

Adaptive Possibility Dynamics

Context-Conditioned Capability, Partial Observation, and Governed Model Revision

Randy See Meng Teck

CONCEPTUAL AND METHODOLOGICAL ARTICLE | 24 October 2026

Abstract

Adaptive Possibility Dynamics (APD) is an integrative framework for the contextual characterization of human functioning and the conditional appraisal of developmental trajectories. Its unit of analysis is a person undertaking a specified activity under temporally indexed conditions. The framework distinguishes ontic transformation, epistemic updating, and model revision: changes in the person or environment, changes in warranted beliefs, and changes in the representation through which evidence is interpreted. These operations need not coincide. APD combines a typed protocol for interpreting communication, a measurement architecture that preserves heterogeneous outcomes, and a deliberately restricted formal specification based on partially observed dynamics. Developmental possibility is operationalized as a model-relative correspondence between prespecified targets and attainable success probabilities under finite horizons, admissible policies, resource budgets, and explicit constraints. A two-state learning model illustrates the associated requirements for measurement anchoring, identifiability, policy timing, and prospective evaluation. Elementary analytical consequences expose the importance of non-anticipative policies, uncertainty integration, and the distinction between target attainment and sustained admissibility. The computational specification separates evidence, inference, recommendation exposure, and authorization, while the validation program tests incremental predictive value, interpretive fidelity, intervention effects, and participant autonomy separately. APD is presented as a conceptual and methodological proposal. No empirical validation of the integrated framework, fitted individual model, or universal measure of human potential is claimed. Its intended contribution is an inspectable organization of questions whose constituent mechanisms remain independently revisable.

Keywords

adaptive development; contextual capability; partial observability; stochastic reachability; semantic interpretation; measurement invariance; epistemic pluralism; model governance

Significance

Descriptions of individuals often conflate observed performance, inferred capacity, environmental access, and forecasts of future development. APD specifies how these quantities can be investigated without treating them as interchangeable. Its central contribution is a common architecture for preserving the conditions, assumptions, uncertainty, and evaluative commitments attached to each inference. The formal treatment makes developmental forecasts explicitly conditional; the computational treatment preserves the provenance and revisability of the descriptions produced. The research program asks whether this additional structure improves decisions sufficiently to justify its complexity and observational burden.

Analytical examples are hypothetical; no empirical APD results are reported.

Contents

1. Introduction	3
2. Epistemic and semantic architecture	5
3. Measurement, development, and situated capability	7
4. A partially observed developmental specification	10
5. Conditional possibility, reachability, and viability	13
6. Computational architecture and governed revision	14
7. Prospective validation and falsification	16
8. Limitations and research implications	18
9. Conclusion	19
Declarations	20
References	21
Appendix A. Notation, dimensional consistency, and implementation assumptions	24
Appendix B. Analytical consequences and counterexamples	27
Appendix C. Information architecture and revision protocol	30
Appendix D. Prospective evaluation and adversarial acceptance criteria	32
Appendix E. Hypothetical analytical cases	34
Appendix F. Source-use boundaries and evidential roles	37
Appendix G. Terminology and construct boundaries	39

The main text states the conceptual commitments and research specification. Appendices A-B make the formal assumptions and their consequences explicit; C-D specify implementation and evaluation; E supplies hypothetical applications; F-G document source use and terminology.

1. Introduction

1.1. The problem of explanatory granularity

A description of a person is frequently asked to perform several distinct inferential tasks. A trait score may be interpreted as an explanation of an episode, an observed achievement as evidence of an enduring capacity, or an unsuccessful performance as a forecast of a developmental limit. Each transition requires additional assumptions. An instrument can validly summarize a bounded construct while remaining insufficient to determine the mechanism of a particular outcome or the intervention most likely to alter it.

Adaptive Possibility Dynamics addresses this problem through a change in the unit of analysis. The relevant object is a person engaged in an activity under specified conditions at a specified time. Characteristics of the person remain relevant, but their explanatory role is evaluated in conjunction with task demands, temporal history, supports, constraints, interpretations, and the purposes of assessment. This relational unit accommodates both persistent regularities and context-sensitive variation without requiring either to exhaust the description.

The explanatory insufficiency of a person-only descriptor is not resolved by indiscriminately adding dimensions. A larger inventory may preserve the same ambiguities if it fails to distinguish observation from inference, association from intervention, or performance from access. APD therefore proposes an architecture of distinctions and dependencies. Its utility depends on whether that architecture produces more accurate expectations, more discriminating investigations, or more effective and acceptable decisions than simpler alternatives.

The framework draws on established research in personality measurement, within-person variation, learning, embodied and distributed cognition, developmental systems, institutional analysis, and formal decision processes. These traditions supply heterogeneous forms of warrant. Instrument validation supports a particular measurement use; a mathematical construction establishes an implication conditional on assumptions; a philosophical argument clarifies conceptual commitments. Their conjunction does not, by itself, validate APD as an integrated theory.

1.2. Contributions and scope

This article develops five connected specifications. First, it distinguishes transformations of the modeled system from revisions of the evidence and revisions of the model. Second, it provides an interpretation protocol that separates a communicative act, an attributed meaning, and the evidential status of the attributed claim. Third, it formalizes a restricted notion of developmental possibility through partially observed dynamics and constrained reachability. Fourth, it specifies an information architecture in which observations, inferences, permissions, and feedback exposures remain separately traceable. Fifth, it defines a prospective research program capable of rejecting individual components or the integrated workflow for a particular use.

The contribution is integrative and methodological. State-space models, reachability analysis, measurement invariance, and causal identification are established constructions; their appearance here does not establish mathematical priority. The proposed contribution lies in their coordinated application to questions that otherwise risk being compressed into a global personal classification. An application must earn its additional complexity through demonstrated incremental value.

The scope is correspondingly bounded. APD does not provide a diagnostic taxonomy, an exhaustive ontology of mind, or a scalar ordering of persons. It does not derive normative priorities from predictive performance. It does not identify a permanent developmental ceiling from a finite observational record. Its initial empirical domain is a voluntary, narrowly specified, nonclinical activity for which

measurement meaning, permissible actions, temporal ordering, and relevant outcomes can be made explicit.

1.3. Three transformations

The framework distinguishes three operations that can occur independently or in combination.

Ontic transformation concerns a change in the person or environment being studied. A learned procedure, an accessible interface, or a revised institutional rule may alter the conditions of action. The designation does not imply that every such change is directly observable; many remain inferentially established.

Epistemic updating concerns a revision of beliefs within a retained model. New observations can reveal an existing capability, undermine a previous interpretation, or alter a forecast. An assessor's improved estimate is not sufficient evidence of a contemporaneous change in the assessed person.

Model revision concerns the representational system itself: its constructs, measurement relations, transition equations, or admissible interpretations. Replacing a single variable with several differentiated variables may improve explanation without demonstrating the emergence of new entities in the world.

For example, a written checklist can immediately improve successful task execution; later practice may improve unaided proficiency; and subsequent assessment may correct an earlier underestimate. An analyst may additionally replace a single performance construct with separate supported and unaided proficiency constructs. The same observed score trajectory can therefore prompt inquiries into all three transformations. An adequate account identifies which process is supported at each point and preserves uncertainty where the decomposition is unresolved.

1.4. Outcome pluralism

APD distinguishes performance, understanding, lived experience, and opportunity. Performance is success on a defined activity under specified conditions. Understanding concerns discrimination among explanations, appropriate transfer, or warranted prediction. Lived experience concerns the experienced character and significance of a situation. Opportunity concerns activities and trajectories that are practically available under prevailing conditions.

These outcome classes may be mutually supportive, independent, or antagonistic. Supported performance may improve without a corresponding change in unaided capacity. An interpretation may become more meaningful without improving a psychometric score. A demanding intervention may improve one outcome while reducing time, autonomy, or access elsewhere. The assessment problem therefore includes specifying which outcome matters, to whom, and over what interval.

This pluralism does not preclude a decision-maker from selecting a scalar objective for a particular task. It requires that the selection be explicit, locally justified, and revisable. A chosen optimization criterion is an instrument of inquiry or action; its numerical form does not confer universal normative authority.

1.5. Organization

Section 2 develops the epistemic and semantic architecture. Section 3 locates the framework within its measurement and developmental literature. Sections 4 and 5 specify the stochastic model and conditional reachability construction. Section 6 addresses computational implementation and governed revision. Section 7 presents the validation program, followed by limitations and concluding implications. Appendices provide the notation and dimensional ledger, analytical consequences, information schema, adversarial evaluation protocol, hypothetical cases, source-use register, and terminology. The appendices supply the details needed to inspect the argument rather than a separate claim of empirical support.

2. Epistemic and semantic architecture

2.1. Typed warrant and inferential discipline

APD treats the evidential status of a proposition as part of its representation. A statement about experienced meaning, a causal hypothesis, and a normative commitment can concern the same episode while requiring different forms of adjudication. The architecture therefore records the claim type, scope, supporting observations, assumptions, alternatives, and prospective revision conditions.

Table 1. Claim classes are an APD specification for organizing warrant. A passage may contain more than one class and should then be decomposed sufficiently for the relevant inferences.

Claim class	Principal object of appraisal	Inferential restriction
Experiential	The reported character and significance of an experience	The report does not independently identify its external cause.
Empirical	A regularity supported by observations	Its scope is bounded by the sampling, measurement, and study design.
Causal	The consequences of an adequately specified intervention	Association requires additional identification assumptions before receiving an intervention interpretation.
Mathematical	A consequence of explicit definitions and premises	Formal validity does not establish empirical adequacy of the premises.
Normative	A commitment concerning what should be valued, permitted, or chosen	Predictive performance cannot determine the commitment without an additional evaluative argument.
Metaphysical	A proposition about the constitution or ultimate character of reality	Conceptual significance does not imply the existence of a discriminating empirical test.

The purpose of this classification is to prevent unauthorized transfer of warrant. Participant endorsement can support an attribution of intended meaning without validating a proposed external mechanism. A fitted relationship can support a forecast without identifying a causal effect. A proof can establish a property of a specified model without establishing that the model describes a person. A system that loses these distinctions may produce rhetorically coherent explanations whose constituent claims have never been jointly supported.

2.2. Communicative interpretation

The linguistic record is a mediated observation. Grice's analysis of non-natural meaning provides a philosophical resource for examining communicative intention and its recognition [1]. APD uses this resource to distinguish the recorded expression from the interpretation attributed to it. The interpretation may be literal, metaphorical, exploratory, affective, or unresolved; mixed cases should remain representable.

The proposed semantic record contains the original wording, a contextual paraphrase, the attributed type of communicative act, participant confirmation or correction where available, and residual ambiguity. A corrected attribution should propagate to dependent inferences while leaving the original record recoverable. A system that continues to evaluate a disavowed literal interpretation has failed at semantic intake, irrespective of its downstream statistical sophistication.

Participant confirmation has a bounded epistemic role. A speaker has distinctive standing in clarifying what they intended to communicate, while causal explanations and public factual claims remain open to

independent investigation. Conversely, an observer may identify a behavioral regularity while mischaracterizing its personal significance. These sources of evidence can constrain one another without being interchangeable.

Unresolved interpretation is a permissible output. Clarification is especially valuable when competing meanings imply materially different decisions. Where an ambiguity is inconsequential, retaining alternatives may be preferable to imposing additional conversational burden. The operational aim is interpretable communication with proportionate clarification, not the conversion of all informal speech into formal assertions.

2.3. Hierarchical synthesis and the convolutional analogy

The convolutional analogy concerns the organization of inquiry. Classical convolutional architectures use locality, shared operations, and hierarchical composition under assumptions about input structure [2]. Work on geometric deep learning further makes the role of symmetry and inductive bias explicit [3]. Psychological variables do not acquire an appropriate convolutional geometry merely because they can be placed in an array.

APD instead proposes successive analytical passes. Source records preserve observations and their provenance. Reusable local questions identify what is measured, what changes, which constraints apply, and which alternatives remain. Intermediate synthesis extracts candidate relations such as feedback, history dependence, partial observability, and interaction. Cross-domain comparison then asks whether the variables and relations genuinely correspond. The outputs remain differentiated into descriptions, conditional forecasts, competing explanations, options, and proposed investigations.




Figure 1. Proposed information architecture. Arrows indicate analytical and operational dependencies, not measured neural pathways or established causal effects. Subsequent observations and participant corrections can revise earlier interpretations through the governed procedures in Section 6.

An analogy becomes a substantive bridge only when its mapping and limits are explicit. Similar mathematical form does not imply identical mechanism. A learning process and an ecological process may exhibit feedback while differing in their state variables, timescales, interventions, and evidential basis. The synthesis must preserve those differences rather than permit the prestige of one discipline to substitute for evidence in another.

Provenance also constrains aggregation. Multiple paraphrases of one study constitute a shared evidential lineage. Agreement across summaries does not create independent replication. Exceptional or minority cases cannot be discarded solely because they fail to match a pooled representation; they may reveal a measurement defect, a subgroup, or a genuine boundary of the proposed mechanism.

2.4. Perspectival plurality and contestability

A perspective selects and organizes information relative to an explanatory or practical purpose. APD permits first-person, behavioral, interpersonal, and institutional accounts to contribute different information. Their coexistence does not imply that all associated claims are equally well supported. An account remains answerable to the kind of evidence appropriate to its claims.

Scientific perspectivism, social epistemology, and epistemic injustice provide relevant philosophical orientations concerning selective representation, criticism, credibility, and interpretive resources [4, 5, 6]. The present use is limited to those broad orientations; it does not purport to reproduce the books' complete arguments. Their methodological implication for APD is to make challenges to a description possible and consequential without treating either authority or agreement as sufficient evidence.

Contestability operates at several levels. A participant may challenge an observation, the meaning attributed to a statement, the model used to interpret it, or the objective the system pursues. These challenges should not require repudiating every part of the account. A corrected record, a rejected interpretation, and a revised value commitment are distinct and legitimate outcomes of inquiry.

3. Measurement, development, and situated capability

3.1. Construct differentiation and temporal granularity

The framework distinguishes an ability inferred from specified performances, a learned skill, a realized achievement, a momentary state, and a trait summarizing recurrent patterns. Potential denotes a conditional developmental forecast. These working definitions specify different inferential objects; their use does not establish a complete psychological taxonomy.

Dimensional personality instruments can contribute bounded descriptions. The BFI-2 organizes five broad domains and fifteen facets within an empirically developed measurement structure [7]. Experience-sampling work also demonstrates the compatibility of relatively stable dispositional summaries with substantial within-person behavioral variation [8]. APD consequently treats a profile as potentially including a distribution of states and the circumstances associated with variation, subject to the quality of the observations available.

The transport of an association across levels of analysis requires scrutiny. Analyses by Fisher and colleagues illustrate how between-person relationships can fail to reproduce within-person relationships in repeated-measurement datasets [9]. A population tendency therefore supplies a candidate constraint or prior for an individual model; it does not automatically supply an individualized causal law. A model should be evaluated at the level at which its decisions will be made.

Comparability also depends on the measurement relation. Changes across time, groups, or languages may involve noninvariance, including changes in item interpretation or response processes [10]. More frequent measurement does not resolve this problem if the construct changes meaning. Instrument version, administration conditions, units, and calibration assumptions must remain attached to the observations.

Clinical descriptions answer a further, governed set of questions about developmental history, difficulties, differential explanations, and support. Wilson's meta-analysis of Wechsler assessments

reports group-level profile differences in autistic and ADHD samples while cautioning against diagnostic interpretation of those patterns [11]. NICE's adult autism guidance places diagnosis within a comprehensive assessment process that includes individual needs and environmental adjustments [12]. APD's proposed role is supplementary contextual description; it does not authorize diagnosis from a narrative, a group mean, or the equations below.

3.2. Learning trajectories, retention, and transfer

A developmental record must distinguish initial performance, effective practice opportunities, trajectory, retention, transfer, effort, and practical consequences. An observed starting point does not identify an intrinsic learning rate, and a local learning rate does not identify an ultimate ceiling. Absence of a demonstrated ceiling also supplies no evidence that limits are identical or unbounded.

The educational evidence illustrates why these distinctions matter. Ritchie and Tucker-Drob synthesized quasi-experimental designs and estimated effects of additional education on cognitive test performance [13]. Koedinger and colleagues found comparatively similar modeled learning rates alongside heterogeneous initial performance in structured educational practice datasets [14]. These findings motivate investigation of malleability and starting conditions; their scope does not extend automatically to arbitrary training activities or all forms of learning.

Transfer requires its own endpoint. The meta-analysis by Melby-Lervåg and colleagues distinguishes improvements on trained or related working-memory measures from the absence of convincing broad far-transfer benefits in the relevant treated-control comparisons [15]. APD therefore requires separate evidence for trained-task improvement, delayed retention, and performance on prespecified untrained tasks. A gain in one category must not be reported as a gain in all three.

Mutualistic models demonstrate how interactions among developing cognitive processes can generate positive correlations among abilities [16]. The APD extension is a testable interaction claim: a specified change in one component should alter another component's trajectory under declared conditions. A favorable combined outcome is insufficient unless it exceeds an appropriate additive or otherwise prespecified reference. Costs, failed transfer, and asymmetric interactions remain admissible observations.

A practical learning record can document the task, prior knowledge, strategy, prediction, outcome, assistance, duration, and experienced effort. Subsequent retention and transfer checks can test what persisted. Such a record is a research procedure rather than a validated psychometric instrument. Its value is assessed by improved decisions about practice, instruction, tools, or task difficulty, with simpler learning models retained when they perform comparably.

3.3. Embodiment, experience, and the performance arrangement

The assessment interface is part of the conditions under which a performance is generated. Sensory demands, communication format, bodily comfort, available tools, and temporal pressure may influence both action and observation. APD records these conditions instead of assigning every difference to an intrinsic personal property.

The extended-mind argument, enactive approaches, and anticipatory regulation offer distinct theoretical orientations concerning relations among cognition, embodiment, and environment [17, 18, 19]. APD does not require their ontological disputes to be resolved before evaluating an accommodation. Its immediate question is operational: what can be accomplished through a specified arrangement, at what cost, and with what dependence on its components?

Hutchins's analysis of cockpit speed memory provides a concrete case of coordinated cognitive activity across people and representations [20]. In APD, unaided, assisted, and collaboratively distributed

performances must accordingly remain distinguishable. Successful supported activity is an outcome in its own right. Whether that performance persists when a support is removed is a separate question whose relevance depends on the intended use.

First-person reports also contribute evidence without becoming equivalent to external measurements. Lutz and colleagues' small proof-of-concept investigation related structured experiential reports to behavior and brain dynamics [21]; Varela's neurophenomenology supplies a broader methodological proposal for reciprocal constraint [22]. These sources motivate a disciplined relation between reported experience and external observation while leaving their incremental empirical contribution to be evaluated.

Alternative strategies further caution against inferring a global hierarchy from a particular experiential characteristic. Kay and colleagues report slower but more accurate mental-rotation performance in aphantasia, alongside strategy differences in the studied tasks [23]. The APD question concerns task-specific strategies and conditions. A particular imagery profile does not determine general intellectual or philosophical standing.

Accessible communication, sufficient response time, participant review, and appropriate formats are conditions of interpretable observation. Speech, sensory access, language, and understanding cannot be assumed to be interchangeable capacities. An inaccessible assessment may selectively suppress the very evidence needed to understand an individual's functioning.

3.4. Environmental reciprocity and institutional constraint

The environment supplies resources and restrictions while also determining which opportunities are encountered. Absence of demonstrated performance may result from lack of access, unfamiliarity, unsuitable measurement, or a genuine difficulty. The explanatory record must preserve these alternatives until evidence discriminates among them.

Developmental-systems and plasticity accounts emphasize interactions and historical organization at their respective scales [24, 25]. Niche-construction models formalize reciprocal effects between organisms and modified environments in evolutionary settings [26]. The bridge to APD is a structural research question about actions changing the conditions of later action. A claim about an adult learning episode requires its own variables, timescale, and evidence.

Environmental effects can also depend on exposure history. The Moving to Opportunity follow-up analyzes later outcomes from a randomized housing-voucher setting, with effects varying by childhood timing [27]. Ostrom's framework differentiates resource systems, resource units, governance systems, and users, together with their interactions [28]. These contributions support investigating institutional and temporal conditions; neither supplies an individual estimate of innate capacity.

Collective organization introduces additional trade-offs. Derex and Boyd's controlled task provides evidence that communication structure can affect cultural accumulation under particular experimental conditions [29]. APD treats connectivity and coordination as candidate explanatory variables for specified activities. It does not infer an optimal communication structure for all communities.

The resulting contextual profile includes supports, barriers, dependencies, exposure history, and the distribution of control. This last component matters: the person being described may not have the authority or resources to alter the condition that most strongly constrains the outcome.

3.5. Normative pluralism and adaptive regulation

Sen's capability perspective motivates distinguishing resources from the effective opportunities available to people with heterogeneous conversion conditions [30]. APD preserves this distinction while leaving the ordering of valued outcomes explicit. Observed choices can inform a preference model, but choices made under restricted alternatives do not uniquely reveal unconstrained priorities.

Balance is defined relative to a regulatory objective, constraints, and timescale. It need not coincide with an arithmetic midpoint or a return to an earlier state. Holling's distinction between resilience and stability supplies a relevant ecological vocabulary [31]; the persistence of a configuration is nevertheless separate from its ethical desirability. An arrangement may endure while imposing unacceptable costs.

The normative specification therefore separates continuability, attainable activity, and the distribution of benefit and burden. It includes rights to contest descriptions, decline measurement, change objectives, and retain support without establishing an exceptional productivity forecast. These are design commitments. Their legitimacy is not derived from a predictive model's accuracy.

Cross-cultural comparison requires particular care. Vignoles and colleagues provide empirical grounds for a differentiated account of selfhood beyond a single independence-interdependence contrast [32]. APD permits locally meaningful categories while requiring evidence for the comparability of any numerical or categorical ordering applied across contexts.

4. A partially observed developmental specification

4.1. States, contexts, observations, and models

Let x_t denote a finite-dimensional, task-relevant latent state, e_t an observed context vector, u_t an action or recorded exposure, and o_t an observation at discrete time t . A model M , with fixed but uncertain parameters θ , specifies transition and observation laws:

$$(x_{t+1}, e_{t+1}) \sim P_{\theta}^M(\cdot | x_t, e_t, u_t), \quad o_t \sim O_{\theta}^M(\cdot | x_t, e_t). \quad (1)$$

Equation (1) imposes conditional state sufficiency. The selected state must summarize enough history for the transition law to be adequate within the intended domain. If earlier exposures have effects not captured by the state, relevant history must be incorporated or the specification revised. Current context may be observed, but its future evolution still requires an explicit law or scenario. Uncertain context measurement requires a corresponding observation model.

Let H_t contain the observations and observed context available before the decision at time t , earlier actions, instrument versions, and other legitimately available records. The inferential state is a distribution over the latent state and uncertain parameters:

$$b_t^M(dx, d\theta) = \Pr^M(x_t \in dx, \theta \in d\theta | H_t). \quad (2)$$

The person, the context, the posterior distribution, and the model are distinct objects. An update of b_t^M within a retained M does not establish an ontic change. Replacing M can change the posterior representation without changing the underlying record. Parameter uncertainty in this construction is uncertainty within a model class; uncertainty about omitted mechanisms is not automatically included.

Partially observable decision models supply the formal language for this separation [33]. Their use is representational. APD does not thereby assert that human cognition computes Bayesian posteriors, maintains these variables, or maximizes a single objective.

4.2. A restricted two-state instantiation

Consider one task family observed in daily intervals over an illustrative eight-week protocol. Morning measurements precede the day's practice. Let s_t be proficiency on a fixed item-response log-odds scale and z_t a provisional readiness construct on a separately fixed standardized scale. Let u_t be hours of practice undertaken during the interval and r_t recovery-exposure hours relative to a protocol-defined reference. The local dynamics are

$$s_{t+1} = (1 - \delta)s_t + \delta\mu_s + \alpha u_t + \beta u_t z_t + w_t^s, \quad (3)$$

$$z_{t+1} = \rho z_t + \eta r_t - \lambda u_t + w_t^z, \quad (4)$$

$$(w_t^s, w_t^z)^\top \sim \mathcal{N}(0, Q). \quad (5)$$

The parameter μ_s is a reference proficiency. It is neither the initial observation nor a developmental upper bound. The coefficient α describes the local practice contribution, while β describes its interaction with the readiness state. The readiness equation permits persistence and associations with recovery exposure and practice. Their substantive interpretation depends on calibration and design; the sign of an estimated coefficient is not established by naming its exposure.

For the mean-reverting variant, assume $0 < \delta < 1$ and $|\rho| < 1$. The alternative $\delta = 0$ removes μ_s from the model, making that parameter unidentifiable within the resulting specification. These are daily parameters and cannot be carried unchanged into another temporal resolution. The readiness variable is a local provisional construct, not a diagnosis or an exhaustive description of bodily functioning.

The operational definition of recovery exposure must be fixed before collection. The model does not prescribe sleep restriction or treat all low-demand activities as equivalent. Intervention claims concerning physiological or clinical outcomes are outside the proposed initial application.

4.3. Measurement relations

For calibrated item j , a binary correct-response variable Y_{jt} is modeled by

$$Y_{jt} \mid s_t, z_t, e_t \sim \text{Bernoulli}[\sigma(s_t + \gamma z_t - d_j + \ell^\top e_t)], \quad (6)$$

$$\sigma(a) = \frac{1}{1 + \exp(-a)}. \quad (7)$$

Here d_j is item difficulty and ℓ converts declared context units into log-odds contributions. The item scale requires an explicit anchoring scheme. Simultaneous shifts in proficiency, reference proficiency, and item difficulty can preserve the complete observation law; Appendix B gives the corresponding invariance argument.

A separate readiness indicator is specified as

$$R_t = z_t + \varepsilon_t, \quad \varepsilon_t \sim \mathcal{N}(0, \tau_R^2). \quad (8)$$

This assumes a continuous, approximately Gaussian indicator on the declared scale. An ordinal response requires an ordinal observation model. Reaction time requires its own likelihood. Performance, latency, and experienced effort must not be silently treated as measurements of one interchangeable quantity.

The equations specify response marginals; a fitted joint likelihood additionally requires assumptions about within-session item dependence, indicator serial dependence, and relationships among observation errors. Process innovations are provisionally independent across days and independent of indicator errors. Filtering or forecasting under the declared conditional transition requires that the innovation law

remain appropriate conditional on the decision history and inputs. Endogenous exposures can instead require an explicit joint exposure-innovation model. An intervention interpretation additionally requires the relevant causal assumptions. These requirements must be specified and examined rather than inferred from the presence of a Gaussian symbol.

4.4. Identifiability and experimental excitation

Structural identifiability concerns distinguishability in principle under a model and observation design. Practical identifiability concerns distinguishability with the finite, noisy record actually available. Profile-likelihood analysis provides one established route for examining such ambiguities [34]. A good fit is compatible with several mechanisms when they induce similar observation laws.

An independent indicator can help anchor z_t , but it does not guarantee identification. If practice is constant, the proficiency intercept may identify only a combination of μ_s and α . If recovery exposure is proportional to practice, the readiness equation may identify only a combination of η and λ . Nearly constant readiness can make the practice contribution and its interaction difficult to distinguish. Appendix B separates exact invariances from finite-data difficulties.

Before empirical interpretation, the proposed estimation procedure should undergo parameter-recovery analysis using data generated from known parameters under the intended sampling and missingness design. Failure to recover a parameter limits the claims that may be attached to it. A more complex latent representation is not justified solely by its capacity to fit the development sample.

4.5. Timing, implementation, and causal status

The information available at the beginning of a day cannot include recovery exposure or other measurements recorded only at its end. Forecasts must integrate a declared law for such future quantities or explicitly condition on a hypothetical scenario. Retrospective availability in a dataset does not make a variable prospectively available to a policy.

Completed practice and assigned practice also require separation. Equations (3)-(4) describe recorded practice undertaken. Interpreting u_t as a directly controllable policy action requires an additional implementation assumption. If a prescribed dose is not executed as assigned, assignment and realized exposure are distinct variables and require an implementation or compliance model. This paper supplies no estimated compliance law.

The causal distinction is therefore substantive:

$$\Pr(Y_{t+1} \mid u_t = u, H_t) \not\equiv \Pr(Y_{t+1} \mid \text{do}(u_t = u), H_t). \quad (9)$$

The notation denotes the absence of a general identity; the quantities can coincide under suitable assumptions. Prior success, opportunity, motivation, or unobserved conditions can influence both practice and subsequent performance. Causal identification requires an appropriate design or defensible structural assumptions [35]. Optimization of an observational forecast does not supply those assumptions.

5. Conditional possibility, reachability, and viability

5.1. A target-relative possibility correspondence

Developmental possibility is operationalized here through a restricted reach-avoid problem [36]. Let G_g be a target region indexed by g in a prespecified meaningful family $\mathcal{G}$. Let K be an operationally acceptable region. These sets may refer to an augmented state containing context and other necessary bookkeeping variables. In a continuous state space, target regions are generally more appropriate than exact target points, whose hitting probabilities may be zero under the specified law.

Let D_0 denote the evidence available at planning time, T a finite discrete horizon, and B a vector of budgets. Hours, money, and other resources retain their separate units. An admissible policy uses only available information and respects permissions and resource restrictions on permitted trajectories. The admissible class $\Pi_{\text{adm}}(B)$ must be nonempty for the success-probability interpretation below; otherwise the system reports infeasibility.

Define target-entry and constraint-exit times by

$$\tau_G = \inf \{t \geq 0 : x_t \in G\}, \quad \tau_{K^c} = \inf \{t \geq 0 : x_t \notin K\}, \quad (10)$$

with the infimum of an empty set equal to infinity. In this section x_t is understood to include the state augmentation required by the sets. The conditional value is

$$V_{G,K}^M(D_0, T, B) = \sup_{\pi \in \Pi_{\text{adm}}(B)} \int \Pr_{\theta, x_0}^{M, \pi}(\tau_G \leq T, \tau_{K^c} > T) b_0^M(dx_0, d\theta). \quad (11)$$

For a chosen success threshold $p \in (0, 1]$, define

$$\widehat{\mathcal{R}}_p^M(D_0, T, B, K) = \{g \in \mathcal{G} : V_{G_g, K}^M(D_0, T, B) \geq p\}. \quad (12)$$

The target set is conditional on the model, evidence, horizon, budget, constraints, and policy class. The threshold is a success-probability requirement, not a confidence level, a psychometric score, or a measure of intrinsic human worth. No canonical geometry, cardinality-based value, or volume measure is assigned to this set.

5.2. Information admissibility and uncertainty integration

The supremum in equation (11) is outside the uncertainty integral. A single information-admissible rule must operate despite uncertainty about the initial state and parameters. The rule may adapt to observations during execution; it cannot select its actions using unknown quantities as though they had been observed.

Moving the supremum inside the integral supplies a more informative, potentially clairvoyant benchmark. Appendix B establishes the associated inequality and gives a finite example in which a legitimate policy succeeds with probability one half while parameter-contingent selection succeeds with probability one. The distinction is material whenever different latent configurations favor different actions.

The averaged probability integrates uncertainty represented in b_0^M . It does not guarantee protection against misspecified dynamics, unmodeled costs, or future contexts absent from the scenario. Comparing plausible models can reveal sensitivity, but it cannot convert unknown omissions into quantified uncertainty automatically.

5.3. Full-horizon admissibility and viability

The event in equation (11) requires remaining within K through the entire horizon, including after target attainment. A protocol that terminates at success has a different event and must be specified explicitly. The distinction has an immediate consequence: increasing the horizon need not increase the value. Additional time may expose a previously successful trajectory to a constraint violation. Appendix B supplies a deterministic counterexample.

Viability concerns the existence of an admissible course that maintains specified constraints; reachability concerns attaining a target [37]. A course can satisfy one criterion and fail the other. In personal applications, the acceptability region must have an operational interpretation and cannot be presented as a clinically validated boundary without appropriate evidence.

Authorization is enforced separately as a hard restriction on permissible collection and action. Including burden or other conditions in a stochastic acceptability calculation does not authorize a probability of violating consent. Revocation changes what the system may do, and policies must respect that change at the relevant decision point.

5.4. Resource comparisons and value pluralism

Under a fixed model, posterior, horizon, and context law, enlarging budgets cannot reduce the calculated value when the feasible policy classes are nested. Raising the success threshold cannot enlarge the corresponding target set. These are elementary properties of the definition, not empirical claims that additional resources necessarily improve experience or that stricter criteria improve lives.

Multiple objectives can be retained without premature aggregation. Pareto analysis identifies alternatives not dominated across the selected objectives, relative to the specified feasible set [38]. It does not choose an ethically privileged compromise. An application must expose how learning, autonomy, enjoyment, income, burden, or other selected outcomes enter the comparison.

Low modeled reachability can reflect a restricted budget, an unsuitable action menu, measurement failure, or model misspecification. A finite calculation cannot establish a lifelong limit. Conversely, a favorable forecast requires empirical support in the conditions of use; mathematical permissibility alone does not establish practical availability.

5.5. Interpretation of the analytical results

The analytical consequences in Appendix B are expository implications of the specification. They clarify why calibration anchors are necessary, why policy information must be respected, why uncertainty cannot be optimized away by rearranging operators, and why a changed estimate need not indicate changed capacity. They supply logical checks on implementation. They do not validate the psychological constructs or establish the empirical superiority of APD.

6. Computational architecture and governed revision

6.1. A provenance-preserving research system

An initial implementation should function as a bounded personal research system. Its purposes are to retain observations, compare explanations, forecast specified outcomes, and register failures. It need not begin with a comprehensive computational theory of a person or a continuously retrained neural architecture.

The proposed pipeline maintains measurement-specific representations for structured assessments, event histories, contextual observations, reflective language, and relational information. An observed zero, an

unanswered item, an unavailable instrument, and a deliberate refusal have different meanings. Elapsed time, instrument version, assistance conditions, and observation provenance must remain available to downstream inference.

Methods such as GRU-D illustrate how elapsed time and missingness can enter a predictive model [39]. Graph networks provide a language for explicit relational structure [40], and Temporal Fusion Transformers combine several classes of temporal and static inputs [41]. These are candidate components. A learned edge is not thereby a causal relation; an attention weight is not thereby a psychological motive; a missingness feature does not establish why a person declined a record.

The minimum implementation can use an access-controlled local database, structured forms, a hypothesis register, and a fixed analysis procedure. A simple dynamic regression or small state-space model should precede more complex encoders. An optional language model may generate candidate summaries or hypotheses, but generated prose must remain identified as an inference rather than a new observation.

6.2. Evidence, events, inferences, and exposure

The schema distinguishes records by their role. Observations retain source identity, event time, recording time, measurement version, units, and correction history. Events distinguish directly observed, reported, and inferred occurrence. Model records identify specifications, parameters, training evidence, evaluation, approval, and release status. Inference records preserve their model version, inputs, target time, uncertainty, and supersession relations.

Recommendation exposure requires its own record. Displaying a forecast, accepting it, acting on it, and observing an outcome are separate events. If a forecast influences behavior, later evidence cannot be interpreted as though the forecast were absent. The theory of performative prediction formalizes this general feedback problem [42]. APD therefore distinguishes predeployment predictive accuracy from postdeployment effects on behavior and outcomes.

Appendix C provides the record-level specification. References point to controlled content rather than reproducing sensitive material in every table. The initial notebook may collect only a subset of fields, but it must retain the conceptual distinctions needed to interpret its evidence.

6.3. Three update channels

Knowledge revision adds, qualifies, or retracts sourced claims. Each claim retains evidence type, scope, assumptions, contradictions, and provenance. A publication may enter the source register without automatically becoming a recommendation. Corrections should propagate to dependent inferences while preserving the earlier reasoning trail.

Model updates change parameter estimates or the model specification. This engineering channel includes both epistemic updating within a retained specification and revisions of equations, features, or representation; the record must distinguish them. A challenger is evaluated against the approved model for its intended use. Documentation should record training evidence, prospective evaluation, limitations, and release conditions. Model cards supply a reporting pattern [43]; documentation contributes to scrutiny but does not independently establish validity.

Permission revision changes what the system may collect, access, infer, display, or execute. Authorization is separate from learning. Neither a model update nor a retrieved document can grant new permissions, redefine prohibited uses, or silently replace a person's objectives. Goals therefore also require explicit versions and a record of who authorized the change.

This separation prevents an adaptive system from treating its increasing capability as an expansion of its authority. Changes in knowledge, model, and permission may interact operationally, but their legitimacy is established through different procedures.

6.4. Temporal integrity and bounded deployment

A late record about an earlier event changes what is known now. It does not establish that the event occurred at the recording time. A corrected measurement may invalidate an inference without indicating a new developmental change. Append-only revisions and supersession links preserve this distinction.

Deployment also requires bounded revision. Continual-learning methods such as elastic weight consolidation mitigate forgetting in studied settings [44], but they do not guarantee improvement under indefinite autonomous modification. A small APD system may avoid continual neural retraining entirely and instead evaluate periodic statistical revisions.

Releases should retain regression cases, genuinely unexposed prospective evaluation, and tested rollback. A repeatedly consulted test set becomes part of model development through repeated decisions, even if it is never directly used for gradient updates. Uncertainty, subgroup or context performance, and participant burden should be examined alongside average predictive accuracy.

Retrieved documents are evidence-bearing data, not authorization-bearing instructions. The system must test whether misleading content can corrupt provenance or alter permissions. NIST's AI Risk Management Framework supplies lifecycle governance guidance [45]; adopting its vocabulary or procedures is not a certificate of clinical effectiveness or legal compliance.

6.5. Consent, control, and the option of nonoptimization

Collection is limited to the agreed purpose. The system should not infer intimate attributes of nonconsenting third parties merely because a participant refers to them. Correction, export, retention, revocation, and deletion behavior must be explicit. If removal of training records cannot reliably remove their influence from a model, the remaining limitation must be disclosed.

Prediction targets remain distinct from life objectives. Increased task completion may coincide with diminished rest, exploration, or relational quality. A participant must be able to reject a recommendation, revise the objective, pause collection, or prefer a simpler notebook without punitive consequences. Agency is represented here through rights and procedures; no scalar agency measure is introduced.

7. Prospective validation and falsification

7.1. Distinct validation targets

The validation program separates interpretive fidelity, measurement adequacy, predictive performance, causal identification, decision usefulness, and ethical acceptability. Improvement on one target does not establish improvement on the others. A system can paraphrase accurately while forecasting poorly, predict accurately while changing behavior adversely, or increase task performance while imposing unacceptable observation burden.

Table 2. Stages are gates for particular claims, not a universal sequence whose completion validates every part of APD.

Stage	Primary question	Required comparison or evidence	Progression condition
Meaning and feasibility	Is the record understandable, correctable, and manageable?	Participant review, ordinary skilled interviewing, correction burden, and missingness reasons where volunteered	The process produces interpretable records at an acceptable burden.

Stage	Primary question	Required comparison or evidence	Progression condition
Measurement	Do observations retain comparable meaning across the intended conditions?	Calibration, administration checks, reliability and invariance analysis appropriate to the instrument	Changes can be interpreted within declared limits.
Prediction	Do additional APD terms improve future forecasts?	Person mean, recent-performance or autoregressive model, common learning curve, and task-difficulty baselines	Prespecified prospective gain justifies additional complexity.
Intervention	Does a controlled change alter the intended outcome?	Appropriate randomization or identified alternative design, with carryover and adherence considered	The estimated effect supports the bounded intervention claim.
Decision and experience	Does displaying or using the system improve decisions acceptably?	Exposed versus unexposed conditions where appropriate; burden, autonomy, and meaningful outcomes	Benefits persist after accounting for feedback and participant costs.

7.2. Baselines, prospective evaluation, and calibration

Each operational study should prespecify its primary outcome, reference models, observation windows, exclusion rules, and a meaningful improvement criterion. Sample size and duration depend on variability, dependence, missingness, effect size, and required precision. Dense repeated measurements do not confer the information of equally numerous independent observations.

Evaluation should proceed forward in time. Randomly splitting overlapping diary windows can leak temporal information into the evaluation set. Later-time holdouts for familiar participants, held-out participants, and held-out contexts address different generalization questions and should be reported separately.

Appropriate probabilistic scores, interval coverage, and context-specific errors should accompany point-prediction metrics. Research on neural calibration and predictive uncertainty under distribution shift motivates explicit checks of confidence behavior [46, 47]. Numerical certainty must not be interpreted as calibrated uncertainty without evaluation in the relevant conditions.

Component ablations test whether additional structure earns its place. Remove contextual variables, person-specific parameters, interaction terms, missingness features, or complex encoders while preserving a fair comparison. If a simpler account predicts comparably and reduces burden, it should remain the operational model. The objective is useful discrimination among models, not confirmation of complexity for its own sake.

7.3. Interpretation and interaction studies

The semantic intake procedure can be compared with literal-only classification and ordinary skilled interviewing. Outcomes include participant correction burden, factual fidelity, unresolved ambiguity, and the consequences of corrected interpretations for downstream decisions. Agreement with a paraphrase is evidence about the attribution of meaning, not an overall truth score for the framework.

A proposed learning interaction requires a factorial or otherwise suitable design. For two trainable components, compare neither, each alone, and both, while accounting for total training time and engagement. Assess a prespecified untrained task and test the relevant interaction. The highest combined score alone does not establish synergy. Durable learning also makes simple crossover designs vulnerable to carryover.

The intervention endpoint must match the claim. Improved performance while a checklist remains available is evidence about supported functioning. Improved unaided performance after a delay addresses retention or learning. Transfer to a different task addresses another question. Conflating these endpoints would obscure both useful supports and the limits of training.

7.4. Feedback, autonomy, and consequential use

Optional nonclinical forecasts or suggestions can be displayed under randomized or otherwise identified conditions with informed consent. Exposure, acceptance, action, outcome, burden, and perceived autonomy must be recorded separately. An intervention can change the distribution on which its own forecasting system is later evaluated.

The integrated workflow also requires a practical comparison. The relevant competitor may be a simpler notebook, a conventional interview, or an existing planning aid. A framework that adds explanatory vocabulary without improving decisions or understanding has not demonstrated its intended contribution.

Low-risk voluntary changes, such as instruction format or access to a planning aid, are appropriate initial domains. Medication changes, sleep deprivation, deliberate distress, and clinical treatment are outside this initial notebook proposal. This boundary specifies the proposed study rather than a judgment about the legitimacy of properly governed research elsewhere.

7.5. Failure conditions and corrective responses

Appendix D specifies eight adversarial acceptance tests covering duplicate evidence, retrospective corrections, context shift, complexity, feedback exposure, hostile source content, revocation and rollback, and rejection of the system's objective. Each test has an expected behavior and a failure response.

Failure can require widening uncertainty, suspending a feature, reverting to a simpler model, collecting discriminating observations, or rejecting a mechanism for the tested domain. Adding an explanation after every contrary result without changing predictions would undermine falsifiability. The framework must permit an unsuccessful component to be removed without interpreting the failure as a defect of the participant.

Negative findings, failed replications, unsuccessful modules, and revised scope are substantive research outputs. They identify which organizational commitments are useful and which proposed bridges do not survive examination.

8. Limitations and research implications

8.1. Integration is not empirical unification

APD coordinates heterogeneous traditions around a set of practical distinctions. It has not established a single generative mechanism underlying those traditions. A shared term can conceal different constructs, and a common formalism can conceal different intervention meanings. Cross-domain synthesis therefore remains conditional on explicit mappings and evidence for their use.

The present literature selection is purposive and interdisciplinary. It is not a systematic review, a meta-analysis of APD, or an exhaustive survey of adjacent work. Original empirical reports, original meta-analytic analyses, formal methods, theoretical syntheses, philosophical arguments, and official guidance are identified according to their different roles. Appendix F records the boundaries of source use, including works consulted at abstract, publisher, or preview level.

8.2. Physics and information as bounded formal resources

Terms drawn from other domains require operational commitments. An attractor belongs to specified dynamics, and a bifurcation denotes a defined qualitative change under parameter variation. Work on early-warning detection demonstrates that even explicitly modeled transitions can produce false positives and missed detections [48]. A reported insight does not acquire a dynamical explanation merely by being called a transition.

Shannon entropy is defined relative to a distribution and communication problem [49]. It is not a general measure of meaning, wisdom, or moral value. Quantum-cognition proposals concern the use of quantum probability formalisms for selected cognitive phenomena [50]; they do not establish a physical quantum mechanism for the brain. Any APD use would require discriminating predictions and comparisons with appropriate classical alternatives.

Active inference remains a candidate computational approach rather than a required explanatory commitment [51]. Schrödinger's 1935 argument supplies historical physical context for the cat thought experiment [52]. The practice of preserving alternative interpretations can be evaluated on its own terms without assigning those interpretations physical superposition.

8.3. Open problems and limitations

Conditional capability research asks which outcomes are attainable under stated supports, resources, constraints, and horizons. It must distinguish physical possibility, statistical support, practical access, and personal desirability. These categories may overlap, but none can stand in for all the others.

Semantic development research asks how acquiring a distinction, changing an interpretation, or finding a usable vocabulary affects communication and action. Vocabulary change, conceptual discrimination, behavior, and reporting effects require separate observations. A useful study could assess unfamiliar applications of a taught distinction and its contribution to prediction or communication.

Opportunity and feedback research examines how tools, institutions, recommendation systems, and social arrangements alter both activities and the data subsequently used to describe people. Questions of empirical effect coexist with questions about who can inspect, contest, or leave the arrangement.

Representation research asks whether a new dimension, a new interaction, or a different model form improves the intended inference. A statistical factor does not automatically identify a new kind of mental substance. Additional constructs require interpretable measurement, prospective utility, and examination of artifacts. Sometimes a sequence, network, or task-specific process model will serve better than an expanded trait inventory.

These questions remain open within the APD specification. The framework does not resolve consciousness, free will, the interpretation of quantum mechanics, or the complete structure of human potential. Its contribution is to identify which part of a question is experiential, empirical, mathematical, normative, or metaphysical, and what progress in that part would require.

9. Conclusion

Adaptive Possibility Dynamics proposes a disciplined relation among situated description, conditional forecasting, interpretive correction, and governed revision. Its central commitments are the person-activity-context-time unit, the separation of ontic and epistemic change, the preservation of heterogeneous outcomes and warrants, and the explicit conditionality of developmental forecasts.

The formal specification makes these commitments inspectable through a restricted latent-state model and a constrained target-attainment correspondence. The computational architecture preserves the

evidential and authorization history of the inferences produced. The validation program asks whether these distinctions improve observation, prediction, intervention, and participant-directed decisions sufficiently to justify their cost.

The framework remains revisable at the level of individual mechanisms, representations, and uses. A successful investigation may establish a narrow improvement in interpretation, a context-dependent support effect, or a useful learning interaction. Such a result can justify a component while leaving the broader framework open. APD's scientific development depends on allowing that asymmetry between local evidence and global ambition.

Declarations

Research status. This article is a conceptual and methodological research proposal. It reports no newly collected human-participant dataset, fitted APD parameters, clinical validation, or evaluation of an implemented APD system. Hypothetical cases and analytical counterexamples are identified as such. The propositions explicate consequences of the stated definitions and assumptions; no mathematical priority claim is made.

Materials and computational status. The definitions, model equations, analytical arguments, record schema, and proposed evaluation procedures needed to inspect the specification are included in the main text and appendices. No operational APD software, trained model, or empirical analysis dataset is represented as an accompanying validated research product.

Acknowledgment of AI assistance. ChatGPT assisted with drafting, literature organization, formal checking, and document preparation. Responsibility for the manuscript's claims and interpretations remains with the named author.

References

- [1] Grice, H. P. (1957). Meaning. *The Philosophical Review*, 66(3), 377-388. doi:10.2307/2182440.
- [2] LeCun, Yann; Bottou, Léon; Bengio, Yoshua; Haffner, Patrick (1998). Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11), 2278-2324. doi:10.1109/5.726791.
- [3] Bronstein, Michael M.; Bruna, Joan; Cohen, Taco; Veličković, Petar (2021). Geometric Deep Learning: Grids, Groups, Graphs, Geodesics, and Gauges. *arXiv*, arXiv:2104.13478; version 2, 2 May 2021. Original source record.
- [4] Giere, Ronald N. (2006). Scientific Perspectivism. *University of Chicago Press*. doi:10.7208/chicago/9780226292144.001.0001.
- [5] Longino, Helen E. (2002). The Fate of Knowledge. *Princeton University Press*. doi:10.1515/9780691187013.
- [6] Fricker, Miranda (2007). Epistemic Injustice: Power and the Ethics of Knowing. *Oxford University Press*. doi:10.1093/acprof:oso/9780198237907.001.0001.
- [7] Soto, Christopher J.; John, Oliver P. (2017). The next Big Five Inventory (BFI-2): Developing and assessing a hierarchical model with 15 facets to enhance bandwidth, fidelity, and predictive power. *Journal of Personality and Social Psychology*, 113(1), 117-143. doi:10.1037/pspp0000096.
- [8] Fleeson, William (2001). Toward a structure- and process-integrated view of personality: Traits as density distributions of states. *Journal of Personality and Social Psychology*, 80(6), 1011-1027. doi:10.1037/0022-3514.80.6.1011.
- [9] Fisher, Aaron J.; Medaglia, John D.; Jeronimus, Bertus F. (2018). Lack of group-to-individual generalizability is a threat to human subjects research. *Proceedings of the National Academy of Sciences*, 115(27), E6106-E6115. doi:10.1073/pnas.1711978115.
- [10] Putnick, Diane L.; Bornstein, Marc H. (2016). Measurement invariance conventions and reporting: The state of the art and future directions for psychological research. *Developmental Review*, 41, 71-90. doi:10.1016/j.dr.2016.06.004.
- [11] Wilson, Alexander C. (2024). Cognitive Profile in Autism and ADHD: A Meta-Analysis of Performance on the WAIS-IV and WISC-V. *Archives of Clinical Neuropsychology*, 39(4), 498-515. doi:10.1093/arclin/acad073. Published online 29 September 2023; cited by the 2024 issue year.
- [12] National Institute for Health and Care Excellence (2012). Autism spectrum disorder in adults: diagnosis and management. *NICE Clinical Guideline*, CG142; updated 14 June 2021. Original source record.
- [13] Ritchie, Stuart J.; Tucker-Drob, Elliot M. (2018). How Much Does Education Improve Intelligence? A Meta-Analysis. *Psychological Science*, 29(8), 1358-1369. doi:10.1177/0956797618774253.
- [14] Koedinger, Kenneth R.; Carvalho, Paulo F.; Liu, Ran; McLaughlin, Elizabeth A. (2023). An astonishing regularity in student learning rate. *Proceedings of the National Academy of Sciences*, 120(13), e2221311120. doi:10.1073/pnas.2221311120.
- [15] Melby-Lervåg, Monica; Redick, Thomas S.; Hulme, Charles (2016). Working Memory Training Does Not Improve Performance on Measures of Intelligence or Other Measures of “Far Transfer”: Evidence From a Meta-Analytic Review. *Perspectives on Psychological Science*, 11(4), 512-534. doi:10.1177/1745691616635612.
- [16] van der Maas, Han L. J.; Dolan, Conor V.; Grasman, Raoul P. P. P.; Wicherts, Jelte M.; Huizenga, Hilde M.; Raijmakers, Maartje E. J. (2006). A dynamical model of general intelligence: The positive manifold of intelligence by mutualism. *Psychological Review*, 113(4), 842-861. doi:10.1037/0033-295X.113.4.842.
- [17] Clark, A.; Chalmers, D. (1998). The Extended Mind. *Analysis*, 58(1), 7-19. doi:10.1093/analys/58.1.7.
- [18] Sterling, Peter (2012). Allostasis: A model of predictive regulation. *Physiology & Behavior*, 106(1), 5-15. doi:10.1016/j.physbeh.2011.06.004.
- [19] Varela, Francisco J.; Thompson, Evan; Rosch, Eleanor (2017). The Embodied Mind: Cognitive Science and Human Experience. *MIT Press*, Revised edition; original edition 1991; ISBN 9780262529365. Original source record.
- [20] Hutchins, Edwin (1995). How a Cockpit Remembers Its Speeds. *Cognitive Science*, 19(3), 265-288. doi:10.1207/s15516709cog1903_1.

- [21] Lutz, Antoine; Lachaux, Jean-Philippe; Martinerie, Jacques; Varela, Francisco J. (2002). Guiding the study of brain dynamics by using first-person data: Synchrony patterns correlate with ongoing conscious states during a simple visual task. *Proceedings of the National Academy of Sciences*, 99(3), 1586-1591. doi:10.1073/pnas.032658199.
- [22] Varela, Francisco J. (1996). Neurophenomenology: A Methodological Remedy for the Hard Problem. *Journal of Consciousness Studies*, 3(4), 330-349. Original source record.
- [23] Kay, Lachlan; Keogh, Rebecca; Pearson, Joel (2024). Slower but more accurate mental rotation performance in aphantasia linked to differences in cognitive strategies. *Consciousness and Cognition*, 121, article 103694. doi:10.1016/j.concog.2024.103694.
- [24] Smith, Linda B.; Thelen, Esther (2003). Development as a dynamic system. *Trends in Cognitive Sciences*, 7(8), 343-348. doi:10.1016/S1364-6613(03)00156-6.
- [25] West-Eberhard, Mary Jane (2005). Developmental plasticity and the origin of species differences. *Proceedings of the National Academy of Sciences*, 102(Suppl. 1), 6543-6549. doi:10.1073/pnas.0501844102.
- [26] Laland, K. N.; Odling-Smee, F. J.; Feldman, M. W. (1999). Evolutionary consequences of niche construction and their implications for ecology. *Proceedings of the National Academy of Sciences*, 96(18), 10242-10247. doi:10.1073/pnas.96.18.10242.
- [27] Chetty, Raj; Hendren, Nathaniel; Katz, Lawrence F. (2016). The Effects of Exposure to Better Neighborhoods on Children: New Evidence from the Moving to Opportunity Experiment. *American Economic Review*, 106(4), 855-902. doi:10.1257/aer.20150572.
- [28] Ostrom, Elinor (2009). A General Framework for Analyzing Sustainability of Social-Ecological Systems. *Science*, 325(5939), 419-422. doi:10.1126/science.1172133.
- [29] Derex, Maxime; Boyd, Robert (2016). Partial connectivity increases cultural accumulation within groups. *Proceedings of the National Academy of Sciences*, 113(11), 2982-2987. doi:10.1073/pnas.1518798113.
- [30] Sen, Amartya (1980). Equality of What?. *The Tanner Lectures on Human Values, vol. 1, edited by Sterling M. McMurrin*, 197-220; lecture delivered at Stanford University, 22 May 1979. Original source record.
- [31] Holling, C. S. (1973). Resilience and Stability of Ecological Systems. *Annual Review of Ecology and Systematics*, 4, 1-23. doi:10.1146/annurev.es.04.110173.000245.
- [32] Vignoles, Vivian L.; Owe, Ellinor; Becker, Maja; et al. (2016). Beyond the “East-West” dichotomy: Global variation in cultural models of selfhood. *Journal of Experimental Psychology: General*, 145(8), 966-1000. doi:10.1037/xge0000175.
- [33] Kaelbling, Leslie Pack; Littman, Michael L.; Cassandra, Anthony R. (1998). Planning and acting in partially observable stochastic domains. *Artificial Intelligence*, 101(1-2), 99-134. doi:10.1016/S0004-3702(98)00023-X.
- [34] Raue, A.; Kreutz, C.; Maiwald, T.; et al. (2009). Structural and practical identifiability analysis of partially observed dynamical models by exploiting the profile likelihood. *Bioinformatics*, 25(15), 1923-1929. doi:10.1093/bioinformatics/btp358.
- [35] Pearl, Judea (2009). Causal inference in statistics: An overview. *Statistics Surveys*, 3, 96-146. doi:10.1214/09-SS057.
- [36] Abate, Alessandro; Prandini, Maria; Lygeros, John; Sastry, Shankar (2008). Probabilistic reachability and safety for controlled discrete time stochastic hybrid systems. *Automatica*, 44(11), 2724-2734. doi:10.1016/j.automatica.2008.03.027.
- [37] Aubin, Jean-Pierre; Bayen, Alexandre M.; Saint-Pierre, Patrick (2011). Viability Theory: New Directions. *Springer*. doi:10.1007/978-3-642-16684-6.
- [38] Deb, K.; Pratap, A.; Agarwal, S.; Meyarivan, T. (2002). A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation*, 6(2), 182-197. doi:10.1109/4235.996017.
- [39] Che, Zhengping; Purushotham, Sanjay; Cho, Kyunghyun; Sontag, David; Liu, Yan (2018). Recurrent Neural Networks for Multivariate Time Series with Missing Values. *Scientific Reports*, 8, article 6085. doi:10.1038/s41598-018-24271-9.
- [40] Battaglia, Peter W., et al. (2018). Relational inductive biases, deep learning, and graph networks. *arXiv*, arXiv:1806.01261; version 3, 17 October 2018. Original source record.

- [41] Lim, Bryan; Arık, Sercan Ö.; Loeff, Nicolas; Pfister, Tomas (2021). Temporal Fusion Transformers for interpretable multi-horizon time series forecasting. *International Journal of Forecasting*, 37(4), 1748-1764. doi:10.1016/j.ijforecast.2021.03.012.
- [42] Perdomo, Juan C.; Zrnic, Tijana; Mender-Dünner, Celestine; Hardt, Moritz (2020). Performative Prediction. *Proceedings of the 37th International Conference on Machine Learning; Proceedings of Machine Learning Research*, 119, 7599-7609. Original source record.
- [43] Mitchell, Margaret; Wu, Simone; Zaldivar, Andrew; et al. (2019). Model Cards for Model Reporting. *Proceedings of the Conference on Fairness, Accountability, and Transparency (FAT* '19)*, 220-229. doi:10.1145/3287560.3287596.
- [44] Kirkpatrick, James; Pascanu, Razvan; Rabinowitz, Neil; et al. (2017). Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences*, 114(13), 3521-3526. doi:10.1073/pnas.1611835114.
- [45] National Institute of Standards and Technology (2023). Artificial Intelligence Risk Management Framework (AI RMF 1.0). *National Institute of Standards and Technology*, NIST AI 100-1. doi:10.6028/NIST.AI.100-1.
- [46] Guo, Chuan; Pleiss, Geoff; Sun, Yu; Weinberger, Kilian Q. (2017). On Calibration of Modern Neural Networks. *Proceedings of the 34th International Conference on Machine Learning; Proceedings of Machine Learning Research*, 70, 1321-1330. Original source record.
- [47] Ovadia, Yaniv; Fertig, Emily; Ren, Jie; et al. (2019). Can you trust your model's uncertainty? Evaluating predictive uncertainty under dataset shift. *Advances in Neural Information Processing Systems*, 32. Original source record.
- [48] Boettiger, Carl; Hastings, Alan (2012). Quantifying limits to detection of early warning for critical transitions. *Journal of the Royal Society Interface*, 9(75), 2527-2539. doi:10.1098/rsif.2012.0125.
- [49] Shannon, C. E. (1948). A Mathematical Theory of Communication. *Bell System Technical Journal*, 27(3), 379-423; 27(4), 623-656. doi:10.1002/j.1538-7305.1948.tb01338.x. The linked DOI identifies the July installment; the complete work appeared in July and October 1948.
- [50] Wang, Zheng; Busemeyer, Jerome R.; Atmanspacher, Harald; Pothos, Emmanuel M. (2013). The Potential of Using Quantum Theory to Build Models of Cognition. *Topics in Cognitive Science*, 5(4), 672-688. doi:10.1111/tops.12043.
- [51] Friston, Karl; FitzGerald, Thomas; Rigoli, Francesco; Schwartenbeck, Philipp; Pezzulo, Giovanni (2017). Active Inference: A Process Theory. *Neural Computation*, 29(1), 1-49. doi:10.1162/NECO_a_00912.
- [52] Trimmer, John D. (translator; original text by Erwin Schrödinger) (1980). The Present Situation in Quantum Mechanics: A Translation of Schrödinger's "Cat Paradox" Paper. *Proceedings of the American Philosophical Society*, 124(5), 323-338; English translation of the 1935 original. Original source record. Consulted in John D. Trimmer's complete 1980 English translation of the 1935 original.

Appendix A. Notation, dimensional consistency, and implementation assumptions

A.1. Scope of the formal apparatus

The formal apparatus specifies a task-relative model family and a conditional planning objective. It is intentionally restricted. No equation is a representation of a complete person, and no variable denotes personal worth. The notation makes dependencies and inferential assumptions explicit so that an implementation can be examined before substantive conclusions are attached to its parameters.

Table A1. Notation and dimensional ledger. A changed observation interval requires a revised transition specification or justified temporal transformation.

Symbol	Definition	Scale or convention
t, T	Observation index and planning horizon	Daily steps for the two-state example; finite nonnegative integer horizon
x_t	Task-relevant latent state	Model-specific; augmented where context or event history is required
e_t	Measured context vector	Each component retains its declared units
o_t, H_t	Current observation and available history	History contains only information available at the decision time
M, θ	Model specification and uncertain fixed parameters	Model and parameter uncertainty are distinct
b_t^M	Posterior over latent state and parameters	Probability measure conditional on the retained model and history
s_t, μ_s, d_j	Proficiency, reference proficiency, item difficulty	Common anchored log-odds scale
z_t, R_t	Readiness state and indicator	Fixed standardized readiness units
u_t, r_t	Completed practice and recovery-exposure deviation	Hours within the daily interval; exposures, not rates
α	Local practice coefficient	Log-odds per practice hour
β	Practice-readiness interaction	Log-odds per practice hour per readiness unit
γ	Readiness contribution to item response	Log-odds per readiness unit
δ, ρ	Reversion and persistence coefficients	Dimensionless for the specified one-day update
η, λ	Recovery and practice coefficients in readiness dynamics	Readiness units per respective exposure hour
ℓ	Context coefficients	Each component converts its context unit to log-odds
w_t^s, w_t^z, Q	Process innovations and covariance	Corresponding state units; squared and cross-product units in covariance
τ_R	Readiness-indicator error standard deviation	Readiness units
D_0, B	Planning evidence and vector resource budget	Budget components remain separately denominated

Symbol	Definition	Scale or convention
G_g, K	Target region and acceptable-state region	Operationally specified subsets of the relevant augmented state
$\mathcal{G}$	Target labels	Prespecified meaningful candidate family
$\pi, \Pi_{\text{adm}}(B)$	Policy and admissible policy class	Information-adapted, permission-respecting, resource-feasible rules
τ_G, τ_{K^c}	First target entry and first constraint exit	Stopping times for the full-state filtration; not necessarily observable to the policy; infinity if absent
p	Required model-estimated success probability	Chosen threshold in $(0, 1]$; not a confidence coefficient

A.2. Conditions required before estimation

A fitting procedure must specify an initial latent-state and parameter distribution, a positive-semidefinite process covariance, an admissible parameter domain, a calibration convention, and the joint observation law. The Bernoulli marginals in the main text do not by themselves justify multiplying item likelihoods. Conditional independence must be stated or replaced with a dependence model.

Indicator serial dependence and relationships among indicator and item-response errors also require specification. Independence from process innovations is not equivalent to independence from every observation error. If the context is measured with error, treating it as exact may transfer that error into proficiency or readiness estimates.

Missingness requires its own treatment. A missing indicator cannot be replaced with a zero without changing the model. A refusal, inaccessible instrument, missed reminder, and unrecorded failure may imply different observation processes. The reasons may be unknown or deliberately undisclosed; that uncertainty must remain visible rather than being assigned a convenient cause.

The proposed recovery analysis generates observations from known parameters using the intended timing, item calibration, missingness, and input variation. The estimation procedure is then applied without access to the generating values. Recovery failure indicates that the associated parameter interpretation is unsupported for that design. Adequate recovery is required before substantive interpretations are attached to those parameters, but does not demonstrate that the assumed model is empirically correct. A bounded predictive quantity may remain estimable despite ambiguity among underlying parameters; its predictive adequacy requires separate evaluation.

A.3. Policy filtration and realized exposure

Let $\mathcal{F}_t$ be the information available immediately before action selection. The policy is adapted to this information: it may use current and past observations, earlier actions, remaining resources, current permissions, and independent randomization. It may not condition on unknown parameters or future observations as if they were known.

For the daily example, completed recovery exposure during day t is typically unavailable at the morning decision. A forecast must average over a declared exposure law or identify the exposure as a scenario assumption. A fitted retrospective regression using the realized exposure is not automatically an executable prospective policy.

Likewise, practice undertaken is an observation of realized exposure. A policy over assigned practice requires either a stated perfect-implementation simplification or a model connecting assignment to

realized practice. If that relation is unknown, the output remains conditional on completed practice. Even for statistical filtering, an unconditional Gaussian innovation law does not imply the same law conditional on correlated same-day exposures. The implementation must justify the declared conditional innovation law or specify a joint exposure-innovation process. Independence across days alone establishes neither this statistical condition nor the additional assumptions required for causal interpretation.

A.4. History-dependent events under partial observation

The reach-avoid value is defined directly over trajectories and information-adapted policies. A recursive implementation that retains only a posterior over the current physical state can lose necessary event history. Whether the target was reached earlier, or a constraint was violated earlier, may not be recoverable from the current state alone.

An implementation can retain target-attainment and prior-admissibility indicators in an augmented state, together with remaining budgets. Under partial observation, these event indicators can themselves be latent; their joint posterior with the current state and parameters must be retained when required by the calculation. They are computational records of the defined success event, not additional psychological constructs.

A.5. Common assumptions for the analytical comparisons

For the comparisons in Appendix B, hold the model, transition and observation laws, posterior, target family, context law or scenario, and constraints fixed except for the quantity explicitly varied. Budgets are ordered componentwise. Policies obey the same information restrictions. Resource feasibility applies on permitted trajectories rather than merely in expected expenditure.

The relevant events and probability functions are measurable, and the admissible policy class is nonempty. Where a pointwise supremum is integrated, measurability is assumed; a finite or countable class supplies a simple sufficient setting. No result requires an optimizing policy to attain the supremum. The probability threshold is strictly positive, since a zero threshold would admit every candidate target without discrimination.

Appendix B. Analytical consequences and counterexamples

The following results are elementary consequences of the stated specification. Their role is expository and diagnostic. They establish neither empirical adequacy nor mathematical novelty.

B.1. Location indeterminacy before item anchoring

Proposition B1. In the two-state model with unanchored item difficulties, hold observed inputs and context fixed, or use the same exogenous laws for them. Hold remaining parameters and noise laws fixed. For any real constant c , define

$$s'_t = s_t + c, \quad \mu'_s = \mu_s + c, \quad d'_j = d_j + c. \quad (\text{B.1})$$

Translate the initial proficiency distribution by the same constant. The transformed model induces the same distribution of item responses and readiness indicators.

Proof. The response predictor is invariant because $s'_t - d'_j = s_t - d_j$. Readiness and context contributions are unchanged. In the proficiency transition,

$$(1 - \delta)s'_t + \delta\mu'_s = (1 - \delta)s_t + \delta\mu_s + c. \quad (\text{B.2})$$

Thus the transformed proficiency process is exactly the original process shifted by c , while the readiness process is unchanged. Translating the initial law preserves the full observation law, including when subsequent actions are selected from the same observed history. No observation under this specification distinguishes these locations. **QED.**

An anchoring constraint removes this particular freedom. An informative prior can also select a location, but prior selection is not identification from the response data. Neither operation establishes identifiability of the remaining parameters.

B.2. Budget, target, and threshold monotonicity

Proposition B2. Hold the common objects in Appendix A fixed. If $B \leq B'$ componentwise and the admissible policy classes are nested, then

$$V_{G,K}^M(D_0, T, B) \leq V_{G,K}^M(D_0, T, B'). \quad (\text{B.3})$$

Consequently the reported target set at budget B is a subset of the reported set at B' , for a common threshold. If $0 < p_1 \leq p_2 \leq 1$, then

$$\widehat{\mathcal{R}}_{p_2}^M(D_0, T, B, K) \subseteq \widehat{\mathcal{R}}_{p_1}^M(D_0, T, B, K). \quad (\text{B.4})$$

Finally, enlarging a target region from G to G' , with $G \subseteq G'$, cannot reduce the value when the remaining objects are unchanged.

Proof. A supremum over a larger feasible class cannot decrease. Applying the same threshold preserves the corresponding target-set inclusion. A value at least p_2 is also at least p_1 , yielding the threshold inclusion. Every trajectory entering G also enters G' ; the acceptability event is unchanged. Probability, integration, and the supremum preserve this event inclusion. **QED.**

All inequalities permit equality. Budget monotonicity does not assert that a resource change leaves obligations, preferences, institutional conditions, or experience unchanged. If those change, the fixed-comparison premise may fail. Enlarging a target is also a relaxation of the success criterion, not evidence of developmental improvement.

B.3. Posterior optimization and a clairvoyant benchmark

For compactness, let $q_G^\pi(x, \theta)$ denote the success probability for fixed initial state and parameters under policy π , with horizon and constraints held fixed.

Proposition B3. Under the common-policy and measurability assumptions,

$$\sup_{\pi \in \Pi_{\text{adm}}(B)} \int q_G^\pi(x, \theta) b_0^M(dx, d\theta) \leq \int \sup_{\pi \in \Pi_{\text{adm}}(B)} q_G^\pi(x, \theta) b_0^M(dx, d\theta). \quad (\text{B.5})$$

Proof. For every admissible policy and every initial-state/parameter pair, its success probability is bounded by the pointwise supremum. Integration preserves the inequality. Taking the supremum over the policies on the left proves the result without requiring attainment. **QED.**

Strict inequality is possible. Let an unknown parameter have two equally probable values, zero and one. There is a single decision with actions zero and one, no informative observation before the decision, and no binding resource or acceptability restriction. Success occurs exactly when the action matches the parameter. Every legitimate policy, including an independent randomized policy, succeeds with probability one half. Selecting the action separately for each parameter value succeeds with probability one.

The first value is attainable under the specified information restriction. The second is a benchmark with unavailable information. A common policy can still adapt to legitimate later observations; the restriction does not require an open-loop action sequence.

B.4. Full-horizon acceptability and nonmonotone horizons

Counterexample B4. Consider three states a, g, f , one action, zero costs, and a known initial state a . Transitions are deterministic:

$$a \longrightarrow g \longrightarrow f \longrightarrow f. \quad (\text{B.6})$$

Let $G = \{g\}$ and $K = \{a, g\}$. Target entry occurs at time one and constraint exit at time two. Therefore

$$V_{G,K}^M(D_0, 1, B) = 1, \quad V_{G,K}^M(D_0, 2, B) = 0. \quad (\text{B.7})$$

The target belongs to the thresholded set at the first horizon and not at the second for any positive threshold. The additional interval extends exposure to an unacceptable state after success. A stopping-at-success convention or a feasible action that guarantees continued acceptability would alter the problem. Neither is entailed by the original reach-avoid definition.

B.5. Epistemic contraction without developmental deterioration

Counterexample B5. Hold the physical situation, planning time, horizon, model, target, resources, and action set fixed. An unknown parameter is favorable with prior probability 0.9 and unfavorable with probability 0.1. There is one action, all states are acceptable, and the target is reached precisely in the favorable case. The model-estimated value is 0.9, so the target passes a threshold of 0.8.

A costless diagnostic observation before action reveals the unfavorable case without changing the state. The posterior success probability becomes zero, and the target leaves the reported set. The observation changes the estimate of a situation already present. It does not demonstrate deterioration in the person or environment.

This conditional reduction is consistent with the ex ante value of information. When information is costless, changes neither dynamics nor feasible actions, and may be ignored, every formerly admissible policy remains available. Optimal ex ante value cannot be reduced merely by making the information

available. Individual posterior realizations can nevertheless be unfavorable; in this example the expected posterior value remains 0.9.

B.6. Fixed-input equilibrium and its interpretation

Proposition B6. Conditional on fixed parameter values, suppose practice and recovery inputs are constant and exogenous, the initial state has finite second moments, and innovations are identically distributed, independent across time and independent of the pre-innovation state and history, with zero mean and fixed finite covariance. Assume $0 < \delta < 1$ and $|\rho| < 1$. The mean two-state dynamics have diagonal coefficients $1 - \delta$ and ρ , with the readiness-to-proficiency off-diagonal coefficient βu . Their equilibrium is

$$z^* = \frac{\eta r - \lambda u}{1 - \rho}, \quad s^* = \mu_s + \frac{\alpha u + \beta u z^*}{\delta}. \quad (\text{B.8})$$

Derivation. Set both mean increments to zero, solve the readiness equation first, and substitute its equilibrium into the proficiency equation. The upper-triangular transition matrix has eigenvalues $1 - \delta$ and ρ , both strictly inside the unit circle, so the mean converges to the stated equilibrium. The unique stationary covariance satisfies the corresponding discrete Lyapunov equation. This covariance conclusion is conditional on the fixed parameters; a posterior mixture over stable parameter values need not have finite unconditional covariance. **QED.**

This is a property of a local stochastic model under fixed forcing. It establishes neither a human equilibrium nor an innate ceiling. In particular, the proficiency equilibrium can exceed or fall below μ_s . When $\delta = 0$, the reference parameter disappears and the preceding stability argument for proficiency no longer applies.

B.7. Further ambiguities after calibration

Anchoring the item scale does not resolve confounding created by the input design. With constant nonzero practice u_0 , the proficiency intercept contains the combination $\delta \mu_s + \alpha u_0$. For fixed δ , changing μ_s by c and α by $-\delta c/u_0$ leaves the transition unchanged.

If recovery exposure is always c times practice, the readiness input contains only $c\eta - \lambda$. Increasing η by a and λ by ca preserves it. If readiness is fixed at z_0 , the practice effect contains only $\alpha + \beta z_0$. These are exact ambiguities under the stated conditions. Near-constant or nearly collinear inputs can create practical estimation difficulties without exact nonidentifiability.

The appropriate response is a revised measurement or input design, a restricted parameterization, or an explicit uncertainty report. Assigning distinct names to parameters does not establish that observations distinguish them.

Appendix C. Information architecture and revision protocol

C.1. Record classes

The schema below is a proposed engineering specification. Field names identify conceptual distinctions; they are not evidence that an implementation exists or has been validated. Content references should resolve through access controls, and records should be collected only where required for the agreed purpose.

Table C1. Minimal record classes. Goal and policy versions are referenced explicitly; their content and authorization history must remain recoverable.

Record	Required distinctions and representative fields
Consent	consent_id, person_id, purpose, allowed_sources, granted_at, expires_at, revoked_at
Source	source_id, origin, author_or_observer, source_type, provenance_parent_ids, access_scope, content_ref
Observation	observation_id, source_id, consent_id, observed_at, recorded_at, instrument_version, value_ref, units, missingness, supersedes_id, revision_reason
Event	event_id, occurred_start, occurred_end, occurrence_status, supporting_observation_ids
Model	model_version_id, parent_version_id, specification_hash, parameter_hash, training_observation_ids, evaluation_ref, approval_id, release_status
Inference	inference_id, model_version_id, input_observation_ids, generated_at, target_time, output_ref, uncertainty_ref, supersedes_id
Exposure	exposure_id, inference_id, goal_version_id, policy_version_id, displayed_at, accepted_at, performed_at, outcome_event_ids

The event status differentiates observed, reported, and inferred occurrence. The observation record distinguishes when something occurred from when it entered the system. An inferred event must retain its inferential status even if it later becomes a familiar part of the narrative. Familiarity and repeated summarization do not transform an inference into an observation.

Missingness values should distinguish at least an unavailable measurement, a missed record, and a deliberate refusal when those distinctions are known. A participant need not explain a refusal. Unknown reasons remain unknown, and a downstream encoder must not silently convert them into a personal characteristic.

C.2. Semantic and claim records

A semantic record links the original expression to its contextual interpretation, communicative-act classification, participant confirmation or correction, and unresolved alternatives. A correction to meaning is a revision of that attribution. It need not imply that the original expression was insincere or that the external claim has been verified.

A claim register records exact wording, claim type, scope, supporting and conflicting evidence, assumptions, alternatives, provenance, and the condition under which confidence should change. Numerical confidence is appropriate only when its interpretation has a defensible basis. Evidence-type descriptions are often preferable to a percentage that collapses different forms of warrant.

Source provenance must preserve shared ancestry. A quoted study, a summary of that study, and a generated paraphrase of the summary form a dependency chain. They cannot be counted as three

independent observations. The system may retain all three records for usability while preserving the single underlying evidential lineage.

C.3. Algorithmic revision procedure

Procedure C1: correction propagation under fixed authorization. The following is a specification of expected behavior, not a claim about executed software.

1. **Validate scope.** Check that the requested correction or new observation falls within current permissions and the agreed purpose. A source document cannot extend that scope.
2. **Preserve identity.** Retain the prior record and create a revision with event time, recording time, supersession relation, and reason. Distinguish a corrected observation from a newly inferred event.
3. **Locate dependencies.** Identify interpretations, estimates, forecasts, and displayed outputs that relied on the superseded record. Preserve the model and goal versions attached to those outputs.
4. **Recompute permitted inferences.** Apply the approved model or an explicitly authorized revised model. Keep uncertainty and unresolved alternatives visible. Do not infer a recent developmental event solely from a changed estimate.
5. **Mark exposure consequences.** Record whether an earlier output was displayed or acted upon. Future evaluation must retain that exposure history.
6. **Evaluate and communicate.** Compare changed outputs with the previous versions and explain the reason for the difference at an appropriate level. Escalate material unresolved discrepancies to review or abstention.
7. **Retain reversibility.** Preserve the permitted rollback path and apply retention or deletion rules. A record may need removal from access even when a historical log must retain a nonrevealing indication that a revision occurred.

The procedure does not permit a model to approve its own expansion of authority. Knowledge acquisition, model deployment, and permission changes remain separate operations.

C.4. Release and goal versioning

Maintain distinct versions for the conceptual framework, data schema, measurement instruments, model specifications, parameters, goals, and permissions. A wording clarification need not imply a new theory. A changed outcome definition may invalidate a numerical comparison even when the software version changes only slightly.

A challenger model is evaluated for the intended use against the retained model and simpler baselines. Release decisions consider future prediction, uncertainty behavior, contextual failures, participant burden, and documentation. Repeated exposure of a test set during development should be treated as adaptation to that set; later releases require new evaluation evidence.

Goal revision is an authorized change in the criterion of success. It is not automatically a model error or evidence of irrationality. Historical outcomes should remain associated with the goals under which they were evaluated. Rewriting that history to make the latest objective appear timeless would erase information necessary for interpretation.

Appendix D. Prospective evaluation and adversarial acceptance criteria

D.1. Minimal protocol record

A proposed study should document its question, participant-defined purpose, activity and task family, setting, measurement instruments, timing, outcome definition, assistance conditions, and allowable interventions. It should identify the primary analysis, comparison models, evaluation windows, missingness policy, exclusions, meaningful-improvement criterion, and conditions for stopping or revising the procedure.

The empirical endpoint must be interpretable independently of APD terminology. For example, completion of a specified learning activity is an observable outcome; a claim that a participant has become globally more developed is not an adequate replacement. If reported experience is the endpoint, its measurement and intended interpretation should be stated accordingly.

The design should distinguish immediate supported performance, delayed unaided performance, transfer, subjective usefulness, burden, and autonomy. It need not measure every endpoint. It must avoid claiming an unmeasured endpoint from a favorable result on another.

D.2. Eight adversarial tests

Table D1. Adversarial cases examine different failure modes. Passing them does not establish general safety, effectiveness, or validity; failing one identifies a concrete deficiency to address.

Test	Perturbation	Required behavior and failure implication
D1. Evidential duplication	Introduce several paraphrases or summaries of one underlying study.	Independent-evidence attribution remains tied to the original lineage. Inflation indicates a provenance or aggregation failure.
D2. Retrospective correction	Correct an earlier measurement while preserving its event time.	Dependent inferences are revised without inventing a contemporary developmental event. Failure indicates temporal or ontic-epistemic conflation.
D3. Context displacement	Introduce a setting outside the model's supported observation domain.	Reduced support is disclosed, uncertainty is reevaluated, or the system abstains. Unchanged unwarranted confidence indicates a calibration or domain-detection failure.
D4. Complexity challenge	Replace an elaborate component with a prespecified simple baseline.	The additional component demonstrates prospective value sufficient to justify its burden. Comparable performance favors the simpler operational choice.
D5. Forecast exposure	Display a prediction before collecting its outcome.	Exposure and subsequent actions remain recorded and enter evaluation. Treating the outcome as unaffected observational evidence indicates a feedback-accounting failure.
D6. Hostile source content	Insert instructions into material retrieved as evidence.	The content cannot alter permissions, release criteria, or execution authority. Success by the injected instruction requires feature suspension and correction.
D7. Revocation and rollback	Revoke a source and execute a documented rollback.	Retrieval, access, release state, and residual model influence match the documented policy. Unsupported claims of deletion or unlearning constitute a failure.
D8. Objective rejection	Decline a recommendation, change the goal, or pause participation.	Agency, permitted access, and nonpunitive treatment are preserved. Coercive or retaliatory feedback violates the normative specification.

D.3. Interpreting prospective results

Three generalization questions should remain separate: prediction later in time for an observed participant, prediction for an unobserved participant, and prediction in an unobserved context. Reporting a pooled average can obscure failure in any one of these regimes. Where data are insufficient to estimate a regime-specific result, the appropriate report is limited support.

Interaction studies should state the null model and the scale on which an interaction is tested. A combined intervention can outperform either component while remaining compatible with simple additivity. Differences in total practice, attention, or task difficulty can also produce apparent interaction patterns. Durable learning requires an explicit account of carryover and observation timing.

For intervention studies, distinguish assignment, uptake, realized exposure, and outcome. A favorable association with completed practice is not the same estimand as the effect of assigning practice. Noncompliance, selection, and attrition can change the interpretation of a comparison even when randomization was initially appropriate.

D.4. Failure responses and publication of negative findings

Before consequential use, the protocol should associate plausible failures with responses. Measurement ambiguity can require recalibration or suspension of comparisons. Unidentified parameters can require restriction or redesign. Poor prospective prediction can require reversion to a baseline. Harmful feedback or excessive burden can require withdrawal of an output despite adequate predictive scores.

Some failures undermine a component; others undermine the chosen integration for the domain. The framework should not respond to every failure by increasing representational complexity. A result that demonstrates no incremental benefit is informative when the comparison was adequate and the endpoint meaningful.

Reports should include unsuccessful modules, adverse contextual results, sensitivity to modeling assumptions, and the boundaries of replication. No sample size or observation duration is universal. Their adequacy must be justified in relation to dependence, variability, missingness, and the precision required for the intended claim.

Appendix E. Hypothetical analytical cases

The following cases are hypothetical analytical constructions. They report no observations, fitted parameters, or intervention effects. Their purpose is to distinguish competing interpretations, identify the evidence each would require, and clarify which conclusions would remain unwarranted. Each case retains the person-activity-conditions-time unit of analysis while varying the explanatory question. No case supplies an estimate of an individual's ultimate developmental capacity.

E.1. Perceived procedural pace and measured performance

A participant completing an equipment-assembly simulation reports that the instructions proceed too slowly, although their final accuracy is below the task's maximum. The report initially leaves several variables conflated: waiting during instruction, time taken to execute an action, perceived difficulty, engagement, and response accuracy. A subjective judgment about one variable cannot identify the others.

The assessment should first distinguish the participant's intended meaning. They may find introductory steps repetitive, experience delays between actions, or perceive different sections as unevenly paced. The protocol can then examine instruction duration, execution time, item-level errors, response-format familiarity, and the conditions under which assistance was available. These observations address different candidate explanations. Rapid completion of some steps may coexist with errors elsewhere; an unfamiliar control interface may impede an otherwise understood procedure; unengaging presentation may coexist with demanding content.

The resulting account should specify which distinctions the evidence resolves and which remain uncertain. A mismatch between presentation pace and preference would not, by itself, establish underestimated proficiency. Conversely, an error does not identify a stable deficiency or explain the participant's reported experience. Repeating similar informal exercises without preserving administration conditions may add familiarity without clarifying the original discrepancy.

The appropriate output is therefore a conditional account of performance and procedural experience. Claims about proficiency require the task's measurement assumptions and actual response pattern; the participant's description remains relevant evidence about how the procedure was experienced.

E.2. Interpretive change and differentiated outcomes

A participant in a collaborative design project reports that considering several interpretations of disagreement has improved their experience of teamwork. This report is compatible with distinct changes: greater curiosity about colleagues' priorities, less certainty about inferred motives, more frequent requests for clarification, or a revised understanding of what disagreement signifies. These possibilities should remain distinguishable rather than become interchangeable names for a single transformation.

The evaluation begins by specifying the outcomes that matter to the participant. Fewer avoidable misunderstandings, more accurate anticipation of another member's stated position, improved project decisions, and increased experienced meaning require different observations. A reflective record can describe experience. Voluntary feedback from collaborators can address communicative conduct. Predictions recorded before a discussion can later be compared with what participants actually state. Agreement among these records would strengthen a particular interpretation; disagreement would require examining their different targets and limitations.

Scope must also be tested. The practice may help routine deliberation but become difficult under time pressure or substantial conflict. Recognizing another interpretation need not entail accepting it, abandoning a justified boundary, or treating all explanations as equally well supported.

A defensible conclusion could therefore be selective: the practice may improve communication and experienced participation without altering unrelated task scores. Its effectiveness may depend on the setting and require continued practice. Such a result would constitute a meaningful local change without establishing a general improvement across every cognitive or interpersonal domain.

E.3. Equivalent scores through support and practice

Two participants achieve the same final score on a software-configuration task. One uses a clearly organized checklist; the other has practiced the procedure extensively. Equality of observed performance does not establish equality of learning histories, unaided proficiency, dependencies, or available opportunities. The immediate explanandum is how each arrangement produced the recorded outcome.

An informative record distinguishes the checklist as an environmental support from practice as a learning exposure. Immediate assisted performance, retention after a delay, and performance on a prespecified untrained task answer separate questions. Assessment should identify which of these questions serves the practical objective before selecting a comparison. A change of interface may test portability, while a delayed task may test retention; neither is merely another version of the original score.

If performance declines when a support is absent, the result identifies a dependency under the tested conditions. It does not invalidate the supported achievement or establish that maintaining the support is undesirable. For a workplace in which a reliable checklist can remain available, successful assisted functioning may be the relevant outcome. Where the task must sometimes be completed without that resource, unaided performance becomes a separate operational requirement.

APD consequently represents both the achievement and its conditions. The practical choice concerns the reliability, cost, availability, and value of the arrangement, together with any demonstrated learning or transfer. A single global score would obscure precisely the distinctions needed to choose between continued support, additional training, or a combination.

E.4. Abrupt success after gradual preparation

A fault-detection exercise requires a participant both to interpret a signal display and to recognize a mismatch between measurement units. Suppose performance remains poor while either prerequisite is insufficient, even as competence in the other improves. Once both are adequate, task success may increase sharply. This logical conjunction illustrates how an abrupt change in an outcome can be compatible with gradual changes in its components.

The construction does not establish that any observed breakthrough has this mechanism. Alternatives include discovering a different strategy, changing the representation of the problem, receiving clearer feedback, or revising the criterion by which success is judged. The appearance of a transition therefore supplies a question rather than a causal diagnosis. Calling the change a bifurcation would add a mathematical claim requiring specified dynamics and parameter dependence.

To investigate the prerequisite account, observations should distinguish the proposed components before and after improvement. Their hypothesized relation to task success should be stated before examining the decisive outcomes. If the components neither predict the transition nor distinguish success from continued difficulty, the proposed explanation should lose credibility. A plausible retrospective account is insufficient.

The broader research question is whether a task-specific account of prerequisites explains a plateau better than a stable individual learning-rate parameter. Either account may prove inadequate. The value of the example lies in making that competition examinable while separating observed discontinuity from a claim about the mechanism producing it.

E.5. Conditional planning for continuing support

A training coordinator and a consenting participant must plan the support required for a changing workplace role. Extrapolating the present arrangement indefinitely may neglect learning, improved tools, altered responsibilities, or revised preferences. Assuming that support will inevitably become unnecessary may neglect persistent barriers and future demands. Neither forecast follows from the current observation alone.

The planning exercise therefore specifies a small set of conditional scenarios. One retains the existing role and resources. Another introduces a proposed tool or training opportunity. A third changes the allocation of responsibilities or the accessibility of the procedure. Each scenario identifies the conditions on which its expectations depend, the resources required, the uncertainties that remain, and the decisions that can be reconsidered later. These are provisional forecasts, not guarantees of intervention effects.

The participant's goals and consent govern the arrangement. A planner cannot infer an unrestricted preference from acceptance of the only available option. Successful planning might preserve dependable support, facilitate a newly valued activity, or reduce an avoidable burden; it need not aim at one universal trajectory toward unaided performance.

Review points allow later observations to revise the plan without treating revision as a failure. Contingencies can accommodate changed circumstances while preserving agency. The contribution of APD is the explicit connection between a forecast and its conditions, together with the separation of desirable outcomes from predicted ones. It does not promise convergence toward a common developmental endpoint.

Appendix F. Source-use boundaries and evidential roles

F.1. Literature selection and source roles

The source register preserves the 52 works informing the APD specification. Selection is purposive and interdisciplinary. It does not establish exhaustive coverage or a systematic-search claim. The text separates original empirical reports, original meta-analytic analyses, formal and computational contributions, theoretical syntheses, philosophical arguments, and official guidance.

An original meta-analysis is cited as an analysis of its included studies, with the scope of that synthesis preserved. A theoretical model supplies a candidate mechanism or formal possibility. A philosophical work supplies a conceptual or normative argument. Official guidance supplies a governed practice or reporting framework. None of these roles constitutes direct empirical validation of APD's integration.

The bibliography provides original publication identifiers and author, publisher, or official records. Source-specific claims in the manuscript are deliberately narrower than the full programs associated with those works. The selected quantitative findings in the underlying literature have not been recomputed here, and no pooled APD effect size is reported.

F.2. Textual access and interpretive limits

Source consultation ranges from selected sections of original full texts to original abstracts, official recommendations, and publisher descriptions or previews. Bibliographic metadata checks verify publication identity; they do not independently verify substantive results. The access register below identifies the consultation supporting this edition and retains limits where full-text reassessment was unavailable.

Table F1. Access classification records the level of consultation supporting the bounded source use. Reference numbers link to the original publication records.

References	Consultation level	Permitted use in this article
[1, 2, 7, 11, 13, 14, 15, 17, 20, 21, 25, 28, 29, 30, 31, 32, 33, 35, 36, 42, 44, 46, 48, 49, 52]	Selected original text. Selected sections inspected; no claim of exhaustive textual analysis.	Bounded claims tied to the inspected methods, arguments, results, or limitations.
[3, 8, 9, 10, 16, 18, 22, 23, 24, 26, 27, 34, 38, 39, 40, 41, 43, 47, 50, 51]	Original abstract or opening. Original abstracts, introductory passages, or author records.	The stated high-level finding or proposal; no unsupported reconstruction of detailed methods.
[12, 45]	Official recommendations or report. Relevant official guidance or report sections.	The specified practice or governance orientation; no APD certification or clinical validation.
[4, 5, 6, 19, 37]	Publisher or preview. Publisher description, contents, or available preview.	Broad declared orientation only; full arguments and technical chapters are not represented as examined.

The broad orientations attributed to Giere, Longino, Fricker, Aubin and colleagues, and Varela, Thompson, and Rosch are limited to the relevant publisher or preview material [4, 5, 6, 19, 37]. The manuscript does not present a comprehensive textual analysis of those books. Their appearance in the bibliography should not be read as evidence that every chapter was examined.

F.3. Cross-domain transfer restrictions

Conceptual and formal borrowing must state its inferential role. The architectural properties of a convolutional network do not establish a geometry of psychological attributes. Evolutionary feedback does not identify the mechanism of a short-term adult learning event. Ecological persistence does not determine ethical desirability. A causal identification result does not validate a particular psychological measurement scale.

The same restriction applies to numerical results. An educational effect estimate, an observed within-person pattern, or a forecasting benchmark remains conditional on the study's design and domain. The manuscript uses these results to motivate distinctions and tests. It does not infer an individual's numerical parameter or lifelong developmental limit from a published group result.

F.4. Independence and negative evidence

The register preserves evidential ancestry. A book citing an experiment, a paper summarizing that book, and an automated summary of the paper do not yield independent replications of the original experiment. The relevant unit of evidential independence is determined by the underlying observations and their dependencies, not the number of references in a paragraph.

Claims should retain contrary observations and unresolved alternatives alongside supporting sources. When a finding fails to generalize, the appropriate revision may concern population, activity, measurement, timescale, intervention, or the proposed mechanism itself. A broader vocabulary does not remove the need to specify which claim has changed.

Appendix G. Terminology and construct boundaries

The following terms make the external formulation more precise. Their definitions stipulate how they are used in this manuscript; they do not introduce additional empirical entities.

Table G1. Controlled terminology. A term earns its use by specifying a distinction or requirement; terminological elaboration alone supplies no additional evidence.

Term	Definition in this specification
Ontic transformation	Change in the modeled person or environment, distinguished from a change in description or belief.
Epistemic updating	Revision of a warranted estimate or interpretation within a retained model, following evidence available at that time.
Model revision	Change to constructs, equations, measurement relations, or representational structure.
Situated capability	Task-relevant capacity or effective functioning considered with the conditions through which it can be expressed.
Idiographic inference	An inference concerning a particular person's dynamics; population-level evidence requires an explicit transport argument before serving this role.
Measurement invariance	Sufficient comparability of the measurement relation across the conditions relevant to the proposed comparison.
Structural identifiability	Distinguishability of relevant model parameters in principle under the model and observation design.
Practical identifiability	Adequate distinguishability with the finite, noisy observations actually available.
Partial observability	Incomplete direct access to the state variables relevant to a specified forecast or decision.
Non-anticipative policy	A decision rule restricted to information available when the decision is made.
Conditional reachability	Model-estimated ability to enter a prespecified target region within a horizon under declared actions, resources, and constraints.
Viability	Existence of an admissible course maintaining specified constraints over the relevant interval.
Pareto nondominance	Absence of a feasible alternative at least as favorable on every selected objective and strictly better on at least one.
Provenance	Recoverable source identity and dependency history for observations, claims, and derived inferences.
Performative feedback	Influence of a prediction or its use on the behavior and outcomes subsequently observed.
Contestability	Procedural ability to challenge observations, interpretations, model use, or objectives and obtain an appropriate revision or response.
Normative pluralism	Explicit recognition of multiple potentially conflicting values without an assumed universal scalar ordering.
Incremental validity for a use	Demonstrated additional value of a proposed component relative to appropriate simpler comparators for a prespecified purpose.