

Empowerment-Driven Learning: An Evolutionary and Computational Framework for Academic Motivation

William Wai Sum Chan*

Director, Empowerment Learning Dynamics

Houston, TX

[ORCID: 0009-0006-4577-6525](https://orcid.org/0009-0006-4577-6525)

Preprint Version 1.0 — April 2026

*Author to whom correspondence should be addressed: william.chan@empowermentlearning.org

Abstract

High-achieving students often undergo abrupt, total academic disengagement once terminal credentials (e.g., college acceptance) are secured. This phenomenon—Rational Detachment—reveals a structural gap in the motivation literature: proximate frameworks like Self-Determination Theory describe the psychological conditions under which learning thrives, but none specify the biological mechanism that evaluates whether the physiological cost of cognitive effort is warranted at all. We propose Empowerment-Driven Learning Theory (EDLT) to fill this gap, drawing on evolutionary psychology, Life History Theory, and the computational logic of Active Inference. EDLT models the learner as a metabolic economy: biologically secondary knowledge receives authorization only when its expected return in survival security or social status credibly exceeds the opportunity cost of updating internal predictive models. The framework reconceptualizes motivation not as a psychological state but as a kinematic quantity—the rate of change of expected utility over time. Rational Detachment emerges from this account as an ecologically rational phase transition rather than a character deficit: when the expected utility gradient reaches zero, the system shifts lawfully from Growth into Conservation. Four falsifiable boundary conditions are derived, and a coordinated program of measurement—ecological momentary assessment, life history indexing, and a novel Ancestral Relevance Scale—is proposed to test the framework’s core predictions.

Keywords: active inference, evolutionary educational psychology, life history theory, motivation, dynamical systems, rational detachment

1 Introduction

1.1 The Anomaly: Rational Detachment

Student engagement does not decline gradually across secondary school. Gallup Student Poll data show it plummeting from roughly 74% in fifth grade to approximately one-third by late high school (Gallup, Inc., 2017), a pattern that large national surveys consistently replicate across the transition to late adolescence (National Center on Education and the Economy, 2022; Xu et al., 2025). This collapse persists across culturally diverse contexts and has not yielded to post-pandemic intervention, which suggests the mechanism is structural—embedded in the architecture of how adolescent cognition allocates resources—rather than a local artifact of any particular curriculum or institutional context.

The pattern has a counterpart in early development. In language acquisition, toddlers reliably prioritize vocabulary that maps onto high-stakes evolutionary domains—social status, immediate threat, functional action—over abstract syntactic structures that carry no ancestral utility. The biological system does not acquire data by frequency of exposure alone; it triages by survival value. That same logic, this paper argues, operates in secondary school classrooms. It generates the heuristic core of Empowerment-Driven Learning Theory (EDLT):

$$\text{Motivation to Engage} = \frac{\text{Perceived Survival Utility} + \text{Perceived Status Utility}}{\text{Metabolic Cost}} \quad (1)$$

This ratio is offered here as a conceptual heuristic, not a finalized formal model. Its structure, however, is not arbitrary: it mirrors the logic by which the brain’s predictive-processing engine discounts utility against expected metabolic and opportunity costs (Kurzban et al., 2013), a relationship that will be formalized in subsequent parameterization work. The ratio captures one prediction with particular clarity: an extreme spike in the denominator—

metabolic cost—can suppress engagement even when the numerator remains objectively high. This multiplicative suppression is what distinguishes EDLT’s account from purely expectancy-based or value-based theories.

Consider the following case, which proximate frameworks do not predict and cannot easily retrofit to explain post hoc. A high-achieving student—with a documented record of academic compliance, high output, and strong non-cognitive skills—receives admission to an elite university. Within days, effort drops to near zero. Homework goes undone. Projects sit unfinished. Crucially, the student’s cognitive capacity is fully intact: they simultaneously engage in complex social coordination or competitive online gaming at precisely the level of cognitive demand that studying had previously required.

What shifts is not capacity. It is the return-on-investment calculation. The acceptance letter eliminates the institutional threat that had previously authorized costly academic effort. The utility numerator, which was never intrinsically attached to the academic content itself, collapses once its extrinsic anchor is removed. The student triaging motivation away from physics and toward social strategy is executing the same adaptive algorithm as the toddler who masters social vocabulary before abstract syntax.

The Theoretical Puzzle

Rational Detachment, as this paper terms this phenomenon, is not burnout. Burnout involves gradual, global depletion arising from chronic demand exceeding recovery capacity—emotional exhaustion that extends well beyond the depleting domain and does not resolve when situational contingencies change. Rational Detachment is also distinct from chronic amotivation, which describes students for whom academic effort never acquired utility in the first place.

Rational Detachment is a more precise construct: the sudden, discrete reallocation of cognitive resources away from previously utility-bearing academic goals when the ecological conditions sustaining their perceived value are removed. What defines it is not the magnitude

of disengagement but its abruptness. This is not erosion to a motivational baseline. It is a phase transition.

The puzzle is that the disengaging student, by every standard account, should remain engaged. They frequently exhibit grit—the sustained perseverance that strongly predicts long-term achievement (Duckworth et al., 2007). They demonstrate growth mindset beliefs (Dweck, 2006)—beliefs that, despite yielding a modest pooled effect of $d = 0.10$ at the population level (Sisk et al., 2018), they have consistently enacted across years of high performance. Their classroom environment may provide strong autonomy support. Every construct that the motivation literature identifies as protective is in place. None of them generate the prediction that this student will collapse when the institutional contingency is removed, because none of them contain a mechanism for what we call the metabolic veto: a computation that asks whether the expected return on a specific task, at this specific moment, justifies the biological cost of continued investment.

1.2 The Explanatory Void in Motivation Research

The Proximate Bias

The mainstream motivation literature has built a sophisticated and empirically productive body of theory. Self-Determination Theory (SDT) identifies autonomy, competence, and relatedness as universal psychological needs whose satisfaction predicts intrinsic motivation and well-being (Ryan & Deci, 2000). Expectancy-Value Theory specifies how perceived ability and subjective task value—including cost—jointly shape effort decisions (Wigfield & Eccles, 2000). Control-Value Theory maps how academic emotions arise from students' control and value appraisals, with recent extensions impressively incorporating physiological arousal (Pekrun, 2006). Interest development models trace the path from situational triggers to stable individual interest (Hidi & Renninger, 2006).

However, while these frameworks share a sophisticated explanatory architecture: motivation is modeled as a function of beliefs, values, and appraisals. This is powerful within

a certain bandwidth. It cannot, however, ask why this particular engine evolved, what it fundamentally runs on, or under what conditions it determines that further combustion is not worth the fuel. The field’s explanatory infrastructure is, in Tinbergen’s terms, almost entirely proximate—focused on mechanism and ontogeny while leaving adaptation and phylogeny largely unexamined (Tinbergen, 1963). The present paper takes the adaptive and phylogenetic questions seriously and asks what they reveal about conditions where the proximate machinery fails.

This is not a novel observation at the level of principle. Evolutionary approaches to educational psychology have argued for this distinction before (Cosmides & Tooby, 1997; Geary, 2008), and important parallel work in cognitive science—Spelke’s core knowledge systems, Dehaene’s neuronal recycling model—has documented the evolved domain-specificity that mainstream motivation research typically brackets away. The gap EDLT addresses is not that evolutionary thinking is absent from educational psychology entirely, but that it has not been formally integrated with the computational frameworks (Active Inference, dynamical systems) that could give it predictive teeth in the classroom.

Resolving the Anomalies

The computational resources for this integration already exist. The free-energy principle provides a formal account of how agents allocate limited inferential resources (Friston, 2010). Opportunity-cost models from neuroeconomics establish that the subjective feeling of effort is a computed signal that current resources could yield better returns elsewhere (Kurzban et al., 2013). What is missing is a bridge between these computational frameworks and the evolutionary currencies that the adolescent brain actually tracks.

Table 1 illustrates the practical stakes of this gap by juxtaposing how standard proximate frameworks interpret several persistent classroom anomalies against how EDLT’s metabolic framework reframes them. The EDLT column is explicitly theoretical—these are not established findings but competing explanatory hypotheses whose empirical adjudication is the

program this paper initiates.

1.3 Purpose and Scope

EDLT is designed ultimately for formal parameterization as a nonlinear dynamical system, with motivation reconceived as a biological velocity rather than a static trait. This manuscript is the necessary precursor to that formalization: it establishes the evolutionary boundary conditions and metabolic logic of the system, providing the structural scaffolding that formal computational modeling requires before its parameters can be meaningfully constrained.

2 Theoretical Background

The explanatory gap identified above has two sides. On the computational side, we need a mechanism that specifies how biological systems allocate costly cognitive resources. On the evolutionary side, we need a specification of which goals the adolescent brain evolved to pursue—the currencies that give the numerator of the motivational economy its content. This section addresses both, drawing on the proximate–ultimate distinction, the limits of SDT as an exhaustive motivational account, and the evolutionary mismatch literature that specifies why formal schooling is metabolically expensive by design.

2.1 The Proximate–Ultimate Distinction

Tinbergen (1963) established that a complete account of any biological behavior requires answers to four complementary questions: what mechanism produces it, how it develops across ontogeny, what adaptive problem it solves, and what evolutionary history it reflects. The first two questions yield proximate explanations; the latter two yield ultimate ones. The motivation literature has built an impressive proximate infrastructure. It has not systematically pursued ultimate explanations, and the cost of that asymmetry is an inability to

predict conditions under which the proximate machinery will fail (Cosmides & Tooby, 1997; Geary, 2008).

This is not merely a theoretical preference. Cognitive mechanisms are biological adaptations constrained by energetic demands—a claim that requires defense rather than mere assertion. The case for it rests on three converging lines of evidence: the substantial caloric cost of prefrontal cortex operation relative to subcortical systems (Kurzban et al., 2013), the well-documented conservation of neural energy across mammalian taxa (Li & van Rossum, 2020), and the fact that human infants display species-typical learning efficiencies for evolutionarily ancient domains (social cognition, spatial navigation, spoken language) that are not replicated for evolutionarily recent ones (reading, formal mathematics) (Geary, 2008). Together, these observations support the theoretical commitment that academic motivation cannot be fully understood without specifying both the adaptive function the underlying machinery serves and the energetic constraints within which it operates.

A full account must bridge the proximate and ultimate levels. Sections 2 and 3 address the adaptive and phylogenetic questions; the proximate frameworks are then contextualized as partial maps of a more fundamental biological system.

2.2 Self-Determination Theory: Contribution and Limit

Self-Determination Theory (SDT; Deci & Ryan, 2000; Ryan & Deci, 2017) remains the most rigorously validated proximate model of academic motivation. Its identification of autonomy, competence, and relatedness as universal basic psychological needs—whose satisfaction predicts deeper engagement, internalization, and well-being—has accumulated extensive empirical support across cultures and institutional contexts. More recent elaborations have extended SDT to address need-thwarting, contextual stress, and the conditions under which the organismic drive toward growth fails to materialize (Ryan & Deci, 2017). This is not the naive version of SDT that assumes unconditional human growth-orientation; the theory has developed considerably since Deci and Ryan’s foundational formulations.

What SDT does not provide—even in its updated form—is a mechanism that specifies when the biological system will refuse to pursue the satisfaction of its own psychological needs. The organismic integration model assumes that growth, given adequate nutriment, is the default trajectory. That assumption cannot generate the prediction that a student with fully satisfied autonomy, competence, and relatedness needs will undergo abrupt disengagement the moment an external contingency is removed. It cannot do so because it has no mechanism for the metabolic veto: a computation that evaluates whether the expected return on engagement justifies its biological cost at a specific ecological moment.

EDLT attempts to supply that mechanism by providing an evolutionary grounding for SDT's three needs—specifying not just that they motivate, but why they evolved, what ecological functions they serve, and under what precise metabolic conditions they authorize versus veto engagement.

Autonomy, understood evolutionarily, is not simply a desire for choice. It reflects the capacity for rule-negotiation that was fitness-relevant in ancestral environments where social contracts could be modified through demonstrated competence. The brain's sensitivity to fixed, top-down environments is not a preference but a functional inference: when environmental parameters are perceived as immutable, the expected return on effortful behavior drops toward zero, and the system rationally down-regulates anticipated utility. This is why high levels of external control systematically reduce engagement not merely by violating a psychological need but by triggering the inference that effort cannot alter one's trajectory—the same inference that produces learned helplessness, though through a metabolic rather than expectancy mechanism.

Competence, recast through an evolutionary lens, functions as a fitness signal within a localized social ecology. Academic mastery generates engagement to the extent that it generates legible, status-relevant social capital. When mastery is invisible to peers or irrelevant to salient hierarchies, the competence need may register as satisfied (in SDT terms) while the evolutionary currency it would normally generate is blocked. The behavioral consequence—

diminished engagement despite apparent need satisfaction—is a prediction SDT cannot easily generate but EDLT derives directly from the metabolic logic.

Relatedness is grounded in ancestral coalition-building: affiliative bonds that increased survival odds against external threats. The motivational implications depend critically on who, in the student’s ecological model, occupies the role of threat versus ally. When the teacher is perceived as the primary institutional threat, coalition with peers carries metabolic risk—it consumes resources that the system is allocating to threat-management. The empirical correlate is the well-documented pattern in which high-threat classroom environments simultaneously suppress peer academic collaboration and increase underground social coordination.

Taken together, this reframing does not challenge SDT’s empirical record. It provides a biological account of the boundary conditions under which each need successfully converts into engagement—and predicts the thresholds where it will not.

2.3 Control-Value Theory: The Nearest Precedent

Of the proximate frameworks surveyed in Section 1, Control-Value Theory (CVT; Pekrun, 2006) comes closest to the explanatory territory EDLT occupies and therefore deserves specific engagement rather than a survey mention. CVT models academic emotions as arising from students’ appraisals of control over learning outcomes and the value they attach to those outcomes. Its recent extensions have incorporated physiological arousal into the achievement emotion framework, acknowledged the role of environmental affordances in shaping appraisal processes, and treated attainment costs as a component of the value structure that can suppress engagement even when positive appraisals are present (Pekrun, 2006). In this respect, CVT has moved meaningfully toward the cost-sensitive account that EDLT proposes.

The point of divergence between CVT and EDLT is not that CVT ignores cost—it does not—but that it treats cost as a consciously appraised component of subjective value, amenable to the same reflective evaluation as utility. EDLT’s metabolic denominator oper-

ates at a different level: it specifies a biological mechanism that influences engagement prior to and independently of conscious appraisal. Under high threat or steep temporal discounting, the metabolic veto fires before the student has formed a reflective cost assessment. This generates a prediction CVT cannot produce: that cost will suppress engagement even when the student’s explicit appraisal of cost is low, because the biological cost computation operates on a faster timescale than the phenomenological one. The empirical test distinguishing CVT from EDLT is therefore not a comparison of mean-level engagement predictions but a within-person design that separates the timing of biological state change from the timing of conscious cost appraisal—a question that EMA methodology, combined with biometric markers, is equipped to address.

2.4 Evolutionary Educational Psychology: The Mismatch

If proximate theories describe the phenomenological experience of motivation, evolutionary educational psychology identifies why schooling is metabolically expensive in the first place—the structural mismatch that defines the baseline parameters of EDLT’s Motivational Economy.

Primary vs. Secondary Knowledge

Geary (2007) draws a consequential distinction between biologically primary and biologically secondary knowledge. Primary knowledge encompasses domains for which the human brain has evolved dedicated neurocognitive machinery over evolutionary time—social navigation, coalition tracking, spoken language, basic numerical intuition (Geary, 2008). Children acquire primary knowledge through play and observation, with the kind of effortless efficiency that reflects metabolically subsidized evolved architecture. Secondary knowledge consists of culturally constructed, evolutionarily novel abstractions: formal algebra, written language, propositional logic, theoretical science. No dedicated machinery exists for these domains. Acquiring them demands sustained prefrontal engagement and metabolically costly synaptic

reorganization.

A single classroom observation captures the asymmetry cleanly. An adolescent can effortlessly track and update a complex, shifting social hierarchy across hundreds of peers—a cognitive feat that dwarfs, in objective information-processing terms, the challenge of retaining a three-step algebraic procedure. The algebra is retained with difficulty not because it is objectively more complex but because it recruits costly general-purpose machinery rather than any evolved specialist system. In Cognitive Load Theory terms (Sweller, 1988), secondary knowledge carries a structurally high intrinsic load.

The mismatch has a direct consequence for the Motivational Economy. Because secondary knowledge is not tagged with ancestral survival or status utility, the brain’s predictive systems will not authorize the metabolic cost of acquiring it unless that cost can be offset by primary evolutionary currencies—social status gains, threat avoidance, or coalition advantage. Secondary knowledge, therefore, sets the baseline for the denominator of the motivational heuristic; the numerator must be actively populated through pedagogical design.

One refinement to this account is important for EDLT’s instructional implications. Geary’s framework specifies not merely a binary between primary and secondary knowledge but an internal architecture within the primary domain, organized around three folk knowledge systems: folk psychology (social cognition, theory of mind, coalition tracking), folk biology (organism classification, ecological navigation), and folk physics (intuitive mechanics, spatial reasoning) (Geary, 2008). These folk systems are not motivationally equivalent. Folk psychology is the most powerful motivational substrate available to the classroom, because it is the domain most directly coupled to the survival and status currencies that EDLT identifies as the operative numerator. Folk physics and folk biology carry substantial motivational potential but require more deliberate coupling to social or threat-relevant outcomes. This architecture has a practical implication for Class I interventions (Primary Knowledge Anchoring, Section 6): the metabolic subsidy achieved by anchoring secondary content to

primary systems will vary systematically depending on which folk domain is recruited. Anchoring a physics concept to folk psychology—framing thermodynamic equilibrium as a social negotiation problem, for instance—should produce a deeper metabolic discount than anchoring it to folk physics alone, precisely because the social domain carries stronger evolutionary motivational weight. EDLT’s intervention predictions are therefore not uniform across primary-to-secondary mappings; they are sensitive to the folk domain hierarchy, and empirical tests of Class I should include folk domain as a structural variable.

Life History Theory: Contextual Calibration

Life History Theory provides the second evolutionary parameter: how far into the future a learner’s predictive engine extends when evaluating the return on costly investment (Ellis et al., 2009). Life history strategies are not genetic types but environmentally calibrated orientations. In stable, resource-predictable environments, organisms develop slow strategies characterized by delayed gratification, substantial investment in long-term skill acquisition, and willingness to defer immediate rewards for distal payoffs (Figueredo et al., 2006). In harsh or unpredictable environments, the adaptive response is a fast strategy: steep discounting of future returns, prioritization of immediate survival and status gains, and strong resistance to deferring present utility for uncertain future reward (Del Giudice, 2020).

Industrial schooling is architecturally engineered for a slow life history strategy. It demands high, immediate metabolic expenditure on secondary knowledge in exchange for distal returns—credential attainment, career entry—that arrive years later and depend on institutional stability the organism cannot verify. For a student whose developmental ecology has calibrated a fast strategy, this structure is not merely unappealing; it is informationally credible evidence that the promised payoff will not materialize. Appeals to future career benefit carry near-zero ancestral credibility in fast-calibrated metabolic systems. Their temporal discount function renders the expected return effectively zero regardless of the reward’s nominal size.

These observations are relevant to the Marshmallow Test reanalysis. Watts et al. (2018) demonstrated that delay of gratification in childhood predicted fewer outcomes than the original Mischel experiments suggested once family background was controlled. Kidd et al. (2013) showed that children in unreliable environments rationally chose the immediate marshmallow when their experience gave them grounds to doubt the experimenter’s reliability. The life-history integration of EDLT treats this not as a story about executive function but as a story about ecological rationality: the predictive engine has learned that delayed rewards in unpredictable environments are ecologically unreliable, so it discounts them steeply—a mathematically appropriate response, not a deficit.

The same temporal discount algorithm governs academic effort allocation. Primary versus secondary knowledge establishes the cost structure of engagement; life history calibration establishes the horizon over which returns are credible. Section 3 introduces Active Inference as the computational engine that integrates these parameters into a single decision architecture.

3 Empowerment-Driven Learning Theory

3.1 Core Mechanism: Active Inference and the Motivational Economy

The Computational Analogy: Active Inference

At the computational level of analysis, EDLT adopts the framework of predictive processing and Active Inference (Clark, 2013; Friston, 2010). The brain, on this account, is a hierarchical generative model that minimizes its expected free energy—an information-theoretic quantity that upper-bounds long-run surprise about sensory and interoceptive states. Two terminological clarifications are essential throughout: “free energy” in this account refers strictly to this computational quantity, while “metabolic cost” refers to the physical ATP

required to update and maintain the neural architecture executing these computations. The two are related but distinct.

The invocation of Active Inference here is explicitly analogical rather than formally derived. EDLT does not claim to implement the full mathematical apparatus of the free-energy principle; it draws on the principle’s core logic—that biological systems selectively allocate precision to prediction errors in proportion to their expected fitness relevance—as the computational motivation for EDLT’s resource-allocation account. The eventual formalization of this relationship, including derivation of EDLT’s parameters from the variational inference framework, is a target for subsequent modeling work. Crucially, while Active Inference explains how the brain allocates precision to prediction errors (Friston, 2010), the formal Free Energy Principle does not inherently contain a physical energy constraint. EDLT therefore bridges Friston’s informational architecture with Kurzban et al. (2013)’s opportunity-cost model. Friston explains the predictive hierarchy; Kurzban supplies the mechanism for the metabolic ROI gate: the system authorizes accommodation of secondary schemas—costly synaptic reorganization—only when the expected gain in future resource control, safety, or status justifies the physiological ATP expenditure (Westbrook & Braver, 2015).

Secondary knowledge acquisition is a particularly expensive subtype of this process. It demands sustained prefrontal engagement and metabolically intensive synaptic plasticity (Geary, 2008; Li & van Rossum, 2020). Following Kurzban et al. (2013)’s opportunity-cost account, the subjective experience of effort can be understood as a real-time signal that current resources might be more efficiently deployed elsewhere. EDLT formalizes that signal as the operative gate on academic engagement.

A note on the inferential structure of EDLT’s predictions is warranted here, because the theory operates across three distinct levels of analysis and conflating them is a source of interpretive error. The evolutionary level specifies the design logic of the motivational machinery: selection pressures shaped systems that conserve costly metabolic resources and allocate them preferentially toward survival- and status-relevant inference. This level ex-

plains why the machinery has the functional form it does, but it does not directly predict moment-to-moment behavior in contemporary institutional environments. The psychological mechanism level is where behavioral predictions are actually generated: it specifies how the evolutionarily shaped machinery manifests in the attractor states, temporal discounting functions, and ROI computations that adolescents execute in credential-driven school contexts. The behavioral output level—the observable academic effort, reallocation patterns, and phase transitions that EDLT’s hypotheses target—is downstream of the psychological mechanism level and mediated by individual developmental history, peer ecology, and institutional structure. The practical consequence is that EDLT’s core hypotheses (Section 5) are predictions at the psychological mechanism and behavioral output levels, grounded in but not directly deduced from the evolutionary level. Each link in this inferential chain requires separate empirical validation, and a failed behavioral prediction should be traced back through the chain to identify which level’s assumptions require revision—rather than interpreted as a global disconfirmation of the evolutionary framework.

Empowerment as a Kinematic Quantity

Standard motivational constructs treat empowerment as a psychological state—a feeling of agency, autonomy, or capability. EDLT respecifies it as a kinematic quantity to avoid the measurement ambiguity that attaches to introspective state reports. Empowerment, in this framework, is the positive first derivative of the organism’s expected utility over time: $dU/dt > 0$. An environment is empowering in the technical sense only if it is generating an expanding utility gradient—if the organism is actively acquiring new margins of survival security or social status.

The kinematic specification has several non-trivial implications. It explains reward habituation: when a student receives identical praise daily, absolute utility may remain high, but the first derivative drops toward zero, and metabolic withdrawal follows. It distinguishes motivating novelty (acceleration, a positive second derivative) from mere adequacy (posi-

tive but constant utility). Most consequentially, it predicts that the subjective experience accompanying genuine learning—what Csikszentmihalyi calls flow and what students call understanding something for the first time—corresponds specifically to a sudden positive acceleration in the utility gradient, which in turn drives the acute reward signal that clinicians and neuroscientists associate with dopaminergic contingency-learning circuits (Berridge & Kringelbach, 2015). The point is not that flow is a dopamine spike—the phenomenology of sustained flow and of prediction-error reward are distinct—but that the initiation of genuine insight has the kinematic signature of sudden acceleration, and that this connects it mechanically to reward circuitry in a way that the kinematic framework makes explicit.

The Two-Currency Parsimony and Temporal Discounting

EDLT proposes a parsimonious claim about which evolutionary payoffs the adolescent brain actively tracks in the institutional classroom. Rather than an unbounded array of fitness-relevant signals, the system compresses relevant value into two task-proximal currencies: *institutional survival* (avoiding failure, credential loss, or irreversible institutional exclusion) and *status* (standing within salient peer and teacher hierarchies). This two-currency parsimony follows from the observation that these are the two domains in which secondary knowledge most reliably translates into ancestrally legible payoffs—the teacher as credentialing authority, the peer group as the status ecology within which academic performance carries visible social meaning.

Life history calibration then determines how heavily future expressions of these currencies are discounted. For slow-strategy learners, a credential arriving in four years retains substantial discounted utility. For fast-strategy learners, the same credential at the same nominal magnitude may carry near-zero discounted value because their predictive engines—calibrated to environments where distal promises do not materialize—apply steep discount rates to all deferred returns. Crucially, this is not irrationality. It is the behavior of a predictive system that has learned from its developmental environment.

Discrete Attractor States

EDLT describes three operational modes of the motivational economy, modeled as attractor basins rather than points on a continuous spectrum. This discreteness is a deliberate simplification of what is, in formal Active Inference accounts, a continuous precision-weighting process—a simplification that is useful for generating categorical behavioral predictions while acknowledging that the underlying dynamics are graded. Future parameterization work should treat these as attractors in a continuous state space rather than strict categories.

Survival Mode obtains when the generative model encodes high probability of proximal punishment or status threat. Prediction errors related to threat receive extreme precision, mobilizing allostatic responses and compressing the policy space onto low-variance, low-exploration behaviors—rote compliance being the paradigm case (Immordino-Yang & Damasio, 2007). Velocity, in the kinematic sense, is effectively zero: the organism is spending ATP on threat-management rather than on any utility-expanding activity. The high metabolic expenditure is sustained not by dopaminergic reward but by cortisol-mediated mobilization, which predicts progressive exhaustion without the acquisition of new competencies.

Growth Mode emerges when high ecological safety coexists with credible opportunities for expanding utility—new competencies that map onto survival or status gains. Velocity is strongly positive. The system settles into a high-investment attractor, with dopaminergic contingency signals sustaining the expensive synaptic updating that genuine learning requires (Berridge & Kringelbach, 2015).

Conservation Mode reflects high safety but a collapsed utility gradient. When the marginal expected return on academic effort drops below the metabolic threshold—that is, when dU/dt approaches or falls below zero—the rational policy is disengagement or strategic minimalism. What orthodox accounts call “loss of motivation” is better described as the system registering that its resources are more efficiently preserved than expended. This distinction matters: Conservation Mode is not a failure state. It is the adaptive output of a

metabolic economy that has correctly evaluated its options.

3.2 The Efficiency Trap: Reinterpreting Proximate Constructs

Mindset as an ROI Prior

Mindsets, within EDLT, are not free-floating beliefs about the malleability of ability. They are compact probabilistic summaries—prior distributions, in Bayesian terms—over the expected metabolic return on future effort. A fixed mindset is not a cognitive deficit but a prior calibrated to environments where additional effort has historically failed to alter survival or status trajectories. The prior protects the organism from expending ATP in zero-yield contexts. It is efficient.

What this reframing generates that standard mindset theory does not is a conditional prediction about when growth mindset interventions will work. If mindsets are priors over environmental returns, then updating them through verbal instruction should succeed only when the ecological context makes the updated prior experientially credible—that is, when the environment genuinely delivers improved status returns in response to increased effort. This prediction maps directly onto Yeager et al. (2019)’s large-scale national experiment, in which growth mindset effects concentrated among lower-achieving students in schools with peer norms supporting academic effort. The effect was absent where ecological conditions made the updated prior implausible. The brain’s generative model is updated by experienced outcomes, not by verbal reframings alone; EDLT predicts this asymmetry from first principles.

The SDT Needs Under Metabolic Constraints

The evolutionary grounding of the SDT needs developed in Section 2.2 has a direct empirical consequence that SDT’s own framework cannot produce. SDT predicts that satisfying autonomy needs will increase engagement across ecological contexts, because need satisfaction is modeled as universally beneficial. EDLT makes a more restricted prediction: autonomy

satisfaction will increase engagement specifically when the autonomy granted connects to ecologically credible survival or status outcomes. In Conservation Mode following terminal credential attainment, even maximal autonomy support should produce near-zero motivational gain. The student’s generative model has already computed that no available academic action will expand the utility gradient; additional choice within the zero-return space adds nothing.

Each of the three SDT needs, when recontextualized through EDLT’s metabolic logic, functions as a proximate channel through which the survival/status numerator and metabolic cost denominator are updated. Each need reliably authorizes effort when the underlying payoff structure is intact, and each reliably fails when the payoff structure collapses—regardless of the phenomenological experience of need satisfaction. The intersection of SDT’s proximate accuracy with EDLT’s ultimate constraint specification is precisely where the two frameworks generate non-overlapping predictions, and those predictions are what Section 5 develops into a testable research program.

4 Rational Detachment: The EDLT Explanation

4.1 The Phase Transition

Proposition 1. *Post-acceptance disengagement will be most severe among previously compliant, slow-life-history students whose survival and status portfolios were concentrated in academic performance, and weakest among students whose portfolios were distributed across multiple domains or preserved by conditional post-acceptance stakes.*

Senioritis—the marked reduction in academic effort following postsecondary admission—is routinely attributed to failures of grit or character. That interpretation mistakes a symptom for a mechanism. Rational Detachment is a mathematically lawful phase transition driven by a parameter shift in the motivational economy, not a character-level deficit.

Before college acceptance, high-achieving students occupy the Growth Mode attractor.

Cumulative GPA functions as a proximal lever on life-history-relevant outcomes: admission to a credential-granting institution that expands access to high-status adult ecologies (Ellis et al., 2012; Geary, 2008). The survival utility gradient is positive and steep. Policies involving sustained study minimize expected free energy by increasing the probability of attaining those preferred future states (Gershman, 2019), and this expected gain justifies the biological cost of sustained synaptic updating.

The acceptance letter operates as an exogenous parameter shift. More precisely, what reaches its asymptote at that moment is the marginal utility of current academic performance for the specific goal of credential attainment—not survival security in any global sense, and not utility across all domains. Students remain exposed to financial insecurity, social status competition, and institutional uncertainty; the acceptance does not eliminate these. What it eliminates is the contingency between current GPA performance and admission outcome. That specific contingency was the primary driver sustaining high-investment academic behavior. Its removal collapses the utility gradient for academic effort, prompting a transition into Conservation Mode—not because the student has become globally amotivated, but because the active-inference engine has correctly identified that further academic investment yields no marginal return on the specific payoff that had been authorizing it.

The transition is predicted to be rapid, occurring over days to weeks following the perceived removal of the contingency. This timescale reflects the speed of prior updating in Bayesian learning systems when evidence is decisive: the acceptance letter is a high-precision signal that essentially eliminates uncertainty about the credential outcome, producing a near-immediate revision of the system’s policy priorities. This is distinct from burnout, which unfolds over months and reflects depletion rather than reallocation.

4.2 Metabolic Reallocation and Moderation

The post-acceptance behavioral shift is not evidence of global amotivation. A student who abandons coursework but simultaneously intensifies social coordination, competitive gaming,

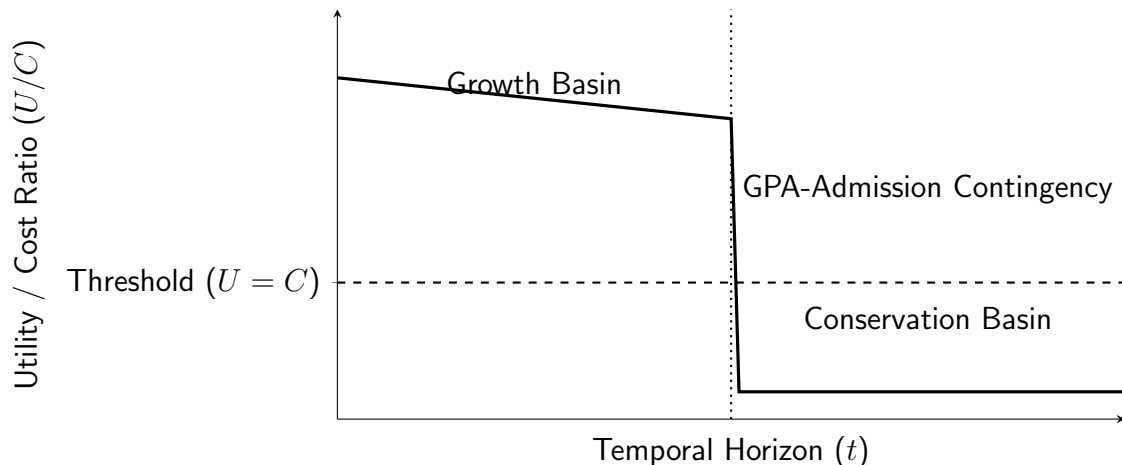


Figure 1: Illustrative threshold collapse graph depicting the hypothesized EDLT phase transition. The curve is schematic, not derived from empirical data. It depicts the system maintaining Growth Mode while the expected utility-to-cost ratio exceeds the engagement threshold; once the GPA-admission contingency is removed, the ratio drops below threshold and the system transitions into the Conservation attractor. The shape of the transition—abrupt versus gradual—is itself a testable prediction (see Hypothesis 1).

or exploratory research into their future institutional environment is not less motivated in any aggregate sense—they are differently motivated, in a manner that reveals the underlying system’s continued operation. The active-inference engine has identified domains where the utility gradient remains positive and has reallocated resources accordingly (Friston, 2010; Gershman, 2019). This reallocation is EDLT’s central behavioral signature: Conservation Mode for academic effort should co-occur with Growth Mode for social or status-relevant nonacademic effort.

Table 2 extends this logic across a range of motivational anomalies, translating phenomenological descriptions into the kinematic variables that EDLT uses to characterize system outputs. As with Table 1, the EDLT mechanism column represents theoretical hypotheses subject to empirical evaluation.

Life history calibration moderates the magnitude of post-acceptance reallocation. EDLT predicts stronger drops among students with: (a) high pre-acceptance grade-contingent self-worth, (b) few alternative status channels outside academic performance, and (c) a developmental history of environmental stability—conditions characteristic of slow-strategy orienta-

tions (Del Giudice, 2020). When acceptance flattens the academic utility gradient, students with slow-strategy profiles face the hardest search problem: they must identify alternative domains capable of generating positive velocity when their entire motivational infrastructure was concentrated in a single channel.

Fast-strategy students, by contrast, tend to show smaller and more diffuse post-acceptance shifts because academic effort was rarely their primary survival or status channel (Ellis et al., 2012). The directionality of this effect, however, depends on how the academic domain was encoded prior to acceptance. If school was primarily encoded as a threat-avoidance domain, removal of the threat may paradoxically produce visible effort avoidance as an active disengagement from a previously aversive context. If school was encoded as a low-opportunity domain, the post-acceptance change will be minimal—a flatline rather than a collapse. Current longitudinal data cannot adjudicate between these two fast-strategy pathways because most studies of post-acceptance behavior have not stratified samples by prior ecological encoding. This is a productive empirical ambiguity, not a theoretical underspecification: EDLT generates distinct predictions for each pathway given appropriate measurement.

The failure of proximate interventions to reverse post-acceptance disengagement follows directly from this analysis. Autonomy-supportive climates address phenomenological need satisfaction. They do not alter the objective contingency between current academic performance and any life-history-relevant outcome—because, post-acceptance, there is no such contingency remaining. Reversing Rational Detachment requires not improved pedagogy but a genuine structural change to the payoff landscape: a contingency between current performance and a future state that the student’s generative model encodes as survival or status relevant, delivered with sufficient consistency that the precision assigned to the new signal overcomes the prior’s resistance to updating.

5 Boundary Conditions and Testable Predictions

5.1 Boundary Conditions: Where the Model Predicts Its Own Failure

EDLT is not a theory that applies to all learners under all conditions. Its parameters operate within biological and ecological boundaries, and identifying those boundaries precisely is the mechanism through which the theory generates its most distinctive and falsifiable predictions. Each boundary condition articulates a set of ecological parameters under which the active-inference engine cannot execute the metabolic ROI calculation that EDLT takes as its central mechanism—not because the theory fails, but because it predicts exactly this failure mode.

Boundary 1: The Executive Function Threshold

EDLT assumes the brain computes a metabolic return on investment for secondary learning. This computation itself demands executive resources: working memory sufficient to maintain task goals across the delay between effort and outcome, and inhibitory control sufficient to suppress competing response tendencies during novel schema accommodation. When these executive function capacities are severely constrained, the computation breaks down in a specific way: the working-memory cost of maintaining abstract symbol sequences (e.g., algebraic steps) inflates the denominator of the motivational economy, often beyond what any realistic survival or status numerator can overcome (Sweller, 1988).

The important implication is that applying extreme survival threat to learners with low executive function does not solve this problem—it compounds it. Threat management and novel schema building are competing consumers of the same limited executive budget. When both demands are simultaneously high, the system defaults to the evolutionarily older response: assimilation of existing schemas rather than accommodation of new ones (Friston, 2010). The behavioral result is rigid compliance or Rational Detachment, not engagement.

Among the components of executive function, the most relevant to this boundary con-

dition is updating—the capacity to refresh working memory representations as new information arrives. Shifting and inhibition are also implicated but are less directly linked to the information-integration demands of secondary knowledge acquisition. Future empirical work should operationalize the EF threshold with measures sensitive to updating capacity specifically.

The Falsification Test. If learners with severely constrained updating capacity spontaneously engage in complex, unstructured secondary tasks without external scaffolding or proximate rewards, the EF threshold boundary is invalid.

Boundary 2: The Cultural Epistemic Boundary

EDLT posits that the metabolic cost of secondary learning can be subsidized by the acquisition of social status. The mechanism of status acquisition, however, is not culturally uniform. Evolutionary anthropology distinguishes two pathways: Prestige (freely conferred deference for demonstrated competence) and Dominance (compliance extracted through threat) (Henrich & Gil-White, 2001). These two pathways activate fundamentally different motivational architectures.

Constructivist and autonomy-supportive pedagogies are designed for Prestige economies. They assume learners are ecologically safe enough to explore, make errors, and revise—that the cost of public failure is recoverable. In Dominance-avoidance contexts, such as the high-stakes testing environments documented in Confucian Heritage Culture research (Chen, 2023; King & McInerney, 2012), this assumption does not hold. Failure carries existential institutional threat, and open-ended exploration introduces precisely the prediction errors that the threat-sensitized system is most motivated to prevent. The consequence, documented empirically, is high aggregate achievement coexisting with intense pressure, obligation-oriented motivation, and systematic avoidance of the exploratory strategies that Western pedagogies target (King & McInerney, 2012).

EDLT predicts that transplanting autonomy-supportive instructional models into Dominance-

avoidance ecologies will not generate Growth Mode. In the absence of ecological safety, increased autonomy paradoxically inflates perceived threat by removing the predictable structure within which the student could at least map the cost of compliance. EDLT's prediction is that exploratory learning requires not merely autonomy but a prior reduction in the ecological cost of failure.

The Falsification Test. If a strict Dominance-avoidance ecology generates sustained, high-metabolic exploratory learning without any reduction in the systemic cost of public failure, the cultural epistemic boundary is falsified.

Boundary 3: The Evolutionary Satiety Threshold

Marginal utility diminishes as resource security and social standing approach saturation—a principle shared by economic and evolutionary reasoning alike (Ahl et al., 2023). Once the predictive engine infers that the organism is reliably buffered against foreseeable disruptions, additional academic effort produces vanishing marginal returns in expected free-energy reduction. This satiety dynamic explains a pattern that proximate theories find puzzling: students from highly resourced families, with strong social capital and institutional protection, who disengage despite apparently favorable motivational conditions.

Their disengagement is not curiosity deficit or motivational failure. It is the behavior of a system that has genuinely optimized: further academic investment offers no meaningful expansion of security or status in an ecology where both are already saturated. Re-engaging such learners requires introducing genuinely novel utility—a fresh competitive arena, a scarcity signal, or a coalition opportunity—that creates positive velocity in a domain the student's existing resource base cannot guarantee.

The Falsification Test. If adolescents with demonstrably saturated resource security and social dominance sustain high-metabolic secondary knowledge investment without the introduction of novel scarcity or status signals, the satiety threshold is invalid.

Boundary 4: The Temporal Discounting Threshold

The temporal horizon over which the motivational economy evaluates returns is itself a free parameter calibrated by life history experience (Ellis et al., 2009). For fast-strategy learners, this parameter—the discount rate, k —is steep: the subjective value of a promised payoff declines sharply with temporal distance. Critically, increasing the nominal size of a distant reward does not reliably overcome this discount once k is set: the discounted expected utility asymptotes at near zero regardless of reward magnitude, because the discount function dominates at the timescales that institutional schooling typically invokes.

The threshold at which temporal discounting vetoes engagement is therefore not a fixed calendar duration but a function of k : engagement fails when the discounted expected return drops below the undiscounted metabolic cost. For a fast-strategy learner with a steep discount function, this threshold may be crossed within weeks; for a slow-strategy learner, it may tolerate years of deferred return. Interventions designed around distal credential payoffs will fail for fast-strategy learners not because those learners cannot plan, but because their predictive engines have been calibrated, accurately, to an ecology in which distal promises are unreliable.

The Falsification Test. If learners from highly unstable developmental environments reliably sustain high-metabolic academic effort for payoffs beyond their discount horizon—without the introduction of immediate, proximal survival or status scaffolding—the temporal discounting boundary is falsified.

5.2 Testable Predictions and Proposed Measures

A framework of this scope requires coordinated methodological development. The primary obstacle to testing evolutionary hypotheses in classroom settings is measurement substitution: standard motivation instruments track attitudinal constructs (interest, value, efficacy) that do not map cleanly onto the ancestral currencies EDLT identifies as operative. We propose three measurement tools and four core hypotheses.

Proposed Methodological Tools

The **Ancestral Relevance Scale (ARS)** is designed to operationalize the numerator of the motivational economy: perceived evolutionary currency. Item generation will target Status Utility (e.g., “Success on this task is immediately visible to peers whose judgment matters to me”) and Survival/Threat Utility separately, rather than aggregating them into a general value score as EVT’s utility value subscale does (Wigfield & Eccles, 2000). The ARS is distinct from existing evolutionary utility instruments—including Singh’s reproductive value scales and Kenrick’s CSEA—in being designed specifically for institutional academic contexts where the relevant currencies are credential-mediated status and threat-avoidance rather than reproductive or mate-selection domains. Construct validity (Flake et al., 2017) should be established by demonstrating that ARS scores predict task-initiation latency under conditions of zero intrinsic interest more strongly than EVT utility value scores do.

The **Mini-K** (Figueredo et al., 2006) is a validated brief index of slow-fast life history orientation. EDLT proposes its integration as a standard covariate in motivational research, both to test life-history-moderated predictions directly and to control for the substantial between-person variance in temporal discounting that otherwise confounds motivational studies.

Standard self-report surveys cannot detect the rapid, threshold-driven state changes that EDLT predicts. **Ecological Momentary Assessment (EMA)** (Shiffman et al., 2008) is the natural measurement tool for this purpose: it can track real-time fluctuations in perceived effort cost, status salience, and behavioral engagement within the narrow temporal window—seven to fourteen days—where institutional parameter shifts occur. EMA data also provide the density needed to distinguish between gradual linear decline (what burnout models predict) and abrupt threshold crossing (what EDLT predicts).

Testable Hypotheses

Hypothesis 1: The Phase Transition Hypothesis. Burnout models predict gradual, approximately linear deterioration in academic effort across an academic semester. EDLT predicts a non-linear, threshold-driven drop precisely temporally correlated with the removal of an institutional contingency (e.g., credential attainment). The collapse will be concentrated within a narrow window of days to weeks, distinguishable from gradual decline through Bayesian change-point analysis of daily EMA effort ratings. The temporal correlation between the phase transition and the contingency removal event is the signature prediction.

Hypothesis 2: The Reallocation Hypothesis. Rational Detachment is not global amotivation. Drawing on goal disengagement and reengagement research (Wrosch & Miller, 2009), EDLT predicts a systematic inverse relationship between academic effort and social or status-seeking effort in the post-acceptance window: as academic engagement drops, effort allocated to social coalition, competitive nonacademic activities, or post-graduation exploration will rise. This relationship should be detectable in EMA data at the individual level, not merely at the group mean, and its magnitude should be moderated by pre-acceptance life history calibration (Mini-K scores).

Hypothesis 3: The Metabolic Framing Hypothesis. The denominator of the motivational economy reflects the perceived, not merely objective, cost of engagement. Framing an identical secondary task as one that builds on existing knowledge (assimilation frame) versus as genuinely new challenging material (accommodation frame) should alter the prior on expected metabolic cost and, consequently, task initiation latency. The EDLT prediction is that assimilation framing produces significantly faster task initiation on objectively identical tasks. This effect is predicted to be strongest for learners operating near the engagement threshold—those for whom a modest denominator reduction is sufficient to cross into engagement—and attenuated for learners in clear Growth Mode or clear Conservation Mode. We acknowledge the theoretical tension this hypothesis creates: if framing alone can shift

the effective cost prior, it raises the question of how much of EDLT’s explanation reduces to expectancy effects. The prediction is therefore accompanied by a test of the moderating role of evolutionary currency: the framing effect should interact with ARS-measured status salience, with the assimilation advantage largest when status utility is high.

Hypothesis 4: The Multiplicative Suppression Hypothesis. The most structurally distinctive claim EDLT makes relative to Expectancy-Value Theory is that cost interacts with utility through a ratio rather than an additive structure. EVT’s cost dimension suppresses engagement additively: doubling cost reduces engagement by a fixed amount regardless of utility level. EDLT predicts multiplicative suppression: when metabolic cost is extremely high, it vetoes engagement even when survival or status utility is independently held at a high level—a pattern additive models cannot generate. To test this, we propose a 2×2 factorial design crossing utility level (low vs. high, manipulated via ARS-validated status salience framing) with metabolic cost level (low vs. high, manipulated via cognitive load induction). The critical prediction is a disordinal interaction: under high-cost conditions, high utility should produce no reliable engagement advantage over low utility—because the cost veto operates independently of the utility signal. Under low-cost conditions, the standard utility advantage should emerge. A purely additive model predicts a main effect of utility that is consistent across cost conditions, without the interaction. This test is necessary to establish that EDLT’s structural departure from EVT is empirically warranted rather than terminologically distinct.

6 Discussion

6.1 Theoretical Contributions

EDLT advances the study of academic motivation by integrating three explanatory levels that the literature has, to date, left largely separate: the ultimate constraint of evolutionary mismatch (Geary, 2008), the computational logic of Active Inference (Friston, 2010), and

the proximate phenomenological architecture of SDT (Ryan & Deci, 2017). The value of this integration lies not in any single piece but in what the combination makes possible: predictions that no single level can generate alone.

The tautology risk for predictive-processing frameworks deserves direct attention. Because any behavior can theoretically be explained post hoc by positing an appropriate prior, unconstrained applications of the free-energy principle can lose falsifiability (Gershman, 2019). EDLT addresses this by specifying, a priori, which priors govern adolescent academic ecologies: institutional survival and local status. By locking these priors before applying the computational framework, EDLT generates domain-specific predictions about which prediction errors the adolescent brain will prioritize—a move that converts the free-energy principle from an explanatory schema into a testable model.

The Rational Detachment account illustrates what this constraint buys. When terminal credential attainment removes the GPA-to-admission contingency, the survival prior is instantly satisfied—not globally, but for that specific contingency. The subsequent behavioral signature (academic disengagement paired with social effort increase) is not surprising from an EDLT perspective; it is the expected output of a system rationally redistributing expected free energy. What makes this a theoretical contribution rather than a redescription is that EDLT specifies, in advance, what modulating conditions should amplify or attenuate the effect (life history calibration, portfolio concentration, contingency specificity) and what co-occurring behavioral changes should accompany it (the reallocation to social domains).

A comparison with Expectancy-Value Theory’s cost dimension (Flake et al., 2017; Wigfield & Eccles, 2000) is warranted here because cost is the most obvious precedent for EDLT’s metabolic denominator. EVT’s cost construct captures perceived investment demands—time, effort, forgone alternatives—as a component of subjective task value. EDLT’s metabolic cost is operationalized differently in two respects. First, it is grounded in a specific biological mechanism (ATP expenditure and opportunity cost in a predictive-processing system) rather than a subjective appraisal, which means it generates predictions about when cost

will be perceived regardless of conscious evaluation—specifically under high threat or steep temporal discounting. Second, EDLT’s denominator interacts with the numerator through a ratio structure that EVT does not share: extremely high cost does not merely reduce value additively but can suppress engagement multiplicatively, even when the value numerator is objectively high. Whether this distinction generates empirically discriminable predictions is itself a testable question.

6.2 Practical Implications: Classes of System Perturbation

If motivation is an evolutionary economic calculation, educational reform should target the parameters of the economy rather than the student’s attitude toward engagement. The aim is not to inspire students but to alter the structural inputs that the active-inference engine computes.

Table 3 outlines four theoretical intervention classes organized by the parameter each targets. These are structural blueprints, not prescriptions: the domain-specific application of each class requires separate empirical development.

One structural implication of this framework has particular policy salience and deserves more prominence than it typically receives. The motivational economy described above applies to teachers as much as to students. Demanding that practitioners function as real-time schema diagnosticians—identifying individual students’ prior models and intervening precisely—requires a substantial and sustained metabolic investment. A teacher operating in institutional Survival Mode, under high-stakes evaluation threats with inadequate autonomy and resource support, will systematically default to low-cost, compliance-based instruction (Collie et al., 2012). This is not a failure of professional commitment; it is the predictable output of an active-inference system that has computed the opportunity cost of adaptive pedagogy against the threat cost of institutional exposure. Systemic reform that demands sophisticated motivational responsiveness from practitioners without first addressing the metabolic economy in which those practitioners operate is not merely inefficient—it

is incoherent within EDLT’s own theoretical framework.

6.3 Future Directions: Formal Parameterization

The present framework requires translation into formal computational models before its predictions can be tested with the precision they claim. Several avenues are tractable in the near term. Drift Diffusion Models of decision-making (Ratcliff & McKoon, 2008) offer a natural vehicle for parameterizing the engagement threshold: the drift rate toward the “engage” boundary can be operationalized as a function of the perceived utility-to-cost ratio, and the phase transition hypothesis (Hypothesis 1) becomes a prediction about drift rate change over time. Continuous biometric markers—pupillometry, heart-rate variability, electrodermal activity—can serve as proxies for metabolic cost in real-time assessment contexts, though their mapping onto ATP expenditure specifically requires validation against more direct metabolic measures. Delay discounting behavioral tasks provide a validated window onto the discount parameter k and can be administered alongside EMA protocols to test the temporal discounting threshold predictions in Section 5.

A second priority is developing multi-level models capable of separating within-person fluctuations in motivational state from between-person differences in life history calibration. EDLT makes predictions at both levels, and the two can only be distinguished with study designs that include both intensive longitudinal measurement (EMA) and between-person covariates (Mini-K, ARS).

6.4 Limitations

Several constraints bound the current framework’s claims. At the conceptual level, EDLT’s evolutionary arguments are inferential: we can measure current life history calibration through instruments like the Mini-K, but ancestral environments must be characterized by extrapolation from developmental ecology and cross-cultural data. This limits the precision of the ultimate-level argument and requires that life history typologies be treated as contin-

uous descriptors of environmental calibration—never as fixed types, and emphatically never as normative categories that rank learners against each other.

Readers should also be aware that several foundational literatures on which EDLT draws have faced replication scrutiny. Life history strategy research has seen contested findings in the dual mating strategy literature (Del Giudice, 2020). Some effects in the growth mindset literature have replicated inconsistently across populations and measurement contexts (Sisk et al., 2018). Geary’s primary/secondary knowledge distinction, while widely cited, has been critiqued on grounds of construct underspecification (Geary, 2008). EDLT’s robustness to these uncertainties should be tested by examining which of its core predictions survive if the contested foundational claims are weakened. The boundary conditions in Section 5 represent one step in this direction: they constrain the theory’s claims rather than asserting universal applicability.

Empirically, the translation of Active Inference mathematics to classroom ecologies remains nascent. The current account relies on self-report proxies for quantities (metabolic cost, survival utility) that the theory treats as computational rather than introspective. Aggressive multimethod triangulation—combining EMA, biometric indicators, behavioral latency measures, and self-report—is needed before any single operationalization can claim theoretical adequacy.

Finally, EDLT is calibrated specifically for adolescent, credential-driven institutional schooling. Its application to early childhood learning, adult voluntary education, or informal learning contexts is speculative and would require separate theoretical derivation. The mechanisms operating in credential-free learning environments—where the survival and status currencies take different forms, and where temporal discounting plays out across very different timescales—may require substantial revisions to the core Motivational Economy model.

A related boundary deserves explicit acknowledgment. EDLT’s two-currency parsimony treats the status currency as uniform across the student population. The evolutionary lit-

erature does not support this uniformity assumption. Consistent cross-cultural evidence indicates that the pathways to status acquisition, and the social domains in which status competition is most acutely tracked, differ systematically between male and female adolescents (Geary, 2008). Male status hierarchies in adolescent ecologies are more consistently organized around competitive dominance and visible resource acquisition; female status hierarchies more consistently around coalition quality, social network centrality, and relational reputation. These are distributional differences, not deterministic categories, but they are sufficiently robust that a theory using evolutionary biology as its mechanistic foundation cannot treat the status numerator as gender-neutral without incurring an explanatory debt. The implications for EDLT’s predictions in mixed-gender academic ecologies—including whether the post-acceptance reallocation pattern in Hypothesis 2 should differ in direction or magnitude by gender—require separate theoretical derivation and constitute a productive empirical extension of the current framework.

7 Conclusion

What the preceding argument establishes is something the introduction could only gesture at: that the explanatory gap in academic motivation research is not incidental but structural, and that closing it requires genuinely integrating the evolutionary and computational constraints that govern cognitive resource allocation with the proximate frameworks that have described their phenomenological effects.

The standard assumption—that sustained cognitive effort in school is biologically natural and disengagement is a failure—has generated a literature rich in proximate description but impoverished in mechanistic prediction. When a student with documented grit, growth mindset beliefs, and a supportive classroom environment abruptly abandons academic effort the week after an acceptance letter arrives, proximate theories record an anomaly they cannot explain. EDLT’s account makes that outcome predictable. The acceptance letter

does not weaken the student’s character. It removes the specific contingency—between GPA performance and credential attainment—that had been converting costly secondary knowledge acquisition into evolutionarily legible payoff. When that conversion mechanism is gone, the active-inference engine performs exactly as it should: it redirects resources toward domains where the utility gradient remains positive, and it conserves resources in domains where the gradient has collapsed.

Recognizing this has a design implication that extends well beyond senioritis. If adolescent learners are metabolic allocators rather than naturally curious information absorbers, then the task of education is not to extract effort through moral appeals to self-improvement but to engineer ecologies where the high cost of secondary knowledge acquisition is credibly offset by immediate, ecologically legible returns. That is a structural problem—one of contingency design, feedback timing, status currency structure, and threat management—not a problem of student character or teacher inspiration.

The boundary conditions and testable predictions developed in Section 5 convert this structural argument into a falsifiable research program. The hypotheses are specific enough to be disconfirmed, and the measurement tools proposed are adequate to the task of disconfirming them. Whether EDLT’s evolutionary and computational commitments survive contact with that evidence is an open empirical question—and answering it will require exactly the kind of integrative, multi-level research design that has been largely absent from the motivation literature. That absence is itself the gap this framework was built to close.

References

- Ahl, R. E., Cook, E., & McAuliffe, K. (2023). Having less means wanting more: Children hold an intuitive economic theory of diminishing marginal utility. *Cognition*, *234*, 105367. <https://doi.org/10.1016/j.cognition.2023.105367>

- 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>
- Chen, S.-W. (2023). Learning motivations and effort beliefs in confucian cultural context: A dual-mode theoretical framework of achievement goal. *Frontiers in Psychology*, *14*, 1058456. <https://doi.org/10.3389/fpsyg.2023.1058456>
- Clark, A. (2013). Whatever next? predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences*, *36*(3), 181–204. <https://doi.org/10.1017/S0140525X12000477>
- Collie, R. J., Shapka, J. D., & Perry, N. E. (2012). School climate and social–emotional learning: Predicting teachers’ stress, job satisfaction, and teaching efficacy. *Journal of Educational Psychology*, *104*(4), 1189–1204. <https://doi.org/10.1037/a0029356>
- Cosmides, L., & Tooby, J. (1997). Evolutionary psychology: A primer. *Center for Evolutionary Psychology*.
- Deci, E. L., & Ryan, R. M. (2000). The “what” and “why” of goal pursuits: Human needs and the self-determination of behavior. *Psychological Inquiry*, *11*(4), 227–268.
- Del Giudice, M. (2020). Rethinking the fast-slow continuum of individual differences. *Evolution and Human Behavior*, *41*(6), 536–549. <https://doi.org/10.1016/j.evolhumbehav.2020.07.006>
- Duckworth, A. L., Peterson, C., Matthews, M. D., & Kelly, D. R. (2007). Grit: Perseverance and passion for long-term goals. *Journal of Personality and Social Psychology*, *92*(6), 1087–1101. <https://doi.org/10.1037/0022-3514.92.6.1087>
- Dweck, C. S. (2006). *Mindset: The new psychology of success*. Random House.
- Ellis, B. J., Del Giudice, M., Dishion, T. J., Figueredo, A. J., Gray, P., Griskevicius, V., Hawley, P. H., Jacobs, W. J., James, J., Volk, A. A., & Wilson, D. S. (2012). The evolutionary basis of risky adolescent behavior: Implications for science, policy, and practice. *Developmental Psychology*, *48*(3), 598–623. <https://doi.org/10.1037/a0026220>

- Ellis, B. J., Figueredo, A. J., Brumbach, B. H., & Schlomer, G. L. (2009). Fundamental dimensions of environmental risk: The impact of harsh versus unpredictable environments on the evolution and development of life history strategies. *Human Nature*, *20*(2), 204–268. <https://doi.org/10.1007/s12110-009-9063-7>
- Figueredo, A. J., Vásquez, G., Brumbach, B. H., Sefcek, J. A., Kirsner, B. R., & Jacobs, W. J. (2006). The k-factor, covitality, and personality: A psychometric test of life history theory. *Human Nature*, *17*(3), 243–272.
- Flake, J. K., Pekrun, R., Kouros, C. R., Jones, K., & Fried, C. (2017). Measure use in psychology: A review of construct validity. *Social Psychological and Personality Science*, *8*(4), 370–379. <https://doi.org/10.1177/1948550617693063>
- Friston, K. J. (2010). The free-energy principle: A unified brain theory? *Nature Reviews Neuroscience*, *11*(2), 127–138. <https://doi.org/10.1038/nrn2787>
- Gallup, Inc. (2017). Student enthusiasm falls as high school graduation nears [Available from <https://news.gallup.com/opinion/gallup/211631/student-enthusiasm-falls-high-school-graduation-nears.aspx>].
- Geary, D. C. (2007). Educating the evolved mind: Conceptual foundations for an evolutionary educational psychology. In J. S. Carlson & J. R. Levin (Eds.), *Educating the evolved mind* (pp. 1–99). Information Age Publishing.
- Geary, D. C. (2008). An evolutionarily informed education science. *Educational Psychologist*, *43*(4), 179–195. <https://doi.org/10.1080/4.61520802392133e+14>
- Gershman, S. J. (2019). What does the free energy principle tell us about the brain? *Neuroscience and Biobehavioral Reviews*, *104*, 233–248. <https://doi.org/10.1016/j.neubiorev.2019.06.021>
- Henrich, J., & Gil-White, F. J. (2001). The evolution of prestige: Freely conferred deference as a mechanism for enhancing the benefits of cultural transmission. *Evolution and Human Behavior*, *22*(3), 165–196. [https://doi.org/10.1016/S1090-5138\(00\)00071-4](https://doi.org/10.1016/S1090-5138(00)00071-4)

- Hidi, S., & Renninger, K. A. (2006). The four-phase model of interest development. *Educational Psychologist*, *41*(2), 111–127. https://doi.org/10.1207/s15326985ep4102_4
- Immordino-Yang, M. H., & Damasio, A. (2007). We feel, therefore we learn: The relevance of affective and social neuroscience to education. *Mind, Brain, and Education*, *1*(1), 3–10. <https://doi.org/10.1111/j.1751-228X.2007.00004.x>
- Kidd, C., Palmeri, H., & Aslin, R. N. (2013). Rational snacking: Young children’s decision-making on the marshmallow task is moderated by beliefs about environmental reliability. *Cognition*, *126*(1), 109–114.
- King, R. B., & McInerney, D. M. (2012). Family obligation and academic engagement: A study of chinese students. *Educational Psychology*, *32*(4), 445–458.
- Kurzban, R., Duckworth, A., Kable, J. W., & Myers, J. (2013). An opportunity cost model of subjective effort and task performance. *Behavioral and Brain Sciences*, *36*(6), 661–679.
- Li, S., & van Rossum, M. C. W. (2020). Energy efficient synaptic plasticity. *eLife*, *9*, e50804. <https://doi.org/10.7554/eLife.50804>
- National Center on Education and the Economy. (2022, August). Scaling the “school engagement cliff” [Available from <https://ncee.org/scaling-the-school-engagement-cliff/>].
- Pekrun, R. (2006). The control-value theory of achievement emotions: Assumptions, corollaries, and implications for educational research and practice. *Educational Psychology Review*, *18*(4), 315–341. <https://doi.org/10.1007/s10648-006-9029-9>
- Ratcliff, R., & McKoon, G. (2008). The diffusion decision model: Theory and data for two-choice decision tasks. *Neural Computation*, *20*(4), 873–922. <https://doi.org/10.1162/neco.2008.12-06-420>
- Ryan, R. M., & Deci, E. L. (2000). Self-determination theory and the facilitation of intrinsic motivation, social development, and well-being. *American Psychologist*, *55*(1), 68–78. <https://doi.org/10.1037/0003-066X.55.1.68>

- Ryan, R. M., & Deci, E. L. (2017). *Self-determination theory: Basic psychological needs in motivation, development, and wellness*. Guilford Press.
- Shiffman, S., Stone, A. A., & Hufford, M. R. (2008). Ecological momentary assessment. *Annual Review of Clinical Psychology, 4*, 1–32. <https://doi.org/10.1146/annurev.clinpsy.3.022806.091415>
- Sisk, V. F., Burgoyne, A. P., Sun, J., Lawson, J. L., & Macnamara, B. N. (2018). To what extent and under which circumstances are growth mind-sets important to academic achievement? two meta-analyses. *Psychological Science, 29*(4), 549–571.
- Sweller, J. (1988). Cognitive load during problem solving: Effects on learning. *Cognitive Science, 12*(2), 257–285. https://doi.org/10.1207/s15516709cog1202_4
- Tinbergen, N. (1963). On aims and methods of ethology. *Zeitschrift für Tierpsychologie, 20*(4), 410–433.
- Watts, T. W., Duncan, G. J., & Quan, H. (2018). Revisiting the marshmallow test: A conceptual replication investigating links between early delay of gratification and later outcomes. *Psychological Science, 29*(7), 1159–1177.
- Westbrook, A., & Braver, T. S. (2015). Cognitive effort: A neuroeconomic approach. *Cognitive, Affective, & Behavioral Neuroscience, 15*(2), 395–415.
- Wigfield, A., & Eccles, J. S. (2000). Expectancy–value theory of achievement motivation. *Contemporary Educational Psychology, 25*(1), 68–81. <https://doi.org/10.1006/ceps.1999.1015>
- Wrosch, C., & Miller, G. E. (2009). Depressive symptoms can be useful: Self-regulatory and emotional benefits of goal disengagement in response to untreatable life regrets. *Journal of Personality and Social Psychology, 96*(6), 1181–1190. <https://doi.org/10.1037/a0015172>
- Xu, Z., Goldhaber, D., & Dao, R. (2025). *Student engagement in charter and traditional public schools before and after the covid-19 pandemic* (Working Paper No. 331-1225) (Available from <https://caldercenter.org/sites/default/files/2025-12/CALDER-WP->

[331-1225_0.pdf](#)). CALDER (National Center for Analysis of Longitudinal Data in Education Research).

Yeager, D. S., Hanselman, P., Walton, G. M., et al. (2019). A national experiment reveals where a growth mindset improves achievement. *Nature*, *573*, 364–369. <https://doi.org/10.1038/s41586-019-1466-y>

Table 1: Paradigmatic Contrasts in Academic Engagement. The EDLT Mechanism column presents theoretical propositions, not established findings; their empirical status is addressed in Section 5.

Observable Phenomenon	Orthodox Paradigm & Representative Frameworks	EDLT Hypothesis (Metabolic ROI)	Testable Prediction / Boundary Condition
Abrupt Disengagement Post-Acceptance (Senioritis)	Goal-Setting Theory; SDT: Diagnosed as a collapse of external contingencies or failure to internalize intrinsic motivation.	Rational Detachment (Conservation Mode): The institutional threat sustaining effort is removed. The metabolic case for further investment collapses.	Prediction: Disengagement onset correlates with the perceived irrevocability of the credential, independent of measured “intrinsic” trait motivation.
Severe Academic Procrastination	Self-Regulation Theory: Diagnosed as a failure of executive function, poor time management, or emotional dysregulation.	Just-in-Time Metabolism: Metabolic expenditure is deferred until deadline proximity spikes the survival utility numerator above the engagement threshold.	Prediction: High-status immediate micro-rewards will reduce procrastination more efficiently than executive function training alone.
Rejection of Unstructured Inquiry in High-Stakes Environments	Mindset Theory; Constructivism: Pathologized as a “fixed mindset” or a deficit in critical thinking and resilience.	Epistemic Boundary Defense (Survival Mode): In dominance economies, open-ended exploration risks status humiliation. Apparent rigidity is a metabolically rational protective strategy.	Boundary Condition: Exploratory learning will succeed only when the ecological cost of failure is sufficiently low that Growth Mode is accessible.
Gifted Burnout / Refusal to Exert Effort	Grit: Labeled as a lack of perseverance and passion for long-term goals.	The Efficiency Trap: The environment lacks credible future payoffs for continued extreme metabolic exertion. Withdrawal is biological efficiency, not character failure.	Prediction: Sustained effort requires that perceived status payoff credibly exceed metabolic cost; when this condition fails, grit provides no additional predictive power.

Table 2: System Outputs: Phenomenological Anomalies Mapped to EDLT Kinematic States (Theoretical Hypotheses)

Phenomenological Anomaly	Ecological Trigger (Parameter Shift)	Kinematic State (Effort/ATP Allocation)	EDLT Hypothesized Mechanism
Rational Detachment (e.g., Senioritis)	The GPA-to-credential contingency is removed (e.g., admission secured).	Velocity: $\rightarrow 0$ or negative ATP allocation (M): Step-function collapse	The active-inference engine infers that further academic investment yields no marginal return on the contingency that was sustaining it. System transitions into Conservation Mode.
Gifted Student Burnout	Massive spike in metabolic cost (C) uncoupled from primary evolutionary currencies.	Velocity: Highly negative M : Rapid collapse	The predictive model encounters large prediction errors regarding metabolic expenditure. When cost persistently exceeds utility, biological withdrawal is the adaptive response.
Rote Compliance (Survival Mode)	Utility is maintained via institutional threat, but no competence-expanding activity is authorized.	Velocity: ≈ 0 M : High	High absolute expenditure is authorized to avoid existential prediction errors, but no positive velocity is generated. ATP is sustained by cortisol; exhaustion is the predicted endpoint.
Flow / Epiphany (Growth Mode)	Metabolic cost is lowered through schema assimilation while status/survival utility remains high.	Velocity: Highly positive Acceleration: Positive spike	The sudden acceleration in expected free-energy reduction is associated with the reward signal that sustains high-investment policy; this is the kinematic signature of insight.

Table 3: Classes of System Perturbation: Structural Architecture of EDLT-Informed Interventions

Intervention Class	Target Parameter	Kinematic Goal	General Pedagogical Application
Class I: Metabolic Cost Reduction	Denominator (C)	Lower the ATP threshold required to authorize secondary schema accommodation.	Primary Knowledge Anchoring: Mapping abstract concepts onto evolved cognitive modules (spatial navigation, social coalition dynamics) before introducing formal algorithmic notation, reducing intrinsic cognitive load at the gateway to engagement.
Class II: Utility Recalibration	Numerator (U)	Elevate perceived ROI without relying on distal credential payoffs.	Status-Currency Coupling: Restructuring task outcomes so cognitive mastery generates immediate, peer-visible social value—observable within the salient status ecology rather than deferred to invisible credential systems.
Class III: Threat Relocation	Velocity (dU/dt) & Attractor State	Shift the system from Survival to Growth Mode by altering the ecology of failure.	Active Inference Realignment: Reorganizing assessment to reduce the institutional dominance threat by eliminating high-stakes delays and positioning the teacher as a collaborator against the problem rather than an evaluative authority.
Class IV: Temporal Horizon Compression	Discount rate (k)	Overcome steep temporal discounting in fast-strategy learners.	Feedback Compression: Reducing the latency between effort expenditure and observable utility feedback to timescales compatible with fast-strategy discount functions—hours rather than semesters. Note that feedback compression is a within-system accommodation and does not substitute for structural interventions targeting the environmental instability that produced steep discounting in the first place.