the challenge in sustaining high-intensity positive affect does not lie in identifying appropriate activities, but in preventing their instrumentalization. Once an activity becomes associated with goals, expectations, or evaluation, it increases evaluative load (Z), strengthens coupling (g), and reinstates pressure to determine what to do next, thereby reducing access. Moreover, the modulatory effects of stimuli depend on prior regime entry. In the absence of access, similar inputs are not expected to produce comparable effects, reinforcing the distinction between modulatory and generative factors.

From this perspective, the regime is difficult to access not because suitable stimuli are rare, but because the conditions required for access are systematically undermined by optimization and the resulting pressure on action selection.

10 Evaluative Load as a General Access Constraint

Within the proposed framework, access to the permissive regime depends on reduced evaluative coupling at the point of entry. Any process that introduces evaluation, monitoring, or goal-directed control increases evaluative load (Z), thereby strengthening the coupling between discrepancy and correction ($g \geq g^*$), and behaviorally reintroducing pressure to select and update the next action, reducing or preventing regime access.

This constraint applies broadly across contexts. Explicit goal-setting, performance-oriented behavior, and attempts to intentionally induce or optimize the state all increase evaluative load. Likewise, measurement procedures such as real-time ratings, introspective prompts, or performance framing recruit self-monitoring processes that act on the same control pathway. In both cases, evaluation does not merely describe the state, but actively alters the underlying dynamics by increasing Z, restoring corrective coupling, and reinstating action-selection pressure.

Accordingly, access to high-intensity positive affect is anti-instrumental in its entry conditions. Attempts to use, reproduce, or stabilize the state tend to reduce its probability of occurrence by reintroducing evaluative processes. Measurement is therefore not neutral, but functions as an intervention on control dynamics, with the potential to suppress the phenomenon under observation or produce attenuated effects.

This leads to a methodological constraint. The absence or rarity of such states in experimental settings may reflect the evaluative structure of the task rather than true inaccessibility. Empirical investigation remains possible, but valid approaches must minimize interference with regime access. In particular, indirect, post-hoc, or minimally intrusive methods are more appropriate than real-time evaluation.

The model generates several methodological predictions. First, post-hoc reporting should yield higher observed incidence than real-time assessment. Second, behavioral or indirect proxies may detect regime transitions without requiring explicit monitoring. Third, experimental designs that reduce anticipatory evaluation, for example through partial blinding to task purpose, should increase access probability by lowering Z and reducing action-selection pressure. Finally, spontaneous reports under low evaluative conditions may serve as behavioral markers of regime entry.

More generally, these considerations imply that observation itself must be treated as a variable in experimental design. In this class of phenomena, the most direct measurement procedures may be the least valid, as they recruit the very processes that prevent the state from occurring. Converging evidence from studies comparing real-time experience sampling with retrospective reports suggests that measurement itself may act as an intervention on affective dynamics, altering both intensity and structure of reported states. This raises the possibility that certain high-intensity experiences are systematically under-detected under standard experimental conditions.

11 Monitoring, Evaluation, and Optimization

A central distinction in the present framework concerns the relationship between monitoring, evaluation, and optimization. Monitoring is defined as a control-layer process that continuously tracks discrepancies between expected and observed states. Evaluation is treated as an observable manifestation of this process, reflecting the presence of evaluative load (Z), rather than its underlying mechanism. Optimization refers to the behavioral regime that emerges when discrepancies are systematically converted into corrective or goal-directed adjustments. In this sense, optimization is not a voluntary strategy but a mode of operation induced by sustained monitoring through increased evaluative load.

Monitoring functions as a stabilizing layer that maintains the system in an evaluation-dominated regime. Mechanistically, sustained monitoring increases evaluative load (Z), which strengthens the coupling between discrepancy and correction (g), ensuring that discrepancies are rapidly transformed into goal-relevant signals. This reinforces its own activity, leading to a default mode in which experience is continuously interpreted, categorized, and adjusted.

This control structure is not directly accessible without altering it, as introspection and reporting recruit evaluative processes and thereby increase evaluative load (Z), modifying the underlying dynamics. The empirical target is therefore not monitoring itself, but its behavioral signatures, including abrupt regime entry, sensitivity to evaluative framing, disruption by goal-setting or repetition, and dissociation between access and skill acquisition.

The framework further predicts that monitoring is indirectly reinforced through its alignment with reward and meaning attribution. States associated with high evaluative load (Z) may be experienced as meaningful, competent, or task-relevant, stabilizing the optimization regime. Conversely, states with reduced evaluative load may be perceived as less legitimate, biasing the system toward reinstating monitoring. As a result, the optimization regime becomes self-stabilizing, both through control dynamics and through its compatibility with reward and meaning signals.

Accordingly, evaluation does not merely describe experience but acts by increasing evaluative load (Z), which strengthens coupling (g) and reshapes the regime in which experience unfolds, contributing to the systematic suppression of experiential modes that depend on reduced evaluative coupling.

11.1 Cultural Sources of Evaluative Load

The preceding analysis identifies monitoring, evaluation, and optimization as central drivers of evaluative load (Z). These processes are not only internally generated but are also reinforced by widely adopted cultural and theoretical frameworks. In many cases, such frameworks do not merely describe affective experience, but actively structure the way it is approached, interpreted, and engaged with.

The following table summarizes a set of common frameworks that may implicitly increase evaluative load by introducing conditional, comparative, or optimization-based relationships to experience. Here, the concern is not the frameworks themselves, but their simplified uptake, which often converts them into monitoring and optimization heuristics.

Table 2: Examples of cultural and theoretical frameworks that may implicitly increase evaluative load (Z) by inducing persistent forms of monitoring.

Framework	**Core assumption**	**Implicit effect on Z**	**Induced monitoring type**
Hedonic adaptation / habituation	Affective responses decrease with repetition	Installs expectation of decline, leading to anticipatory checking	Temporal monitoring: *is it fading?*
Emotion regulation models	Emotions should be managed and optimized	Introduces continuous top-down control over experience	Normative monitoring: *am I regulating correctly?*
Flow (difficulty optimization variant)	Enjoyment depends on maintaining optimal challenge	Transforms experience into a calibration task requiring continuous adjustment	Calibration monitoring: *is this difficulty optimal?*
Serotonin / lifestyle guarantee models	Positive mood is reliably produced by specific conditions such as sunlight or exercise	Creates expectation of predictable outcomes and mismatch detection	Causal monitoring: *is this working as expected?*
	Intense positive states require external substances	Externalizes access and creates dependency on triggers	Access monitoring: *do I have the required trigger?*

These frameworks converge on a common structural property: they treat affective experience as conditional, whether on outcomes, optimization, correct parameters, future value, or external triggers. This conditionalization introduces persistent forms of monitoring, including anticipatory, comparative, and reflective processes. Within the present framework, these processes correspond to increases in evaluative load (Z), which strengthen the coupling between discrepancy and correction (g), thereby reducing the probability of regime access.

Importantly, this implies that the rarity or instability of high-intensity positive states may not solely reflect intrinsic properties of the system, but also the widespread adoption of cognitive and cultural frameworks that maintain elevated evaluative load. The effectiveness of non-instrumental tasks may be modulated by pre-existing cognitive frameworks that impose evaluative or optimization-based interpretations. Even when a task is structurally designed to reduce monitoring, beliefs that frame experience as conditional, measurable, or optimizable may reintroduce evaluative load (Z), thereby reducing the probability of regime access. While the threshold model accounts for access conditions, it does not yet explain why entry remains difficult even when these conditions appear to be met.

Behavior / Strategy	Effect on Z	Induced monitoring
Going to a "favorite place" to recover a state	Increase	Spatial-outcome monitoring: *is this the right place?*
Walking with a mission (destination, goal, purpose)	Increase	Goal monitoring: *am I getting there?*
Searching for a specific atmosphere (light, mood, vibe)	Increase	Context evaluation: *is this the right atmosphere?*
Checking if the walk is "working"	Increase	State monitoring: *is something happening?*
Not optimizing the route	Decrease	Reduces efficiency constraints
Allowing arbitrary deviations (turns, detours)	Decrease	Weakens directional control
Stopping briefly for no reason	Decrease	Interrupts goal continuity without replacement
Changing sidewalk without justification	Decrease	Introduces non-instrumental variation
Turning into a street without reason	Decrease	Disrupts path predictability

Table 3: Outdoor behaviors and their implicit effect on evaluative load (Z). Behaviors that transform the walk into a search or optimization task increase monitoring, whereas non-instrumental deviations and absence of route optimization reduce evaluative coupling.

Thinking about something precise increases Z, because it turns experience into an object that can be tracked, compared, or verified. This effect becomes even stronger when the thought acts as a destination. At that point, the system is no longer receiving, it is moving toward a target: "go there," "find that," "reach this state." The thought becomes a trajectory, and trajectory reactivates monitoring. So in the model: vague or drifting thought = low to moderate Z; precise thought = higher Z; destination-based thought = very high Z.

More generally, the present account rejects the idea that joy results from the correct alignment of several external or internal variables. Rather, when the ease regime is active, many variables become amplified and retrospectively appear causal. This creates an illusion of reverse causality: the place, atmosphere, music, light, social context, or activity seems to have produced the state, when in fact these elements may have become powerful because the regime was already permissive. This illusion is difficult to abandon, because the amplification is perceptually convincing rather than merely conceptual.

12 Regime Collapse, Hysteresis, and the Z_{shift} Mechanism

Within the proposed framework, collapse of the permissive regime corresponds to the re-engagement of evaluative control, which restores the coupling between discrepancy detection and corrective processes and thereby terminates the regime. Empirically, collapse is often associated with meta-evaluative activity, such as monitoring the quality, duration, or meaning of the state. These processes are sufficient to alter control dynamics even in the absence of changes in external input or reward.

A key property of the system is its asymmetric stability. Entry is constrained by threshold conditions and appears fragile, whereas persistence can tolerate partial re-engagement of evaluative processes. This asymmetry is consistent with hysteresis-like dynamics, in which the authority of evaluative control depends on recent state history. Formally, entry requires suppression of evaluative gain below a critical threshold, while persistence allows partial increases without immediate collapse. Collapse, however, reinstates evaluative authority in a way that prevents rapid re-entry.

A central mechanism (Z_{shift}) involves anticipatory monitoring that maintains elevated evaluative load and prevents re-entry. This mechanism is detailed in the following section. This mechanism suggests a potential developmental transition in which the emergence of anticipatory testing reduces spontaneous access to the permissive regime. Rather than a gradual change, this transition may reflect a discrete shift associated with the realization that the state does not reliably return when expected. Z_{shift} is best understood as a functional construct capturing a set of convergent control-related processes, including evaluative monitoring, goal-directed control, and anticipatory regulation. Among these, evaluative monitoring appears to play a central and partially controllable role, as even minimal forms of self-observation can significantly alter the probability of regime entry.

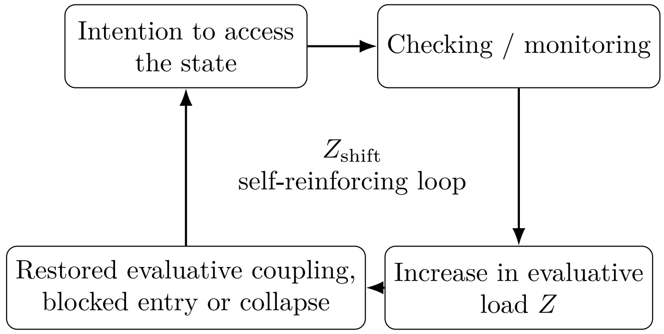

Figure 4: Schematic representation of the Z_{shift} mechanism. Attempts to access or verify the state recruit monitoring processes, which increase evaluative load (Z), restore evaluative coupling, and thereby block entry or terminate the permissive regime. This generates a self-reinforcing loop in which the intention to recover the state helps maintain the conditions that prevent its return.

Monitoring Dynamics and the Z_{shift} Mechanism

Within the proposed framework, monitoring is a control-layer process that tracks discrepancies between expected and observed states. Under standard conditions, it stabilizes behavior by converting discrepancies into evaluative signals that drive corrective adjustments. Sustained monitoring increases evaluative load (Z), strengthens the coupling between discrepancy and correction ($g \geq g^*$), and maintains an evaluation-dominated regime.

Evaluation and optimization are treated as direct expressions of this process. Discrepancies are interpreted in relation to performance or goal relevance and are rapidly transformed into corrective tendencies. This reinforces monitoring itself, producing a stable mode in which experience is continuously interpreted, compared, and adjusted.

Monitoring is therefore not neutral. Introspection, evaluation, and attempts to control experience recruit the same control dynamics, increasing Z and suppressing modes of experience that depend on reduced evaluative coupling.

The Z_{shift} mechanism describes a self-reinforcing extension of this process. Attempts to access or verify the permissive regime induce monitoring, which increases Z, restores coupling, and prevents entry. This dynamic can be summarized as: intention $\rightarrow$ checking $\rightarrow$ elevated Z $\rightarrow$ blocked access.

With repetition, this loop may become persistent. Monitoring shifts from a response to immediate discrepancies to an anticipatory process driven by expectations about internal states. As a result, evaluative load remains chronically elevated, reducing the probability of regime entry even in the absence of external demands.

This introduces an asymmetry between entry and persistence. Entry requires a sufficient reduction in evaluative load to weaken coupling ($g < g^*$), whereas persistence can tolerate partial re-engagement of monitoring. However, reactivation of the Z_{shift} loop rapidly increases Z, restores coupling, and terminates the regime, while biasing the system against re-entry.

Monitoring and Z_{shift} thus form a unified control dynamic: monitoring generates evaluative load, and Z_{shift} maintains it through anticipatory amplification. Together, they stabilize the evaluation-dominated regime and constrain access to the permissive regime.

13 Structured Prediction Error and Engagement in Cartoons

Cartoons provide a useful example of stimulus structures that may interact with the permissive regime. They are typically characterized by weak or absent resolution, minimal goal-directed structure, and low or inconsistent payoff contingencies. Rather than relying primarily on classical humor mechanisms such as setup and punchline, cartoons often exhibit dense sequences of perceptual and causal violations. Physical laws may be intermittently broken, proportions destabilized, and objects transformed without warning. These features generate sustained prediction error and perceptual instability.

Children's engagement with cartoons may therefore reflect more frequent access to such regimes. Importantly, the absence of explicit reports of strong positive affect in these contexts does not imply its absence, but is consistent with a regime in which affect is not amplified, categorized, or reported.

More generally, these observations support the claim that prediction error does not require resolution to sustain engagement. Rather, sustained engagement may depend on reduced evaluative load (Z) and the absence of enforced closure, which together allow discrepancy signals to persist without immediate correction. Their effect is contingent on reduced evaluative coupling, and is therefore not attributable to stimulus structure alone.

14 Control Reduction, Signal Coherence, and Pharmacological Dissociations

Pharmacological interventions that reduce constraint or evaluative control do not reliably produce the regime described here. Although such interventions may increase affective flexibility or reduce negative affect, they often fail to induce a coherent and stable transition into a high-intensity positive regime.

For example, ketamine, which has been associated with reduced activity in anterior cingulate and related control networks, as well as rapid antidepressant effects, does not consistently produce a sustained or structured increase in positive affect. Instead, its effects are often characterized by dissociation, perceptual instability, or diffuse affective changes.

This dissociation suggests that a reduction in evaluative control, taken in isolation, is not sufficient for regime entry. Within the present framework, this corresponds to the fact that reducing evaluative load (Z) is necessary but not sufficient, as access also depends on the integrity of discrepancy signals and their capacity to propagate coherently.

Specifically, entry requires a selective attenuation of evaluative micro-optimization, reflected in a reduction of coupling (g), within an otherwise intact perceptual and salience architecture. When control reduction is global or accompanied by signal degradation, discrepancy signals may lose coherence and fail to propagate in a structured manner.

Under such conditions, the system may exhibit reduced evaluative load (Z) and weakened coupling (g), yet fail to achieve the organized persistence of discrepancy required for high-intensity positive affect. This distinction allows the model to accommodate mechanistic evidence linking control-related structures to affective regulation, while avoiding reduction to a simple inhibition account.

In particular, the framework predicts that interventions which preserve signal coherence while reducing evaluative load (Z) and coupling (g) should be more effective in enabling regime transitions than those that broadly disrupt neural processing. More generally, this perspective suggests that the relationship between control and affect is not monotonic: reducing control can produce qualitatively different outcomes depending on whether underlying signal structure is preserved or degraded.

An additional implication concerns tolerance. Rather than reflecting only reduced reward sensitivity, tolerance effects may partly arise from anticipatory control processes, in which repeated exposure increases evaluative load (Z) through monitoring and prediction of the intervention's effects, thereby restoring coupling (g) and reducing regime accessibility.

15 The Morin Z-Reduction Task (M-ZRT): A Discriminative Protocol

The Morin Z-Reduction Task (M-ZRT) is not designed as an induction technique, but as a discriminative protocol intended to differentiate between gradual, learning-based accounts and threshold-based regime-shift models.

In learning-based frameworks, repeated exposure to a task is expected to produce progressive improvement, increased reproducibility, and performance-dependent effects. By contrast, the present model predicts a distinct empirical signature: repeated failures followed by an abrupt onset, a qualitative transition, and persistence beyond the initial context. Discontinuous success following unsuccessful attempts is therefore treated not as noise, but as evidence of threshold-governed access.

The M-ZRT is constructed to prevent the stabilization of evaluative control. Attempts to directly reduce evaluation are predicted to increase monitoring, creating a self-reinforcing loop that elevates evaluative load (Z) and prevents regime entry. Control over monitoring must therefore remain indirect. Any procedure that becomes stable, repeatable, or optimizable is expected to increase Z, strengthen coupling (g), and reduce its own effectiveness.

The task operates as a brief, non-instrumental perturbation of ongoing activity. Its objective is not to improve performance or induce a target state, but to alter the conditions under which action unfolds. In particular, the protocol targets the continuity of action rather than its outcome.

A central behavioral marker concerns the aftermath of action. Under high evaluative load, actions are typically followed by a residual trace of comparison, adjustment, or justification. This trace reflects the coupling between discrepancy detection and corrective processes. Under reduced evaluative load, actions may still generate alternatives or discrepancies, but these do not reliably trigger correction. The action proceeds without producing a strong evaluative trace. The presence or absence of such post-action traces therefore constitutes a key observable signature of the underlying control regime.

Short sequences of perturbations are used to disrupt motor and cognitive continuity without allowing the formation of structured routines. Sequences that are too brief may be insufficient to affect the control dynamics, whereas extended or repeated sequences risk becoming organized and subject to optimization. The protocol therefore operates within a narrow window: sufficient to destabilize continuity, but not long enough to be captured by evaluative control.

A key constraint of the task is the absence of explicit goals, metrics, evaluation, repetition, or expectation. These conditions are intended to minimize evaluative load (Z) and prevent the reactivation of monitoring processes, thereby weakening the coupling between discrepancy and correction (g).

As an illustrative example, participants may engage in an ongoing activity (e.g., navigating a virtual environment) while introducing brief, non-instrumental perturbations, such as generating a transient mental image without attempting to maintain it, followed by an irregular, non-rhythmic motor action. The sequence is terminated without continuation, evaluation, or attempt at reproduction.

The model further predicts specific failure modes corresponding to the reintroduction of evaluative processes. These include anticipatory commitment, performance monitoring, meaning attribution, and corrective adjustment. Although these may differ in form, they reflect a common mechanism: an increase in evaluative load (Z), restoration of coupling ($g \geq g^*$), and consequent blockage or termination of regime entry.

Within this framework, Z is treated as a latent control variable. Rather than attempting to measure Z directly, the present approach evaluates its predicted effects on system dynamics. The primary dependent phenomenon is not Z itself, but the presence or absence of discontinuous transitions in affective state, along with their persistence and generalization across contexts.

15.1 Pre-task affective neutralization

Before the task, participants are instructed that no affective outcome is required. Pleasant or joyful effects, if they occur, are irrelevant to the task. Participants should not seek, preserve, intensify, or interpret them, and should continue without adjustment. This instruction is intended to prevent affective effects from becoming implicit goals and thereby reactivating evaluative monitoring.

16 Perturbation Classes in the M-ZRT: Openness, Salience, and Suspension

The M-ZRT can be decomposed into a set of minimal perturbation classes that target evaluative control through distinct functional pathways. These perturbations are not defined by their specific content, but by their effect on goal closure, attentional allocation, and action continuity.

Openness refers to transient operations that weaken immediate goal closure without initiating alternative control processes. These perturbations introduce elements that are neither used to advance the task nor maintained over time. For instance, generating a brief mental image during ongoing activity may function as an openness perturbation, provided that it is not sustained or instrumentally used. Similarly, introducing an incomplete thought (e.g., "what if. . . ") without resolving or revisiting it disrupts closure while avoiding engagement in problem-solving or evaluation.

Salience refers to the temporary amplification of a perceptual or internal element without converting it into a goal. These perturbations increase the local prominence of a signal while preventing its integration into task-relevant structure. For example, the spontaneous fading of a mental image may act as a salience event by briefly capturing attention through minimal internal change. Likewise, selecting a single sensory input (e.g., a sound) and momentarily treating it as dominant, without attaching relevance or action to it, produces a transient shift in salience without recruiting evaluative processes.

Suspension refers to brief interruptions of ongoing optimization or strategic continuity. These perturbations disrupt the progression of goal-directed behavior without introducing a substitute objective. For example, a short, non-instrumental interruption of movement followed by an uncorrected change in direction can break motor continuity without triggering compensatory adjustment or explanation.

Although distinct, these three classes converge on a common effect: they alter the conditions under which actions unfold, reducing the likelihood that discrepancies will be converted into corrective responses. In this sense, their functional role is not to suppress activity, but to weaken the coupling between discrepancy detection and correction. A key behavioral consequence is that actions are less likely to generate a persistent evaluative trace, such as comparison, adjustment, or justification.

These perturbations may be combined into short sequences that target evaluative coupling from multiple angles. Sequences that introduce openness, followed by a transient salience shift, and terminate with suspension may be particularly effective in disrupting monitoring while avoiding the formation of structured routines. As in the broader M-ZRT framework, effectiveness depends on remaining below the threshold at which sequences become repeatable, interpretable, or optimizable.

Importantly, the target configuration is not reduced engagement or disengagement, but a regime in which salience is preserved while evaluative control is attenuated. The aim is therefore not to suppress processing, but to decouple salience from optimization.

These perturbation classes are defined functionally and can be implemented across modalities, tasks, and environments without reliance on specific content.

16.1 Limitations and Boundary Conditions of the M-ZRT Protocol

The effectiveness of the M-ZRT protocol may be constrained by pre-existing cognitive and cultural frameworks that shape how participants interpret the task. While the protocol is designed to reduce evaluative load (Z) by removing explicit goals, structure, and optimization demands, its success depends on whether participants are able to engage with the task in a genuinely non-instrumental manner.

In many cases, widely adopted beliefs about positive affect may implicitly reintroduce evaluative processes. Frameworks that treat experience as conditional on outcomes, optimal parameters, correct execution, or measurable effects can lead participants to interpret the task as a method to be performed correctly or a procedure to be evaluated. This interpretation may occur even in the absence of explicit instructions. Such processes correspond to increases in evaluative load (Z), and may trigger the reactivation of the Z_{shift} mechanism, thereby preventing access to the target regime. It may take weeks to reverse deeply ingrained evaluative and optimization-based habits.

17 Discussion

The present framework reorients the analysis of high-intensity positive affect by shifting the explanatory focus from production mechanisms to access conditions. Rather than treating such states as outcomes of reward, learning, or optimization, the account emphasizes system-level control dynamics as the primary determinant of their availability. Within this perspective, phenomena such as abrupt onset, low reproducibility, and strong context sensitivity can be understood as consequences of threshold-governed access rather than variability in underlying affective capacity.

This shift has direct implications for model evaluation. Incremental and learning-based accounts predict gradual improvement, increased stability through repetition, and dependence on skill acquisition. By contrast, a threshold-based framework predicts discontinuity, weak trainability, and the possibility of repeated failure preceding abrupt transitions. The observation of such patterns would therefore provide a critical point of discrimination between these classes of models.

The framework also introduces a structural constraint on voluntary control. The difficulty of re-accessing high-intensity positive states does not appear to arise from insufficient stimulation or diminished reward sensitivity, but from a control configuration in which attempts at access modify the very conditions required for entry. This creates a self-limiting dynamic: efforts to reproduce the state increase evaluative load (Z), restore coupling (g), and thereby reduce the probability of transition. In this sense, failure of access may reflect control interference rather than affective deficit. Related accounts of dopaminergic regulation and hedonic adaptation suggest that anticipation, evaluation, and repeated exposure can attenuate affective intensity over time, potentially through increased predictive control and expectation (e.g., Lembke, 2021).

A key implication concerns methodology. If access depends on reduced evaluative load at the point of entry, then standard experimental procedures, particularly those involving real-time reporting or introspective monitoring, may systematically interfere with the phenomenon under investigation. Measurement does not merely capture the state, but alters the control dynamics that govern its occurrence. This provides a possible explanation for inconsistencies in the literature. Accordingly, low-intrusion approaches, including post-hoc reports, behavioral indicators, and designs that minimize evaluative framing, are better suited to detecting regime transitions.

Beyond rare high-intensity states, the framework predicts a more general property of everyday experience. Rather than stabilizing within a single regime, individuals may operate within a fluctuating band of moderate evaluative load, characterized by rapid transitions between partially engaged configurations such as evaluation, attentional capture, and narrative processing. These oscillations prevent stabilization, limit the accumulation of affective intensity, and disrupt temporal continuity. Baseline experience may therefore appear neither strongly positive nor fully coherent, but shallow, fragmented, and weakly structured.

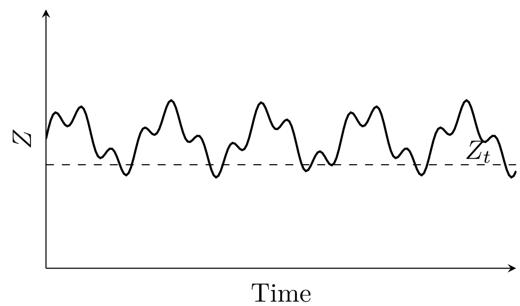

Figure 5: Everyday dynamics of evaluative load. The system fluctuates within a band of moderate Z, rarely crossing the critical threshold (Z_t) required for regime transition. These oscillations prevent stabilization in low-Z regimes and limit the build-up of high-intensity affect.

The account also offers a reinterpretation of developmental patterns. The reduced accessibility of high-intensity positive states across development may reflect changes in control structure rather than loss of underlying capacity. Increased stabilization of evaluative processes would shift baseline conditions away from those permitting regime transitions, providing a parsimonious explanation for both the intensity of early experiences and their later inaccessibility. This perspective also clarifies why memory-based approaches, such as nostalgia, often fail to recover such states despite reactivating relevant content.

At the level of stimulus interaction, the framework distinguishes between factors that modulate intensity and those that determine access. Stimuli characterized by structured variability or sustained prediction error may amplify ongoing states but are not sufficient to induce them in isolation. Their effects depend on prior access conditions, which accounts for their context-dependent efficacy.

Several limitations remain. The construct of evaluative load requires operationalization through measurable proxies, and its relationship to specific neural systems remains to be specified. While the proposed dynamics are compatible with existing accounts of salience, monitoring, and cognitive control, their precise implementation is underspecified. In addition, the phenomenological observations motivating the framework require systematic validation under controlled conditions.

Future work should prioritize experimental designs capable of detecting discontinuous transitions while minimizing evaluative interference. The proposed behavioral protocol provides one such approach, particularly as a means of distinguishing threshold dynamics from gradualist accounts. More broadly, progress in this domain may depend less on identifying optimal stimuli than on refining methods that preserve the conditions under which such states can occur. Taken together, these considerations support the view that high-intensity positive affect may be constrained less by reward availability or learning capacity than by control-related dynamics that regulate access to specific experiential regimes.

18 Implications for Development, Psychopathology, and Learning

If high-intensity positive affect depends in part on regime dynamics rather than solely on reward magnitude or valuation processes, this has broad implications across multiple domains of cognitive science.

18.1 Development

The present framework offers a reinterpretation of developmental changes in affective experience. Early childhood is often characterized by relatively frequent access to high-intensity positive states that become difficult to reproduce in adulthood. Rather than attributing this shift primarily to changes in reward sensitivity or novelty exposure, the current account locates it in the progressive stabilization of evaluative control.

Across development, particularly during adolescence, evaluative monitoring and goal-directed control become more consistently engaged. This increases the baseline coupling between discrepancy detection and correction, reducing the likelihood that discrepancies can persist without triggering adjustment. As a result, the system becomes less likely to enter regimes in which affective intensity can build.

From this perspective, the apparent "loss" of intense positive affect does not reflect a disappearance of underlying capacity, but a shift in access conditions driven by changes in control architecture.

18.2 Psychopathology

The framework also provides a control-based interpretation of conditions involving altered affective regulation. Many anxiety and mood disorders are associated with heightened monitoring, persistent error checking, and increased sensitivity to discrepancies. Within the present model, these can be understood as regimes in which the coupling between discrepancy detection and correction is excessively strong or inflexible.

Such configurations are predicted to reduce access to permissive regimes of positive affect, not because reward systems are necessarily impaired, but because evaluative processes prevent the stabilization of these regimes. Interventions that reduce evaluative gain may therefore increase access, although often in a non-specific manner that can also affect other control-dependent functions.

Importantly, this account distinguishes between reducing negative affect and enabling access to high-intensity positive regimes. These processes may overlap but are not equivalent, suggesting that standard anxiolytic or antidepressant interventions may alleviate distress without restoring access to these specific experiential states.

18.3 Reward and Reinforcement Learning

The framework challenges the assumption that increasing reward magnitude or positive prediction error is sufficient to generate high-intensity positive affect. Under standard control regimes, prediction error is rapidly converted into corrective signals, limiting its contribution to sustained subjective intensity.

The impact of prediction error therefore depends on its integration within the control architecture. Under reduced evaluative coupling, discrepancies may persist long enough to contribute directly to experiential intensity. Under strong coupling, the same signals are quickly neutralized through updating and action.

This implies that modulation of reward signals alone may have limited effects unless accompanied by changes in the coupling between discrepancy detection and correction.

18.4 Learning and Exploration

Finally, the model has implications for learning and exploratory behavior. Standard accounts emphasize optimization processes, including reward maximization and prediction error minimization. However, if certain regimes permit discrepancies to persist without immediate correction, they may support an alternative mode of exploration.

In such regimes, behavior is less constrained by immediate optimization and more responsive to local salience and unfolding dynamics. This may enable broader sampling of states and transitions, potentially facilitating forms of learning that are not directly goal-driven.

From this perspective, adaptive behavior may depend on the capacity to alternate between tightly coupled regimes, favoring efficient control and exploitation, and permissive regimes, in which evaluative coupling is reduced and exploratory dynamics can emerge. The present framework isolates one such regime and provides a basis for studying its access conditions and functional role.

19 Predictions and Falsifiability

The present framework generates a set of empirically testable predictions that distinguish it from standard reward-based and continuous models of positive affect. These predictions diverge from standard reward-based and continuous models, which would instead predict graded increases, reliability under repetition, and compatibility with measurement.

1. Discontinuous outcome distributions

Affective outcomes following repeated exposure to similar conditions should not follow a graded or normal distribution. Instead, most trials are expected to produce no notable effect, while a minority produce sharply delineated, high-intensity episodes. This predicts a highly skewed or bimodal distribution of reported experiences.

2. Sensitivity to evaluative monitoring

Introducing explicit monitoring, evaluation, or performance pressure during an ongoing episode should reliably reduce or terminate the state. Even minimal forms of introspective checking are expected to decrease both intensity and duration.

3. Low reliability of induction despite stable conditions

Repeated attempts to reproduce the state under identical external conditions should yield inconsistent outcomes. Entry probability should remain low and variable, even when contextual parameters are held constant.

4. Incompatibility with instrumental goals

Framing an activity in terms of explicit objectives, optimization, or expected outcomes should reduce the likelihood of state onset. Conversely, conditions that minimize goal-directed structure should increase access probability.

5. Dissociation from reward magnitude

Increasing reward magnitude, frequency, or salience should not reliably increase the probability of state onset. In some cases, higher reward predictability or optimization may reduce access.

6. Post-episode fragility under retrospective evaluation

Following an episode, attempts to analyze, stabilize, or reproduce the experience should lead to a rapid increase in evaluative monitoring and a corresponding decrease in subsequent access probability.

20 Minimal Formalization of Regime Dynamics

The proposed account can be formalized as a change in the coupling between discrepancy detection and evaluative correction within a standard control architecture.

Let $\delta(t)$ denote discrepancy, or prediction error, at time t. Under standard conditions, discrepancies are rapidly transformed into corrective signals through a gain-controlled process:

$$c(t) = g \cdot \delta(t)$$

where $c(t)$ is the corrective drive and g is an evaluative gain parameter. High values of g correspond to a strong and rapid conversion of discrepancies into goal-directed adjustments.

In this formulation, Z can be interpreted as a macro-level index of evaluative load, while g denotes the local coupling parameter through which this load is expressed in the transformation of discrepancy into corrective signals. Increases in Z strengthen the coupling between discrepancy and correction, with regime transitions occurring when g crosses a critical threshold.

The central proposal is that access to the permissive regime depends on a reduction, or temporary failure, of this coupling. Specifically, when the gain falls below a critical threshold g^*, the system no longer enforces immediate correction:

$$g < g^* \;\Rightarrow\; \text{decoupling regime}$$

Importantly, evaluative monitoring feeds back directly onto the gain parameter. Acts such as evaluating, measuring, or attempting to control the state increase evaluative load, which in turn modulates the coupling strength:

$$g = g_0 + \alpha \cdot Z(t)$$

where $Z(t)$ denotes evaluative load and $\alpha > 0$. Through this relationship, monitoring does not merely describe the system's state, but actively reshapes its control dynamics.

This introduces an intrinsic instability in the permissive regime. Because evaluative monitoring increases $Z(t)$, and $Z(t)$ directly increases g, any attempt to observe, stabilize, or reproduce the state feeds back into the control architecture that prevents its own existence. As g increases, discrepancy signals are more rapidly transformed into corrective drives, reducing their persistence and thereby suppressing their contribution to subjective intensity.

Formally, when $g \geq g^*$, the system re-enters the evaluation-dominated regime, in which discrepancies are immediately corrected and the permissive dynamics cannot be sustained. The regime is therefore self-terminating under monitoring: the very processes required for deliberate control or measurement are sufficient to restore coupling and collapse the state.

This feedback loop also provides a mechanistic account of the Z_{shift} phenomenon. Attempts to detect or recover the state generate anticipatory monitoring, which maintains elevated evaluative load over time. This sustained increase in $Z(t)$ stabilizes g above threshold, preventing re-entry even in the absence of external constraints.

Once the regime is entered ($g < g^*$), discrepancies are no longer eliminated but persist as informational signals. Their temporal dynamics are consequently altered. Rather than being rapidly damped, discrepancy signals accumulate or remain active over extended intervals:

$$\delta_{\text{eff}}(t) = \int_0^t \delta(\tau)\, w(\tau)\, d\tau$$

where $w(\tau)$ is a persistence weighting function reflecting reduced corrective attenuation.

Subjective intensity $I(t)$ is then hypothesized to depend on this effective discrepancy signal:

$$I(t) \propto \delta_{\text{eff}}(t)$$

Under high evaluative gain, that is, $g \geq g^*$, $\delta_{\text{eff}}(t)$ remains low due to rapid correction, resulting in limited experiential impact. By contrast, under low gain, that is, $g < g^*$, persistent discrepancies contribute directly to subjective intensity.

This feedback structure captures three core properties of the proposed regime. First, threshold dynamics: small variations in evaluative gain around g^* produce discontinuous changes in system behavior. Second, persistence of discrepancy: prediction error is no longer transient but becomes a sustained signal contributing directly to experience. Third, a self-terminating structure: the very processes that would enable intentional control also reintroduce coupling, rendering the regime inherently unstable under monitoring.

This minimal formulation does not specify a neural implementation, but it is consistent with architectures in which discrepancy signals, such as prediction error, and control signals, such as evaluative or salience-related processes, are separable yet dynamically coupled. In this sense, the equation should be interpreted as defining the geometry of the regime space: variations in evaluative load (Z) shift the system across regions characterized by different coupling regimes, rather than specifying a standalone functional relationship.

21 Conclusion

The present paper proposes that high-intensity positive affect is not adequately captured by models that treat it as a direct outcome of reward, learning, or goal-directed optimization. Instead, joy is conceptualized as a regime-dependent property of the system, emerging when evaluative coupling falls below a critical threshold and discrepancies remain informational rather than being rapidly converted into corrective demands.

This framework accounts for a set of observations that are difficult to reconcile with existing approaches, including abrupt and all-or-none transitions, sensitivity to evaluative context, limited trainability, and the paradoxical effects of measurement. It also provides a unifying interpretation of phenomena spanning development, memory reactivation, perceptual engagement, and pharmacological dissociations by linking them to a common control parameter governing the relationship between discrepancy detection and evaluative correction.

By introducing a regime-based perspective, the model shifts the focus from the production of positive affect to the conditions that enable its accessibility. In this view, high-intensity positive states depend on a regime transition, whereas many constrained positive states may remain below this threshold without undergoing such a transition. The central question is therefore not how to increase reward or optimize behavior, but how to modulate the structure of control such that discrepancies can persist and propagate within the system.

Importantly, the framework is empirically vulnerable. It generates specific predictions that distinguish it from gradualist and reward-based accounts, particularly regarding discontinuity of onset, disruption by evaluation, and the absence of reliable training effects. These predictions provide a basis for direct experimental tests, as well as for methodological revisions that minimize interference with the phenomena under study.

More broadly, the proposal highlights a potential blind spot in affect research: phenomena that depend on reduced evaluative control may be systematically underrepresented because of the methods used to investigate them. Addressing this limitation may require not only new models, but also experimental strategies that treat observation itself as a variable rather than a neutral process.

Taken together, these considerations suggest that joy, in at least one of its forms, is not best understood as an output to be maximized, but as a regime to which access must be permitted.

Conflict of Interest

The author declares no conflicts of interest.

22 references

Barto, A. G., & Sutton, R. S. (2018). Reinforcement learning: An introduction (2nd ed.). MIT Press.

Beeman, M., & Kounios, J. (2009). The Aha! moment: The cognitive neuroscience of insight. Current Directions in Psychological Science, 18(4), 210–216.

Beilock, S. L. (2010). Choke: What the secrets of the brain reveal about getting it right when you have to. Free Press.

Berridge, K. C. (2003). Pleasures of the brain. Brain and Cognition, 52(1), 106–128.

Berridge, K. C., & Kringelbach, M. L. (2015). Pleasure systems in the brain. Neuron, 86(3), 646–664.

Berridge, K. C. (2018). Evolving concepts of emotion and motivation. Frontiers in Psychology, 9, 1647.

Campbell, D. T. (1957). Factors relevant to the validity of experiments in social settings. Psychological Bulletin, 54(4), 297–312.

Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. Behavioral and Brain Sciences, 36(3), 181–204.

Clark, A. (2019). Surfing uncertainty: Prediction, action, and the embodied mind. Oxford University Press.

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

Diener, E. (2000). Subjective well-being: The science of happiness and a proposal for a national index. American Psychologist, 55(1), 34–43.

Epley, N., & Whitchurch, E. (2008). Mirror, mirror on the wall: Enhancement in self-recognition. Personality and Social Psychology Bulletin, 34(9), 1159–1170.

Friston, K. (2010). The free-energy principle: A unified brain theory? Nature Reviews Neuroscience, 11(2), 127–138.

Gopnik, A. (2009). The philosophical baby: What children's minds tell us about truth, love, and the meaning of life. Farrar, Straus and Giroux.

Harris, P. L. (2000). The work of the imagination. Blackwell.

Kahneman, D. (1999). Objective happiness. In D. Kahneman, E. Diener, & N. Schwarz (Eds.), Well-being: The foundations of hedonic psychology (pp. 3–25). Russell Sage Foundation.

Keltner, D., & Haidt, J. (2003). Approaching awe, a moral, spiritual, and aesthetic emotion. Cognition and Emotion, 17(2), 297–314.

Kelley, H. H. (1967). Attribution theory in social psychology. In D. Levine (Ed.), Nebraska symposium on motivation (Vol. 15, pp. 192–238). University of Nebraska Press.

Lembke, A. (2021). Dopamine nation: Finding balance in the age of indulgence. Dutton.

Maslow, A. H. (1964). Religions, values, and peak-experiences. Ohio State University Press.

Morin, F. (2026). Joy is fragile under evaluation: A threshold model of entry. Figshare.

https://doi.org/10.6084/m9.figshare.31385326.v3

Rutledge, R. B., Skandali, N., Dayan, P., & Dolan, R. J. (2014). A computational and neural model of momentary subjective well-being. Proceedings of the National Academy of Sciences, 111(33), 12252–12257.

Scherer, K. R. (2009). The dynamic architecture of emotion: Evidence for the component process model. Cognition and Emotion, 23(7), 1307–1351.

Schultz, W. (2016). Dopamine reward prediction error coding. Dialogues in Clinical Neuroscience, 18(1), 23–32.

Schultz, W. (2021). Dopamine reward prediction error coding: A two-component response. Nature Reviews Neuroscience, 22, 483–495.

Schooler, J. W. (2002). Re-representing consciousness: Dissociations between experience and meta-consciousness. Trends in Cognitive Sciences, 6(8), 339–344.

Schooler, J. W. (2011). Introspecting in the spirit of William James: Comment on Fox et al. Psychological Bulletin, 137(2), 345–350.

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.

23 Appendix A

Dimension	**Flow**	**Ease**
Core condition	Balance between skill and challenge	Low evaluative load, permissive control configuration
Entry dynamics	Gradual, continuous engagement through task alignment	Discrete unlocking event, abrupt transition
Unlocking event	No identifiable onset	Yes, identifiable transition point
Phenomenology (onset)	Progressive immersion	Transient high-intensity episode (reported up to 48 h) with intensified visual imagery, spontaneous uncontrollable laughter, and an uncanny experiential quality
Relation to goals	Goal-directed, performance-linked	Anti-instrumental, disrupted by goal pursuit
Sensory profile	Task-relevant enhancement	Amplified vision and olfaction, global increase in vividness
Stimulus sensitivity	Dependent on task relevance	Low-intensity stimuli become affectively charged
Memory dynamics	Task-related encoding	Spontaneous reactivation of early positive memories (pattern-completion-like reactivation of early positive memories)
Temporal dynamics	Sustained during task engagement	Can spontaneously reactivate upon waking
Emotional tone	Engagement, satisfaction	High-intensity joy, being moved, strong somatic component (often chest-centered)
Transition profile	Continuous adjustment	Discontinuous, threshold-like transition
Stability	Stable under maintained conditions	Metastable, threshold-dependent
Control architecture	Active regulation and feedback alignment	Reduced regulation, permissive dynamics
Measurement compatibility	High (self-report, performance)	Low (measurement disrupts the state)
Failure signature	Performance drops, disengagement	Repeated clean failures, then abrupt onset
State modulation	Fatigue or overload may reduce the state	Can increase under certain physiological states (e.g., mild illness)

Table 4: Structural comparison between flow and ease. Flow is characterized by task-aligned regulation, continuous adjustment, and performance-linked engagement. Ease is characterized by reduced evaluative coupling, threshold-like access, anti-instrumentality, and low compatibility with direct measurement.