

# HBC-Glicko: A Bloom-Constrained Adaptive Assessment Architecture with Uncertainty-Aware Progression

Wissal EL Fougour\*<sup>ORCID</sup>, Mohamed Erradi<sup>ORCID</sup>

TIMS Laboratory-S2IPU-Higher Normal School of Tetouan (ENS), Abdelmalek Essaâdi University (UAE), Morocco

**Abstract**—Adaptive assessment systems typically model learner ability as a single continuous latent trait and select items accordingly. Although effective for efficient measurement, such models rarely encode hierarchical cognitive progression as a formal constraint within the learner model itself. In educational domains where higher-order performance depends on sufficiently stable prerequisite knowledge, this omission may permit pedagogically incoherent trajectories, including premature advancement or misleading readiness judgments produced by compensatory aggregation across cognitively distinct levels, where strength at one level can mask weakness at another. This study pursues three objectives: 1) to formalize an adaptive assessment architecture in which Bloom's revised taxonomy operates as a non-compensatory structural constraint on learner progression, meaning that readiness must be established at each prerequisite level independently, without any area of strength substituting for insufficient evidence elsewhere; 2) to derive the principal theoretical properties that follow from this design; and 3) to situate the proposed architecture relative to existing rating-based learner models. To address these objectives, the study proposes Hierarchical Bloom-Constrained Glicko (HBC-Glicko), a theory-driven measurement architecture for formative adaptive assessment. Instead of representing learner state as a single scalar estimate, HBC-Glicko models it as a band-specific vector with band-level uncertainty. Progression is regulated through anchor-based readiness thresholds and confidence-bound decision rules, such that advancement depends on both estimated performance and evidential stability at prerequisite levels. The study formalizes the learner model, within-band routing logic, threshold construction, promotion and prerequisite reinforcement rules, and derives the principal architectural properties. The contribution is conceptual and architectural rather than empirical, establishing a formally specified foundation for subsequent simulation and empirical investigation.

**Keywords**—Adaptive assessment; Bloom's revised taxonomy; Glicko rating system; learner modeling; non-compensatory progression; uncertainty modeling; formative assessment

## I. INTRODUCTION

Adaptive assessment systems typically rely on learner models that represent students digitally to personalize content, pathways, or task selection [6], [51], [61], [32]. In many adaptive learning systems, learner performance data are used to dynamically adjust assessment difficulty and guide next-item selection during the testing process [22]. Although effective for

difficulty calibration, such designs do not structurally encode hierarchical cognitive progression as conceptualized in Bloom's revised taxonomy. Recent work continues to show both the value and the difficulty of representing cognitive demand accurately: revised Bloom's taxonomy remains a hierarchical framework for organizing assessment objectives [22], yet evidence suggests that cognitive-process classifications cannot be reliably reduced to action verbs alone [30], and many assessment instruments still concentrate heavily on lower cognitive levels rather than spanning the full hierarchy [42].

This creates a structural risk for adaptive systems: when heterogeneous evidence is aggregated into a single performance representation, strong performance at one cognitive level may mask instability at another, allowing progression decisions that are psychometrically efficient but pedagogically incoherent. Existing rating-based adaptive systems, including Elo and Glicko variants, support dynamic learner estimation and efficient online updating, and Glicko-based approaches explicitly model confidence in the current rating [1], [58]. However, recent reviews of learner modeling and adaptive assessment still describe a landscape organized around model families, algorithms, and personalization mechanisms rather than architectures that formalize Bloom hierarchy itself as a non-compensatory progression constraint [6], [22]. Although prior work has shown that Elo-based learner modeling can be effectively applied in adaptive educational systems [39], these approaches generally remain single-dimensional and do not govern advancement through explicit uncertainty-bounded progression rules. Bloom-related work in this line has also treated cognitive taxonomy primarily as a way of characterizing task complexity rather than as a formal control structure for adaptive progression [41].

To address this gap, this study proposes the Hierarchical Bloom-Constrained Glicko (HBC-Glicko) architecture, in which learner ability is represented as a band-specific vector aligned with Bloom's cognitive structure, and advancement is controlled through anchor-derived readiness thresholds combined with confidence-bound constraints. In this way, progression depends not only on estimated ability but also on evidential stability at prerequisite levels. The contribution is conceptual and architectural rather than empirical: the study formalizes a measurement architecture and derives its principal theoretical properties, establishing a foundation for subsequent simulation and empirical investigation.

\*Corresponding author.

More specifically, the present work pursues three objectives:

O1. To formalize an adaptive assessment architecture in which Bloom's revised taxonomy operates as a non-compensatory structural constraint on learner progression, rather than as a descriptive classification of task difficulty.

O2. To derive the principal theoretical properties that follow from combining band-specific learner representation, Glicko-based uncertainty modeling, and confidence-bound progression rules within a single measurement framework.

O3. To situate the proposed architecture relative to existing rating-based learner models and clarify the structural distinctions it introduces with respect to cognitive hierarchy, uncertainty quantification, and progression control.

The remainder of the study is organized as follows. Section II develops the theoretical rationale grounding the architecture. Section III presents the overall system architecture. Sections III-A through III-E formalize the learner model, within-band adaptation, anchor-based thresholds, hierarchical decision rules, and inactivity-driven uncertainty growth, collectively addressing the first objective. Section IV derives the principal architectural properties of the design (Objective 2). Section V situates HBC-Glicko relative to existing approaches and discusses its theoretical implications (Objective 3). Sections VI and VII address limitations and directions for future research.

## II. BACKGROUND AND THEORETICAL RATIONALE

### A. Adaptive Assessment and the Limits of Scalar Learner Models

Adaptive formative assessment systems personalize item delivery by combining online learner-state estimation with sequential item selection, commonly operationalized through IRT-based computerized adaptive testing, Elo-style rating updates, and Bayesian knowledge tracing. In computerized adaptive testing (CAT), learner performance is typically summarized by a single latent proficiency parameter  $\theta$  under unidimensional item response theory [62], an approach whose feasibility continues to be demonstrated in large-scale digital applications [9].

However, this scalar abstraction assumes that performance can be adequately explained along a single proficiency continuum, an assumption that holds only when the measured construct is approximately unidimensional [65]. When tasks differ in cognitive process requirements, collapsing performance into one dimension produces compensatory interpretations in which strengths in one type of task mask weaknesses in another. Multidimensional adaptive frameworks address this limitation by estimating multiple latent traits, including approaches that represent uncertainty in multivariate ability estimation through confidence regions [31]. A similar limitation appears in rating-based learner models: educational adaptations of Elo efficiently update learner and item estimates after each interaction [58] but generally summarize learner state as a single scalar rating without encoding variation in cognitive demand [1].

Knowledge-component and skill-hierarchy models provide a further alternative to scalar learner models by representing learner performance in relation to discrete skills or prerequisite

structures rather than as a single global proficiency estimate [12], [10], [38], [63]. Their main strength is that they can support adaptive sequencing by tracking evidence of mastery across specific components.

However, their dependency structure usually concerns concepts, procedures, or skills rather than Bloom-level cognitive-process demands. From a progression-control perspective, the key limitation is architectural: these models may estimate what a learner is likely to know or perform, but they do not define when the learner has demonstrated sufficiently stable competence at one cognitive level to justify advancement to the next. To our knowledge, existing IRT-, Elo-, and knowledge-component-based architectures do not treat Bloom's hierarchy as an internal non-compensatory constraint on the learner model itself. When cognitive levels appear, they function as item descriptors or reporting categories rather than as confidence-bounded rules governing progression. This limitation motivates learner representations that preserve distinctions in cognitive demand while retaining online updating and uncertainty quantification.

### B. Revised Bloom Taxonomy and Hierarchical Dependency

Bloom's revised taxonomy (RBT) distinguishes a cognitive-process dimension (Remember, Understand, Apply, Analyze, Evaluate, and Create) and a knowledge-type dimension (Factual, Conceptual, Procedural, and Metacognitive), supporting constructive alignment between objectives, instruction, and assessment evidence [3], [28]. However, classification validity may weaken when cognitive level is inferred from surface cues such as verbs rather than from the actual mental operations required by the task [30]. Assessment instruments often fail to reflect the intended balance of cognitive demand: examinations may become dominated by lower-level recall [35], or higher-order demands may increase in ways that raise concerns about fairness [53].

Deeper understanding is better captured when assessment design includes integrative and evaluative demands [34]. This is especially consequential for rating-based learner models, where items are reduced to a single difficulty parameter that may obscure meaningful differences in cognitive demand.

For the present work, RBT is not treated as a rigid staircase but as a plausible dependency structure in which higher-order performance is generally more likely when foundational representations are sufficiently established. This interpretation is consistent with classic validation work suggesting cumulative relations across broad groupings rather than strictly between every adjacent pair [29], and with more recent evidence showing that higher-level objectives may depend on lower ones through indirect pathways [60]. From a cognitive perspective, this is plausible because more complex reasoning places greater demands on working memory and becomes more feasible when prior knowledge is organized into schemas that reduce cognitive load [50], [45].

### C. Uncertainty-Aware Progression in Rating-Based Systems

Elo-style rating systems are widely used in adaptive educational systems because they update learner and item parameters online with modest computational cost [24], [39],

[36], [2], [40]. However, standard Elo represents learner state as a single point estimate without an explicit uncertainty parameter, which can produce unstable early estimates and, in adaptive settings that jointly update learners and items, rating variance that inflates over time [7]. These models also generally ignore the lag between successive interactions, treating adjacent observations as if no learning or forgetting occurred between them [12], [56], [10], [38].

These limitations motivate uncertainty-aware approaches that attach reliability information to ability estimates [21]. One strategy replaces the constant update factor with a decreasing uncertainty function that reduces rating sensitivity as interactions accumulate [2], [57]. The Glicko rating system formalizes this idea by introducing rating deviation (RD) as an explicit uncertainty parameter: high RD permits larger updates under limited evidence, inactivity, or cold start, whereas decreasing RD stabilizes the estimate and supports more conservative adaptive decisions [17].

However, the deeper limitation of Elo and Glicko variants in cognitively structured domains is not merely technical but representational. These systems originate from paired-comparison rating contexts in which the primary objective is to estimate relative performance or predict outcomes. When transferred to education, however, the rating system must support a different kind of decision: whether the learner is pedagogically ready to progress. This requires more than adding educational labels to a generic rating scale. It requires a learner model whose structure reflects educational constraints, including cognitive demand, prerequisite stability, and uncertainty-aware progression.

Recent extensions have introduced multivariate Elo and Glicko representations that estimate separate ratings for knowledge components [1], [2]. These approaches improve diagnostic resolution, but they retain a primarily estimation-oriented design: the multiple dimensions represent knowledge components whose ratings are updated separately, and in some cases cross-informed, rather than an ordered cognitive hierarchy in which advancement is governed by explicit readiness conditions. Thus, when hierarchy is present, it mainly supports parameter estimation rather than constraining progression decisions. This distinction is central to the present work: HBC-Glicko is not simply Bloom-aware estimation, but Bloom-constrained progression.

In scalar Elo and Glicko variants, learner state is encoded as a single rating, so the system has no internal mechanism to detect whether strong performance at one cognitive level is masking weakness at another. In a Bloom-structured formative setting, this can permit advancement on the basis of aggregated performance while foundational knowledge remains unstable, producing a progression decision that is psychometrically plausible but pedagogically incoherent.

#### D. Binary and Continuous Evidence in Adaptive Assessment

The representation of learner evidence should reflect the cognitive complexity of the task being assessed. For lower Bloom levels, particularly Remember and Understand, binary scoring (correct/incorrect) is computationally efficient and aligns with conventional objective formats. However, although

multiple-choice questions can theoretically assess higher-order thinking, empirical work shows they often concentrate on lower-order demands in practice [42], [52], [48], [64], [33]. Higher Bloom levels require richer evidence because responses involve reasoning, justification, or production that cannot be reduced to binary judgments. Rubric-based scoring addresses this by assessing performance against explicit criteria and levels of quality [23], [8].

Rubric-based assessment poses a scalability challenge in formative contexts where consistent human scoring of large response volumes becomes impractical. Large language models offer a promising solution: recent studies suggest that LLMs can support rubric-based scoring at scale with acceptable reliability for formative use [11], [14], [16], [47]. These considerations support a mixed evidence design, with binary scoring for lower Bloom levels and rubric-based continuous evidence for higher-order cognitive performance.

Taken together, the preceding sections suggest that adaptive progression in hierarchical cognitive domains cannot be governed by scalar performance estimates alone. The system must account for hierarchical dependency, uncertainty in learner estimates, and differences in evidence type across task formats. These requirements motivate an explicit control logic that regulates advancement conservatively rather than treating progression as a simple consequence of observed success. The next section presents the system architecture and formal control design of the proposed HBC-Glicko system.

### III. SYSTEM ARCHITECTURE OVERVIEW

To clarify how the proposed components interact, this section presents the overall system architecture of HBC-Glicko (Fig. 1). The architecture provides a high-level view of the information flow from the calibrated, Bloom-tagged item bank to item selection, scoring, rating updates, and progression decisions. This overview is intended to orient the reader before the formal definitions and control rules introduced in the following sections.

#### A. Operational Deployment Flow

To make the architecture more operationally explicit, Fig. 2 presents a practical deployment cycle for HBC-Glicko in an LMS-supported formative assessment context. The cycle begins with LLM-assisted item generation and Bloom-label suggestion, followed by instructor review and confirmation to ensure that each item is pedagogically appropriate and correctly assigned to a cognitive band and category. Validated items then enter a calibrated Bloom-banded item pool with anchor sets used for progression thresholds. During assessment, the system routes the learner within the current band, selects and scores items, and commits band-specific rating updates at the session level. Progression decisions are then made through confidence-bound rules: the learner is promoted only when readiness is sufficiently stable, whereas prerequisite reinforcement is triggered when current-band performance falls below the expected entry standard. This deployment view clarifies that HBC-Glicko functions not as an autonomous replacement for instructional judgment, but as a supervised adaptive layer that combines

LLM-supported item preparation, instructor validation, and uncertainty-aware progression control.

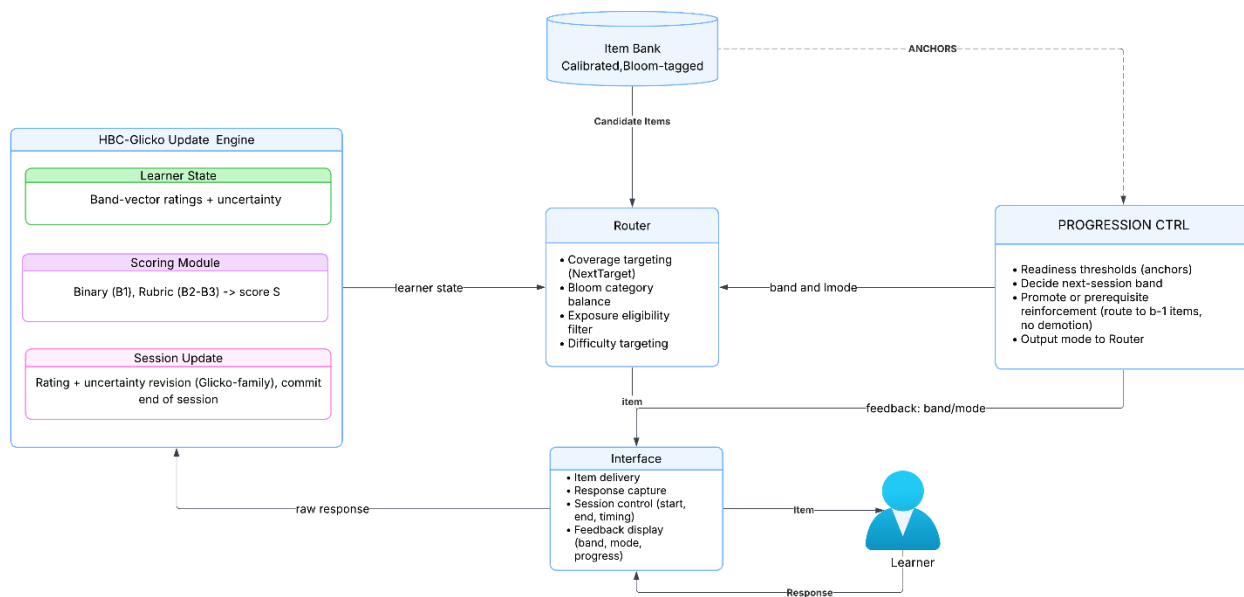


Fig. 1. Overview of the HBC-Glicko architecture.

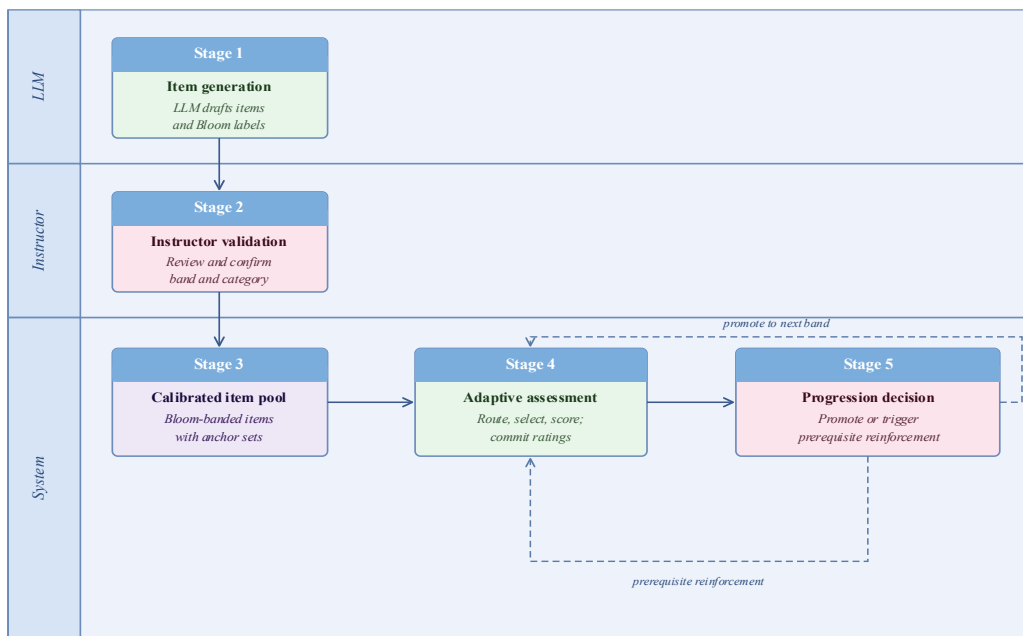


Fig. 2. HBC-Glicko deployment cycle. Dashed arrows indicate conditional system-level transitions after Stage 5.

The proposed system integrates five components, each developed in the following subsections: the formal learner model (Section III-A), within-band adaptation (Section III-B), anchor-based progression thresholds (Section III-C), hierarchical decision rules for promotion and prerequisite reinforcement (Section III-D), and inactivity-driven uncertainty growth (Section III-E).

1) *Formal learner measurement model:* This section specifies the formal representation of a learner within the

proposed adaptive measurement architecture. The learner is defined as a structured latent state with band-specific performance parameters aligned with Bloom-type cognitive bands. The purpose is to define 1) the learner state variables, 2) their uncertainty representation, and 3) the structural assumptions that prevent compensatory masking across bands. This section provides definitions only. Update equations, thresholds, and decision rules are specified in subsequent sections.

For reference, a complete summary of the notation used throughout the following sections is provided in Appendix A (Table III).

2) *Band set and indexing*: Let  $n$  index learners,  $t$  index rating periods (sessions), and  $b \in \{1, 2, 3\}$  index cognitive bands. The bands represent increasing levels of cognitive complexity, following the revised Bloom taxonomy [3]. Although the revised taxonomy distinguishes two dimensions — cognitive processes and knowledge types — this system adopts the cognitive-process dimension only, consistent with common practice in empirical assessment research [59], [35], [5].

Low band:  $\mathcal{K}_{\text{Low}} = \{\text{Remember, Understand}\}$ , Mid band:  $\mathcal{K}_{\text{Mid}} = \{\text{Apply}\}$ , High band:  $\mathcal{K}_{\text{High}} = \{\text{Analyze, Evaluate}\}$ .

The conventional distinction in educational practice divides these categories into lower-order thinking skills (Remember, Understand, and Apply) and higher-order thinking skills (Analyze, Evaluate, and Create) [28]. While useful for curricular planning, this two-tier partition does not distinguish between retrieval, comprehension, and contextual transfer, which are diagnostically different in adaptive assessment. The present architecture groups the categories according to the type of cognitive operation involved. The Low-band combines Remember and Understand as representational processes in which the learner accesses or restates stored knowledge. The Mid-band isolates apply as the transition from accessing knowledge to using it under unfamiliar conditions. The High band combines Analyze and Evaluate as processes that require coordinating and judging multiple knowledge elements simultaneously, imposing high working memory demand [49]. This partition captures three qualitatively distinct modes of cognitive engagement while preserving sufficient evidence per band for stable rating estimation in short formative sessions.

Create is excluded because it introduces constraints that differ in kind from the other levels. Creative production tasks require the learner to generate an original artifact rather than respond to a structured prompt, which means that item difficulty cannot be stably pre-calibrated on the unified rating scale: the effective demand of a Create task depends on the scope and ambition of the learner's response, not solely on fixed item properties. Additionally, while LLM-assisted scoring is used in higher bands for evaluating reasoning quality against explicit rubric criteria, scoring creative output requires additional appraisal of novelty and integrative originality, dimensions for which automated scoring consistency remains less established [8]. Incorporating Create would therefore require modifications to both the item calibration model and the scoring framework, constituting a substantive extension left for future work.

Each item in the pool carries two fixed labels: a band ( $b$ ) and a Bloom category within that band. These labels are assigned through a two-stage process. First, an LLM classifier suggests both the band and the category. Second, the instructor reviews and confirms the suggestion before the item enters the pool. Once confirmed, neither label is changed by the system at runtime.

This dual labeling serves two purposes. The band label supports the non-compensatory structure of the learner model. The within-band category label supports the within-band coverage mechanism.

3) *Learner state variables*: For each learner  $n$  at the rating period  $t$ , the latent performance state is represented by  $\mathbf{r}_{n,t} = (r_{n,1,t}, r_{n,2,t}, r_{n,3,t})$ , where  $r_{n,b,t} \in \mathbb{R}$  denotes performance in band  $b$ . Each estimate is associated with a rating deviation  $RD_{n,b,t} > 0$ , which captures uncertainty. The learner measurement state is

$$\mathcal{S}_{n,t} = \{(r_{n,b,t}, RD_{n,b,t}) : b \in \mathcal{B}\}.$$

Consistent with the theoretical rationale developed in Section II, learner state is modeled as a vector rather than as a single scalar proficiency parameter. This choice reflects the view that performance depends on prior knowledge organization and on the cognitive demands imposed by different tasks, rather than on a single context-free ability alone [49], [45]. High performance in a higher band, therefore, cannot substitute for insufficient performance in a lower band. A non-compensatory representation is retained to preserve pedagogical coherence in progression decisions and to prevent apparent success on more complex tasks from masking weakness in foundational knowledge.

4) *Unified scale assumption*: All band ratings are expressed on a common continuous metric scale. This enables coherent definition of item difficulty parameters, readiness thresholds, and uncertainty bounds on a single axis, while preserving band separation through the evidence structure.

Let  $d_i \in \mathbb{R}$  denote the difficulty (location) parameter of item  $i$  on the same scale as  $r_{n,b,t}$ . The architecture assumes that expected performance is a monotonic function of the rating difference  $r_{n,b,t} - d_i$ . The explicit response function and update rules are provided in Section III-B. Here, the key assumption is scale compatibility: learner and item parameters share a metric that supports probabilistic interpretation.

## B. Within-Band Adaptation

This section specifies the within-band probabilistic update mechanism. Each cognitive band operates as an independent rating process governed by a Glicko-style uncertainty-aware model [18], [19]. Interactions with items assigned to the band  $b$  update only the learner parameters ( $r_{n,b,t} | RD_{n,b,t}$ ) associated with that band and the parameters of the interacted item.

1) *Response model*: Consider an item  $i$  in band  $b$ , with difficulty  $d_{i,t} \in \mathbb{R}$  and rating deviation  $RD_{i,t} > 0$ . Let  $r_{n,b,t}$  denote learner  $n$ 's rating in the band  $b$ . Following the Glicko formulation, define.

$$g(RD_{i,t}) = \frac{1}{\sqrt{1 + \frac{3q^2 RD_{i,t}^2}{\pi^2}}}, q = \frac{\ln(10)}{400}$$

The expected score of the learner  $n$  on item  $i$  at time  $t$  is

$$E_{n,i,t} = \frac{1}{1 + 10^{-\frac{g(RD_{i,t})(r_{n,b,t} - d_{i,t})}{400}}}$$

Observed performance is represented on a common scale  $S_{n,i,t} \in [0,1]$ . This common scale is retained to preserve comparability across bands while allowing the observable response format to remain aligned with the cognitive demands of the task. In Band 1, where tasks primarily target retrieval and foundational understanding, responses are binary, with  $S_{n,i,t} \in \{0,1\}$ , so  $E_{n,i,t}$  corresponds to the probability of a correct response. In Bands 2 and 3, where performance depends more on the quality of application, analysis, justification, or evaluation, rubric-based scores are normalized to  $[0, 1]$ . In those bands,  $E_{n,i,t}$  is interpreted as the expected performance level on the same scale.

This mixed-evidence design is intended to ensure that the granularity of scored evidence remains congruent with the level of cognitive complexity being assessed. Lower-band performance can often be adequately captured through correctness, whereas higher-band performance requires attention to gradations in the quality of reasoning and knowledge use. The purpose of this design is therefore pedagogical as well as technical. By matching the scoring format to the nature of the targeted performance, the system can support formative interpretations that distinguish between foundational correctness and higher-order quality of reasoning, which is important for making progress intelligible to learners across increasing levels of task complexity [37], [54].

The prediction error is then defined as

$$\delta_{n,i,t} = S_{n,i,t} - E_{n,i,t}.$$

This quantity represents the deviation between observed and expected performance and serves as the input to the subsequent update step.

2) *Learner rating update*: Within band  $b$ , the learner rating is updated as

$$r_{n,b,t+1} = r_{n,b,t} + \omega_{n,b,t} \delta_{n,i,t}$$

where,  $\omega_{n,b,t} > \mathbf{0}$  denotes a Glicko-style adaptive weight. The update magnitude increases with learner uncertainty in the band  $b$ , decreases with item uncertainty through  $g(RD_{i,t})$ , and depends on the interaction-information term  $E_{n,i,t}(1 - E_{n,i,t})$ , which reaches its maximum when the predicted outcome is most uncertain. Band updates are strictly independent, such that an interaction in band  $b$  affects only the learner state  $(r_{n,b,t}, RD_{n,b,t})$  for that band and the associated item parameters.

3) *Within-band category routing and item selection*: Within a session, the active band  $b \in \{\text{Low, Mid, High}\}$  remains fixed. The routing mechanism operates only within the current band and determines, at each step  $t$ , which in-band Bloom category should be sampled next.

a) *Phase 1: Target Category Determination*: Phase 1 determines the next category to be sampled within the active band. To guide this decision, the instructor specifies a target sampling distribution over the categories in  $\mathcal{K}_b$ , denoted  $\pi_b(k)$ ,

such that  $\pi_b(k) \geq 0$  and  $\sum_{k \in \mathcal{K}_b} \pi_b(k) = 1$ . By default, this distribution is uniform, although it may be adjusted to reflect pedagogical priorities. During the session, the system maintains, for each category  $k \in \mathcal{K}_b$ , a delivery counter  $c_t(k)$  representing the number of scored items delivered up to step  $t$  in band  $b$ , with total in-band exposure given by  $t_b = \sum_{k \in \mathcal{K}_b} c_t(k)$ . The expected count for the category  $k$  under the target distribution is  $\hat{c}_t(k) = t_b \pi_b(k)$ , and the corresponding coverage deficit is defined as  $\Delta_t(k) = \max(0, \hat{c}_t(k) - c_t(k))$ . The next target category is then selected as the most under-sampled category,  $k_t = \arg \max_{k \in \mathcal{K}_b} \Delta_t(k)$ . In the Mid band, where  $\mathcal{K}_{\text{Mid}} = \{\text{Apply}\}$ , this step is trivial, since  $k_t = \text{Apply}$  at every step. If multiple categories have the same deficit, priority is given to the category with the larger remaining item pool, and any further ties are resolved deterministically to ensure reproducibility.

b) *Phase 2: Adaptive Item Selection Within the Target Category*: Once the target category  $k_t$  has been determined, the router forms the candidate set  $I_t^{\text{cand}} = \{i : b(i) = b, k(i) = k_t\}$ . Because  $k_t$  is selected in Phase 1, this set is already aligned with the remaining within-band category demand. To preserve measurement independence within a session, only the first scored attempt on an item is allowed to update ratings. Let  $H_s$  denote the set of items already scored in the session  $s$ . The eligible set is then  $I_t^{\text{elig}} = \{i \in I_t^{\text{cand}} : i \notin H_s\}$ . If  $I_t^{\text{elig}} = \emptyset$ , secondary constraints may be relaxed in a deterministic manner while preserving the band constraint. This rule is motivated by a conservative cognitive rationale: repeated responses to the same item within a session may still draw on information that remains readily accessible in working memory, whereas a first scored attempt in a later session is more plausibly treated as a new retrieval event from long-term memory [4], [13]. This distinction is also consistent with retrieval-practice research showing that immediate repetition and later retrieval do not have the same implications for retention [43], [25]. Within the eligible set, the next item is selected using a probability-targeting principle. Let  $E_{n,i,t}$  denote the expected performance of the learner on item  $i$  at step  $t$ , and let  $p^* \in (0,1)$  denote the target performance level. In Band 1,  $E_{n,i,t}$  is the probability of a correct response under a binary outcome model. In Bands 2 and 3, where scores are rubric-based and normalized to  $[0,1]$ , it represents the expected performance level. The selected item is therefore  $i_t^* = \arg \min_{i \in I_t^{\text{elig}}} |E_{n,i,t} - p^*|$ . If several items are equally close to  $p^*$ , the tie is broken using the information function  $Info(i) = g(RD_{i,t})^2 E_{n,i,t}(1 - E_{n,i,t})$ , and the item with the highest value is selected.

4) *End-of-session rating commitment*: Although provisional parameter adjustments may be computed during a session for routing purposes, official rating updates are committed only at the end of the session. All scored interactions from the session are aggregated, and the band-specific learner parameters  $(r_{(n,b)}, RD_{(n,b)})$  are updated from their pre-session values using the complete set of session evidence.

This separation serves both a statistical and a pedagogical purpose. Statistically, it avoids sequential double counting

within the same session and preserves coherence in parameter estimation. Pedagogically, it reduces the risk that transient within-session fluctuations are overinterpreted as stable change in learner state. By committing ratings only after the full session has been observed, the model treats learner performance as session-level evidence rather than as a sequence of isolated micro-updates, which is more consistent with the formative aim of basing progression decisions on sufficiently consolidated patterns of performance rather than on momentary variation.

### C. Anchor-Based Progression Thresholds

1) *Entry anchor sets and reference difficulty definition:* In educational measurement, anchor items are used to preserve scale stability across forms and administrations by maintaining item and ability estimates on a shared metric [26], [55]. The present architecture adopts this principle to define stable progression thresholds across Bloom bands. For each transition from band  $b$  to band  $b + 1$ , an entry anchor set  $\mathcal{A}_{b+1} \subset \mathcal{I}_{b+1}$  is defined as a fixed subset of calibrated items drawn exclusively from band  $b + 1$ , where  $\mathcal{I}_{b+1}$  denotes the item pool associated with that band. Each anchor item  $j \in \mathcal{A}_{b+1}$  is characterized by a calibrated difficulty parameter  $d_{b+1,j}$  on the unified Glicko scale, so that anchor difficulties share the same latent metric as learner band-specific ratings. Because these anchor sets are predefined for progression purposes and are not dynamically altered by session-level adaptation, they serve as stable structural reference points for entry-level demand in the target band. The reference difficulty associated with entry into band  $b + 1$  is defined as the mean anchor difficulty,

$$d_{b+1}^{ref} = \frac{1}{|\mathcal{A}_{b+1}|} \sum_{j \in \mathcal{A}_{b+1}} d_{b+1,j}$$

This quantity represents the typical calibrated difficulty of entry-level tasks in band  $b + 1$  and provides a fixed scale location against which readiness criteria are defined. In pedagogical terms, it operationalizes the level of demand that a learner should be prepared to meet before progression is warranted. To preserve the interpretability of progression thresholds over time, anchor item parameters are treated as fixed after calibration and excluded from routine item-parameter updating. In this way, anchor-based reference locations remain effectively stable on the unified rating scale, ensuring that subsequent readiness decisions retain a consistent measurement meaning.

2) *Probability-defined readiness thresholds:* Progression from band  $b$  to band  $b + 1$  is defined in probabilistic terms. Let  $p_{b \rightarrow b+1} \in (0,1)$  denote the target expected success probability required for entry into band  $b + 1$ . This parameter specifies the minimum expected performance required on representative entry-level tasks of the target band.

Readiness is defined relative to the reference difficulty  $d_{b+1}^{ref}$  established above. Let  $r_{n,b}$  denote learner  $n$ 's ability estimate in band  $b$ . On the unified rating scale, the expected outcome of the learner on an item of difficulty  $d$  is given by

$$E(r_{n,b}, d) = \frac{1}{1 + 10^{-\frac{r_{n,b} - d}{400}}}$$

The readiness criterion requires expected performance at the target-band reference difficulty to equal the designated transition probability:

$$E(r_{n,b}, d_{b+1}^{ref}) = p_{b \rightarrow b+1}$$

Solving for the rating value that satisfies this condition yields the progression threshold

$$T_{b \rightarrow b+1} = d_{b+1}^{ref} + 400 \log_{10} \left( \frac{p_{b \rightarrow b+1}}{1 - p_{b \rightarrow b+1}} \right)$$

Thus,  $T_{b \rightarrow b+1}$  represents the minimum band-  $b$  rating required for a learner to achieve the target expected success probability on calibrated entry-level tasks from band  $b + 1$ . Progression is therefore defined as a predictive readiness condition on the unified scale rather than as a retrospective summary of past performance. This predictive interpretation is consistent with cognitivist accounts of learning in which successful future performance depends on the interaction between prior knowledge organization and the cognitive demands of upcoming tasks, rather than on past score accumulation alone [49], [45].

In contrast to standard CAT item selection, which typically targets items near the current ability estimate to maximize measurement information, progression thresholds in the HBC system are designed to express readiness for cognitively more demanding work. Accordingly, transition probabilities are specified above the chance level and increase across bands:

$$0.50 < p_{Low \rightarrow Mid} < p_{Mid \rightarrow High} < 1.$$

This reflects the design assumption that progression to higher cognitive bands should require stronger evidence of readiness. Pedagogically, this means that advancement into more demanding bands is treated as a guarded instructional decision, since the cost of premature progression increases as task complexity and knowledge coordination demands rise [49]. The specific values of  $p_{b \rightarrow b+1}$  are treated as design parameters to be calibrated empirically rather than fixed a priori.

### D. Confidence-Bound Decision Rules for Band Progression

1) *Rating uncertainty and confidence bounds:* To incorporate estimation uncertainty into band progression, lower and upper uncertainty bounds are defined as

$$LCB_{n,b} = r_{n,b} - z_b RD_{n,b}, \quad UCB_{n,b} = r_{n,b} + z_b RD_{n,b}$$

where,  $z_b > 0$  controls the width of the interval. The lower bound provides a conservative estimate of learner ability in band  $b$ , whereas the upper bound provides a corresponding optimistic estimate. As  $RD_{n,b}$  increases, the interval widens, which discourages progression under high uncertainty. As  $RD_{n,b}$  decreases, the interval narrows, allowing readiness to be evaluated more decisively. This formulation makes readiness conditional not only on the estimated level, but also on the stability of the evidence supporting that estimate.

The multiplier  $z_b$  is set to 1.65, corresponding to a one-sided 95% confidence level under the distributional assumptions of

the Glicko framework [18]. Promotion therefore requires 95% confidence that the learner's true ability exceeds the readiness threshold, whereas reinforcement is triggered only when even a 95th percentile estimate falls below the entry standard. This moderately conservative threshold is consistent with the formative orientation of the architecture. A stricter bound, such as  $z_b = 2.00$ , would systematically delay promotion for learners who have already demonstrated sufficient mastery, moving them away from their zone of proximal development without proportionally reducing misclassification risk. The architecture can tolerate a less restrictive threshold because prerequisite reinforcement [Section III-D(3)] provides a built-in safeguard. If promotion proves premature, the system can detect instability and redirect the learner to lower band items. In summative contexts, where no such remediation loop exists, a stricter value would be more appropriate. More generally,  $z_b$  is treated as a tunable design parameter whose optimal value should be calibrated empirically for each deployment context. Although a uniform multiplier is adopted here for simplicity,  $z_b$  could in principle be differentiated across Bloom bands. A stricter bound at higher cognitive levels would reflect the greater pedagogical cost of premature advancement into more demanding tasks, with band-specific values determined through pilot calibration against observed promotion rates and downstream performance outcomes. The relationship between  $RD$  and the progression threshold is structural. Because  $LCB_{n,b} = r_{n,b} - z_b \cdot RD_{n,b}$ , higher uncertainty directly lowers the conservative readiness estimate and makes the promotion condition  $LCB_{n,b} \geq T_{b \rightarrow b+1}$  harder to satisfy. A learner with a high point estimate but high  $RD_{n,b}$  is therefore treated as less ready than a learner with the same point estimate and lower uncertainty. This is pedagogically intentional: promotion requires not only high estimated performance but also stable evidence within the prerequisite band. Formally,  $T_{b \rightarrow b+1}$  is not a free parameter, but a derived quantity computed from  $d_{b+1}^{ref}$  and  $p_{b \rightarrow b+1}$  according to the anchor-based threshold construction in Section III-C.

2) *Promotion rule: conservative and coverage-constrained*: Let  $T_{b \rightarrow b+1}$  denote the readiness threshold for progression from band  $b$  to band  $b + 1$ , and let  $LCB_{n,b}$  denote the lower confidence bound of learner  $n$ 's rating in band  $b$ . Promotion is granted only when the learner satisfies the conservative readiness condition.

$$\text{Promote } b \rightarrow b + 1 \iff LCB_{n,b} \geq T_{b \rightarrow b+1}$$

To ensure that progression is not driven by performance in only a subset of within-band categories, promotion is additionally subject to a minimum evidence requirement over the category set  $\mathcal{K}_b$ . Let  $N_{n,b}(k)$  denote the number of distinct scored item administrations obtained by learner  $n$  in category  $k \in \mathcal{K}_b$ , and let  $m_b$  denote the minimum required evidence per category. The full promotion rule is therefore

$$\begin{aligned} \text{Promote } b \rightarrow b + 1 \\ \iff (LCB_{n,b} \geq T_{b \rightarrow b+1}) \\ \wedge (\forall k \in \mathcal{K}_b: N_{n,b}(k) \geq m_b) \end{aligned}$$

This rule makes progression both conservative and non-compensatory. The confidence-bound condition reduces the risk

of premature advancement under high uncertainty, while the coverage condition ensures that promotion reflects sufficient evidence across all categories within the band rather than strong performance in only a limited subset.

3) *Prerequisite reinforcement as a stability safeguard*: Although promotion is irreversible, subsequent responses may indicate that the learner is not performing stably at the level expected for the current band. To detect such cases conservatively, the system uses the upper confidence bound of the band-specific rating. Prerequisite reinforcement is triggered when

$$UCB_{n,b} < T_{b-1 \rightarrow b}$$

where,  $T_{b-1 \rightarrow b}$  denotes the entry threshold for band  $b$ . This condition indicates that, even under an optimistic estimate, the learner does not meet the readiness standard for the current band. When this occurs, the system temporarily prioritizes items from band  $b - 1$  to reinforce prerequisite knowledge. However, the learner's structural band membership remains unchanged, and no demotion is performed. Because  $UCB_{n,b}$  increases with  $RD_{n,b}$ , reinforcement is less likely to be triggered when uncertainty is high, which reduces sensitivity to sparse or noisy evidence.

Together, the promotion and reinforcement rules define three decision zones. Promotion occurs when  $LCB_{n,b} \geq T_{b \rightarrow b+1}$ , subject to the coverage condition. Prerequisite reinforcement occurs when  $UCB_{n,b} < T_{b-1 \rightarrow b}$ . Otherwise, the learner remains in band  $b$  without structural change.

### E. Inactivity-Driven Uncertainty Growth

To reflect increasing uncertainty when no recent evidence is observed, rating deviation is allowed to grow during periods of inactivity. For learner  $n$  in band  $b$ , let  $t_{n,b}$  denote the number of completed sessions since the learner last produced at least one scored attempt in that band. The effective pre-update uncertainty is defined as:

$$RD_{n,b}^{pre*} = \min \left( RD_{\max}, \sqrt{(RD_{n,b}^{pre})^2 + c_b^2 t_{n,b}} \right)$$

where,  $c_b > 0$  controls the rate of uncertainty growth under inactivity and  $RD_{\max}$  is the maximum admissible rating deviation. This formulation allows uncertainty to increase as previously observed evidence becomes less informative over time, while preventing unbounded growth.

Because Band 1 is oriented toward retention-sensitive recall, the informativeness of prior evidence may decay more rapidly there than in higher cognitive bands [15], [46], [44], [27]. Accordingly, the inactivity-growth parameters are constrained such that

$$c_1 \geq c_2 \geq c_3,$$

so that uncertainty can grow at least as quickly in Band 1 as in Bands 2 and 3. If no new evidence is observed in band  $b$ , the band rating remains unchanged, and only the uncertainty term is updated through the inactivity-growth rule.

IV. ARCHITECTURAL PROPERTIES OF HBC-GLICKO

Table I summarizes the main theoretical properties that follow from the HBC-Glicko design. These properties are structural consequences of the learner-state representation,

threshold formulation, confidence-bound decision rules, and inactivity update mechanism. They clarify what the architecture guarantees by construction, without implying empirical validation.

TABLE I. MAIN THEORETICAL PROPERTIES OF THE HBC-GLICKO ARCHITECTURE

Property	Architectural basis	Theoretical implication
Non-compensatory progression	Learner state is represented as a band-specific vector, and promotion is evaluated separately within each prerequisite band.	Strong performance in one band cannot compensate for insufficient evidence in another band.
Conservative promotion under uncertainty	Promotion depends on the lower confidence bound rather than the point estimate alone.	For the same rating estimate, greater uncertainty reduces effective readiness and makes promotion less likely.
Criterion-referenced progression	Promotion thresholds are defined by fixed anchor difficulties and target transition probabilities.	Readiness is judged against a common formal standard rather than a session-specific rule.
Asymmetric decision rules	Promotion uses the lower confidence bound, whereas prerequisite reinforcement uses the upper confidence bound.	The system is conservative at both boundaries: advancement requires sufficiently strong evidence, while remediation requires sufficiently weak evidence.
Inactivity-dependent uncertainty	Inactivity increases rating deviation without directly changing the stored rating.	After long inactivity, progression decisions become more cautious until new evidence is collected.

V. DISCUSSION

The central contribution of HBC-Glicko is to reposition Bloom's revised taxonomy from a descriptive framework to a control principle within adaptive assessment. In many adaptive systems, the cognitive level is used to classify items or structure reporting, but it does not directly govern when a learner may advance. In the present architecture, Bloom-structured bands are embedded in the learner model itself, and progression depends on whether the learner has accumulated sufficiently strong and sufficiently stable evidence within the prerequisite band. Scalar learner models absorb performance across cognitively different

tasks into a single estimate, allowing strength in one area to mask instability in another.

HBC-Glicko prevents this by representing learner state as a non-compensatory band vector and evaluating readiness within bands rather than across them [1], [62].

Table II situates HBC-Glicko relative to selected adaptive learner modeling approaches. None of the compared systems combines multidimensional Bloom-structured representation, explicit uncertainty quantification, and hierarchy used as a non-compensatory control constraint on progression within a single architecture.

TABLE II. COMPARATIVE POSITIONING OF HBC-GLICKO RELATIVE TO SELECTED ADAPTIVE LEARNER MODELING APPROACHES

Feature	Standard Elo	Glicko	Multivariate Elo	Multivariate Glicko	CAT / UIRT	Skill-hierarchy Elo	Bloom-aware Elo	HBC-Glicko
	[39]	[17], [18]	[2]	[1]	[62]	[63]	[20]	(proposed)
Learner model type	Scalar	Scalar	Vector (KC)	Vector (KC)	Scalar	Multi-level (global per-skill) +	Vector (Bloom)	Vector (Bloom bands)
Explicit uncertainty parameter	No	Yes (RD)	Not explicit	Yes (RD per KC)	Yes (SEM)	No	No	Yes (RD per band)
Hierarchy represented	No	No	Skill components	Skill components	No	Skill hierarchy	Bloom levels	Bloom bands (ordered)
Hierarchy as control constraint	No	No	No	No	No	Estimation only	Estimation only	Yes
Non-compensatory progression	No	No	No	No	No	No	No	Yes
Anchor-based transition thresholds	No	No	No	No	No	No	No	Yes
Confidence-bounded progression rule	No	No	No	No	No	No	No	Yes
Explicit prerequisite reinforcement	No	No	No	No	No	No	No	Yes
Evidence differentiated by cognitive band	No	No	No	No	No	No	No	Yes

**Note.** KC = knowledge component. RD = rating deviation. SEM = standard error of measurement. UIRT = unidimensional item response theory. “Estimation only” indicates that the hierarchy informs parameter estimation or cross-level propagation but does not function as a constraint on learner progression. “Not explicit” indicates that the feature may be partially present or achievable through indirect mechanisms but is not formalized as a distinct architectural component. “Indirect (SEM)” reflects that CAT quantifies estimation precision through the standard error of measurement, which serves a different function from an explicit uncertainty parameter governing progression decisions.

A second contribution is that uncertainty is given a genuinely pedagogical function. In conventional Glicko, rating deviation determines update magnitude. In HBC-Glicko, uncertainty also determines what the system should do next: promotion depends on a lower-confidence interpretation of readiness, whereas prerequisite reinforcement is triggered only when performance remains weak even under an optimistic interpretation. This asymmetry makes the architecture deliberately conservative at transitions between cognitive bands, where instructional errors are most costly. Uncertainty does not merely slow or accelerate adaptation but directly constrains progression, extending its role from parameter updating to pedagogical decision-making.

To our knowledge, Hamzah and Sosnovsky (2023) are the only published work to directly combine a rating-based learner model with Bloom's revised taxonomy, and no prior work has combined a Glicko-based system with Bloom's taxonomy as a non-compensatory progression constraint. Yudelson et al. (2019) demonstrate that Elo can exploit skill hierarchies for few-shot estimation, and Abdi et al. [1] extend Glicko to multivariate knowledge components, but neither structure cognitive-process levels as a non-compensatory constraint on the learner model. In Hamzah and Sosnovsky (2023), Bloom levels organize learner modeling through cross-level propagation, so evidence at one level can influence estimates at another. In HBC-Glicko, bands are treated as non-compensatory decision spaces, and readiness for band  $b + 1$  must be established from within band  $b$  itself. The distinction is therefore between Bloom-aware estimation and Bloom-constrained progression: the former improves the structure of the learner model; the latter governs advancement. For formative assessment, where sequencing and readiness are central, this is the more consequential step.

HBC-Glicko is not novel simply because it is multidimensional, since multivariate rating models already exist, nor because it incorporates uncertainty, since that is a defining feature of Glicko-based systems [1]. Its contribution lies in combining multidimensionality and uncertainty with a cognitively ordered progression rule. The dimensions are not parallel knowledge components estimated for predictive accuracy but ordered bands that participate directly in advancement logic, making the architecture more interpretable pedagogically than a generic multivariate model.

The contrast with computerized adaptive testing is equally important. CAT selects items that maximize information at the current ability estimate and achieves acceptable precision with few observations [62]. HBC-Glicko asks whether the learner has demonstrated enough stable evidence at the current cognitive

band to justify movement to the next. This makes the architecture less efficient psychometrically but more coherent in formative contexts where progression should respect hierarchical cognitive dependency. HBC-Glicko should not be positioned as a replacement for CAT but as an alternative design philosophy for settings in which instructional ordering matters as much as measurement efficiency.

These advantages remain theoretical. The practical value of the architecture will depend on the quality of Bloom-level classification, the stability of anchor calibration, and the appropriateness of parameter choices governing thresholds, confidence penalties, and inactivity growth. The architecture nevertheless provides a formally specified and internally coherent foundation for future empirical work, with progression rules whose theoretical properties have been clarified in Section IV and whose design rationale is grounded in both measurement theory and cognitive theory. The value of HBC-Glicko will ultimately depend on how well those theoretical advantages translate into better progression decisions in practice.

#### Parameter Sensitivity and Calibration of Decision Behavior.

The proposed architecture makes progression decisions explicit by formalizing them through anchor-derived thresholds, confidence bounds, category-coverage requirements, and inactivity-driven uncertainty growth. This explicitness introduces parameter sensitivity, but it also constitutes an important strength of the model: the strictness of progression is no longer implicit, hidden in an opaque adaptive algorithm, or determined only by aggregated performance scores. Instead, it is represented through interpretable control parameters that can be inspected, adjusted, and empirically calibrated. The transition probability  $p_{b \rightarrow b+1}$ , the confidence multiplier  $z_b$ , the minimum evidence requirement  $m_b$ , and the inactivity-growth parameter  $c_b$  collectively determine the balance between adaptive responsiveness and instructional caution.

Misconfiguration of these parameters may shift the system toward predictable forms of decision error. If thresholds or confidence bounds are too conservative, HBC-Glicko may delay promotion for learners who are already capable of engaging with higher-order tasks. If they are too permissive, the system may allow premature progression before prerequisite competence is sufficiently stable. Similarly, an overly strict category-coverage requirement may produce unnecessary repetition, whereas an overly weak requirement may permit advancement from narrow evidence. Inactivity-growth parameters also require careful calibration: excessive growth may discount prior evidence too rapidly, while insufficient growth may cause the system to overtrust outdated performance. These risks do not undermine the architecture; rather, they indicate that HBC-Glicko should be treated as a tunable decision framework whose operating characteristics must be calibrated to the instructional context.

Future simulation and empirical work should therefore examine how alternative parameter settings affect promotion accuracy, reinforcement frequency, time-to-progression, learner workload, and downstream performance. Such analyses would make it possible to identify robust operating ranges and determine whether parameters should be band-specific. For example, stricter confidence bounds may be appropriate in

higher cognitive bands, where premature progression has greater pedagogical cost, whereas lower bands may require more responsiveness to avoid unnecessary repetition. In this sense, parameter sensitivity is not merely a limitation but a necessary design feature of an uncertainty-aware formative assessment architecture.

## VI. LIMITATIONS

The most immediate limitation of this study is that HBC-Glicko remains entirely conceptual. No simulation studies, pilot deployments, or empirical datasets are presented to support the theoretical propositions outlined in Section IV. As a result, the practical advantages of the architecture cannot yet be confirmed beyond theoretical reasoning, and empirical testing remains a necessary next step before any performance claims can be made.

A second concern relates to the quality of Bloom-level classification. The architecture assumes that items are reliably tagged with the correct band and category labels. However, prior research has shown that cognitive-level classification is prone to error when inferred from surface features rather than the actual mental operations a task requires. Misclassified items could compromise the integrity of the band structure and produce progression decisions that are formally correct but educationally misleading.

A further limitation concerns parameter calibration. The architecture introduces several design parameters, including transition probabilities, confidence-bound multipliers, inactivity-growth rates, and minimum evidence requirements per category. Although these parameters make progression decisions explicit and adjustable, their optimal values cannot be established theoretically alone. Without simulation-based or empirical calibration, inappropriate parameter settings may make the system either overly conservative, delaying justified advancement, or overly permissive, allowing premature progression. Future work should therefore examine parameter sensitivity systematically and identify robust operating ranges across Bloom bands, domains, and learner populations.

Finally, the current architecture covers only three cognitive bands and does not incorporate the Create level or the knowledge-type dimension of Bloom's revised taxonomy. This restricts its applicability in learning contexts where generative or metacognitive performance is central. Extending the architecture to these dimensions represents a natural direction for future theoretical development.

## VII. DIRECTIONS FOR FUTURE RESEARCH

Several research directions follow directly from the architecture proposed in this study. The most immediate priority is simulation-based validation. Agent-based simulations in which synthetic learners with known band-specific ability profiles interact with HBC-Glicko would make it possible to evaluate promotion accuracy, sensitivity to item mislabeling, convergence of band-specific ratings, and the effects of alternative parameter settings under controlled conditions. Such simulations would provide a necessary first test of the decision rules before deployment in live educational settings.

A second priority is empirical validation in authentic formative assessment contexts. Deploying HBC-Glicko alongside a conventional adaptive system would allow comparison of learning outcomes, progression patterns, and instructor judgments of diagnostic usefulness. It would also generate the real-world response data needed to calibrate anchor difficulties and transition probabilities.

A third direction concerns scoring validity in higher bands. Because the architecture relies on rubric-based evidence for more complex cognitive performance, and because such scoring may be partially supported by LLMs, future work should examine agreement between human and LLM-based ratings across Bloom levels and assess how any systematic divergence affects rating stability and progression decisions. This question is especially important given the mixed evidence design introduced in Section II-D.

Finally, the architecture could be extended in two theoretically important directions: inclusion of the Create level and incorporation of knowledge-type distinctions within bands. Both extensions could improve representational fidelity, but they would also increase model complexity and calibration demands.

## VIII. CONCLUSION

This study introduced HBC-Glicko, a measurement architecture for formative adaptive assessment that embeds Bloom's revised taxonomy as a non-compensatory structural constraint within a Glicko-based learner model. The architecture formalizes learner state as a band-specific vector with explicit uncertainty, regulates progression through anchor-based readiness thresholds and confidence-bound decision rules, and models evidential decay under inactivity.

Three objectives were addressed. The study formalized cognitive hierarchy as a control principle governing adaptive progression rather than a descriptive classification of tasks. It derived the principal theoretical properties that follow from this design by construction. And it situates HBC-Glicko relative to existing rating-based approaches, clarifying that its contribution lies in combining multidimensional representation, uncertainty quantification, and non-compensatory progression within a single cognitively ordered architecture.

The contribution remains intentionally conceptual. Its practical value will depend on the reliability of item classification, the stability of anchor calibration, and the appropriateness of parameter configuration — conditions that can only be assessed through simulation and empirical deployment. These constitute the immediate next steps for this research program.

## CONFLICT OF INTEREST

The authors declare that there are no conflicts of interest related to this research.

#### AUTHORS' CONTRIBUTIONS

W.E.F. conceptualized the study, developed the HBC-Glicko architecture, formalized the theoretical model, and wrote the manuscript.

M.E. supervised the research, contributed to the conceptual refinement of the work, and provided critical revisions.

All authors reviewed and approved the final version of the manuscript.

#### FUNDING

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

#### ACKNOWLEDGMENT

The authors extend their heartfelt gratitude to the National Center for Scientific and Technical Research (CNRST) for crucial support through the "PhD-Associate Scholarship – PASS" program. This support played a pivotal role in advancing the research and ensuring the successful completion of this study.

#### REFERENCES

- [1] S. Abdi, H. Khosravi, and S. Sadiq, "Modelling learners in adaptive educational systems: A multivariate Glicko-based approach," in Proc. 11th Int. Learning Analytics and Knowledge Conf. (LAK '21), 2021, pp. 496–502. doi: 10.1145/3448139.3448189.
- [2] S. Abdi et al., "A multivariate Elo-based learner model for adaptive educational systems," in Proc. 12th Int. Conf. Educational Data Mining (EDM 2019), 2019, pp. 228–233.
- [3] L. W. Anderson and D. R. Krathwohl, Eds., *A Taxonomy for Learning, Teaching, and Assessing: A Revision of Bloom's Taxonomy of Educational Objectives*. New York, NY, USA: Longman, 2001.
- [4] A. Baddeley, "Working memory: Theories, models, and controversies," *Annu. Rev. Psychol.*, vol. 63, pp. 1–29, 2012. doi: 10.1146/annurev-psych-120710-100422.
- [5] K. Banujan et al., "Revolutionising educational assessment: Automated question classification using Bloom's taxonomy and deep learning techniques," *Int. J. Educ. Dev. Using Inf. Commun. Technol.*, vol. 19, no. 3, pp. 259–278, 2023.
- [6] F. Böck et al., "Learner models: Design, components, structure, and modelling," *User Model. User-Adapted Interact.*, 2025. doi: 10.1007/s11257-025-09434-4.
- [7] M. Bolsinova, B. Gergely, and M. J. S. Brinkhuis, "Keeping Elo alive: Evaluating and improving measurement properties of learning systems based on Elo ratings," *Brit. J. Math. Stat. Psychol.*, 2025. doi: 10.1111/bmsp.12395.
- [8] S. M. Brookhart, "Appropriate criteria: Key to effective rubrics," *Frontiers Educ.*, vol. 3, Art. 22, 2018. doi: 10.3389/educ.2018.00022.
- [9] W. Capon et al., "Validating an adaptive digital assessment of youth mental health needs: A cross-sectional study," *npj Digit. Med.*, 2026. doi: 10.1038/s41746-026-02374-2.
- [10] H. Cen, K. R. Koedinger, and B. Junker, "Learning factors analysis: A general method for cognitive model evaluation and improvement," in *Intelligent Tutoring Systems (Lecture Notes Comput. Sci., vol. 4053)*, M. Ikeda, K. D. Ashley, and T. W. Chan, Eds. Berlin, Germany: Springer, 2006, pp. 164–175. doi: 10.1007/11774303\_17.
- [11] L.-H. Chang and F. Ginter, "Automatic short answer grading for Finnish with ChatGPT," *Proc. AAAI Conf. Artif. Intell.*, vol. 38, no. 21, pp. 23173–23181, 2024. doi: 10.1609/aaai.v38i21.30363.
- [12] A. T. Corbett and J. R. Anderson, "Knowledge tracing: Modeling the acquisition of procedural knowledge," *User Model. User-Adapted Interact.*, vol. 4, no. 4, pp. 253–278, 1994. doi: 10.1007/BF01099821.
- [13] N. Cowan, "Working memory underpins cognitive development, learning, and education," *Educ. Psychol. Rev.*, vol. 26, no. 2, pp. 197–223, 2014. doi: 10.1007/s10648-013-9246-y.
- [14] A. Divya, V. Haridas, and J. Narayanan, "Automation of short answer grading techniques: Comparative study using deep learning techniques," in *Proc. 5th Int. Conf. Electrical, Computer and Communication Technologies (ICECCT)*, 2023, pp. 1–7.
- [15] H. Ebbinghaus, *Memory: A Contribution to Experimental Psychology*, H. A. Ruger and C. E. Bussenius, Trans. New York, NY, USA: Columbia Univ. Press, 1913 (Original work published 1885).
- [16] O. Fagbohun et al., "Beyond traditional assessment: Exploring the impact of large language models on grading practices," *J. Artif. Intell. Mach. Learn. Data Sci.*, vol. 2, no. 1, pp. 1–8, 2024. doi: 10.51219/JAIMLD/oluwole-fagbohun/19.
- [17] M. E. Glickman, "The Glicko system," *Tech. Rep.*, Harvard Univ., Cambridge, MA, USA, 1995.
- [18] M. E. Glickman, "Parameter estimation in large dynamic paired comparison experiments," *J. Roy. Stat. Soc. C*, vol. 48, no. 3, pp. 377–394, 1999. doi: 10.1111/1467-9876.00159.
- [19] M. E. Glickman, "Dynamic paired comparison models with stochastic variances," *J. Appl. Stat.*, vol. 28, no. 6, pp. 673–689, 2001. doi: 10.1080/02664760120059219.
- [20] A. Hamzah and S. Sosnovsky, "Modelling student knowledge in blended learning," in *Adj. Proc. 31st ACM Conf. User Modeling, Adaptation and Personalization (UMAP '23)*, 2023, pp. 76–80. doi: 10.1145/3563359.3597412.
- [21] A. D. Hofman et al., "Tracking with (un)certainly," *J. Intell.*, vol. 8, no. 1, Art. 10, 2020. doi: 10.3390/jintelligence8010010.
- [22] A. Ihich et al., "A systematic review on assessment in adaptive learning: Theories, algorithms and techniques," *Int. J. Adv. Comput. Sci. Appl.*, vol. 15, no. 7, 2024. doi: 10.14569/IJACSA.2024.0150785.
- [23] A. Jönsson and G. Svingby, "The use of scoring rubrics: Reliability, validity and educational consequences," *Educ. Res. Rev.*, vol. 2, no. 2, pp. 130–144, 2007. doi: 10.1016/j.edurev.2007.05.002.
- [24] E. N. Kandemir et al., "Adaptation of the multi-concept multivariate Elo rating system to medical students' training data," in *Proc. 14th Learning Analytics and Knowledge Conf. (LAK '24)*, 2024. doi: 10.1145/3636555.3636858.
- [25] J. D. Karpicke and H. L. Roediger, "Expanding retrieval practice promotes short-term retention, but equally spaced retrieval enhances long-term retention," *J. Exp. Psychol. Learn. Mem. Cogn.*, vol. 33, no. 4, pp. 704–719, 2007. doi: 10.1037/0278-7393.33.4.704.
- [26] M. J. Kolen and R. L. Brennan, *Test Equating, Scaling, and Linking: Methods and Practices*, 3rd ed. New York, NY, USA: Springer, 2014. doi: 10.1007/978-1-4939-0317-7.
- [27] N. Kornell, V. C. Rabelo, and A. T. Klingebiel, "Forgetting curves: New data and a new model," *J. Mem. Lang.*, vol. 125, Art. 104303, 2022. doi: 10.1016/j.jml.2022.104303.
- [28] D. R. Krathwohl, "A revision of Bloom's taxonomy: An overview," *Theory Into Practice*, vol. 41, no. 4, pp. 212–218, 2002. doi: 10.1207/s15430421tip4104\_2.
- [29] R. P. Kropp, H. W. Stoker, and W. L. Bashaw, "The validation of the taxonomy of educational objectives," *J. Exp. Educ.*, vol. 34, no. 3, pp. 69–76, 1966. doi: 10.1080/00220973.1966.11010940.
- [30] T. M. Larsen et al., "Probing internal assumptions of the revised Bloom's taxonomy," *CBE—Life Sci. Educ.*, vol. 21, no. 4, Art. ar66, 2022. doi: 10.1187/cbe.20-08-0170.
- [31] J. Li et al., "A new multidimensional computerized testing approach: On-the-fly assembled multistage adaptive testing based on multidimensional item response theory," *Mathematics*, vol. 13, no. 4, Art. 594, 2025. doi: 10.3390/math13040594.
- [32] M. Liu et al., "Investigating the effect of an adaptive learning intervention on students' learning," *Educ. Technol. Res. Dev.*, vol. 65, pp. 1605–1625, 2017.

[33] Q. Liu et al., "Multiple-choice questions (MCQs) for higher-order cognition: Perspectives of university teachers," *Innov. Educ. Teach. Int.*, vol. 61, no. 4, pp. 802–814, 2024. doi: 10.1080/14703297.2023.2222715.

[34] R. Lourdasamy, P. Magendrian, and C. M. Fonceca, "Analysis of cognitive levels of questions with Bloom's taxonomy: A case study," *Int. J. Softw. Innov.*, vol. 10, no. 1, pp. 1–22, 2022. doi: 10.4018/IJSI.297922.

[35] T. Muhayimana, L. Kwizera, and M. R. Nyirahabimana, "Using Bloom's taxonomy to evaluate the cognitive levels of Primary Leaving English Exam questions in Rwandan schools," *Curriculum Perspectives*, vol. 42, no. 1, pp. 51–63, 2022. doi: 10.1007/s41297-021-00156-2.

[36] M. Pankiewicz and M. Bator, "Elo rating algorithm for the purpose of measuring task difficulty in online learning environments," *e-mentor*, vol. 5, no. 82, pp. 43–51, 2019. doi: 10.15219/em82.1444.

[37] D. Parmigiani et al., "Formative assessment in higher education: An exploratory study within programs for professionals in education," *Frontiers Educ.*, vol. 9, Art. 1366215, 2024. doi: 10.3389/educ.2024.1366215.

[38] P. I. Pavlik Jr., H. Cen, and K. R. Koedinger, "Performance factors analysis: A new alternative to knowledge tracing," in *Artificial Intelligence in Education*, V. Dimitrova, R. Mizoguchi, B. du Boulay, and A. Graesser, Eds. Amsterdam, Netherlands: IOS Press, 2009, pp. 531–538.

[39] R. Pelánek, "Applications of the Elo rating system in adaptive educational systems," *Comput. Educ.*, vol. 98, pp. 169–179, 2016. doi: 10.1016/j.compedu.2016.03.017.

[40] R. Pelánek et al., "Elo-based learner modeling for the adaptive practice of facts," *User Model. User-Adapted Interact.*, vol. 27, no. 1, pp. 89–118, 2017. doi: 10.1007/s11257-016-9185-7.

[41] R. Pelánek, S. Silovská, and J. Řihák, "Complexity and difficulty of items in learning systems," *Int. J. Artif. Intell. Educ.*, vol. 32, no. 4, pp. 1003–1048, 2022. doi: 10.1007/s40593-021-00252-4.

[42] H. Ravand et al., "Measuring cognitive levels in high-stakes testing: A CDM analysis of a university entrance examination using Bloom's taxonomy," *Frontiers Educ.*, vol. 10, Art. 1644811, 2025. doi: 10.3389/educ.2025.1644811.

[43] H. L. Roediger and J. D. Karpicke, "Test-enhanced learning: Taking memory tests improves long-term retention," *Psychol. Sci.*, vol. 17, no. 3, pp. 249–255, 2006. doi: 10.1111/j.1467-9280.2006.01693.x.

[44] R. Sacripante et al., "Forgetting rates of gist and peripheral episodic details in prose recall," *Memory*, vol. 31, no. 3, pp. 328–343, 2023. doi: 10.1080/09658211.2022.2052521.

[45] M. Schneider and B. A. Simonsmeier, "How does prior knowledge affect learning? A review of 16 mechanisms and a framework for future research," *Learn. Individ. Differ.*, vol. 122, Art. 102744, 2025. doi: 10.1016/j.lindif.2025.102744.

[46] F. Sense et al., "An individual's rate of forgetting is stable over time but differs across materials," *Topics Cogn. Sci.*, vol. 8, no. 1, pp. 305–321, 2016. doi: 10.1111/tops.12183.

[47] R. Sreedhar et al., "Comparing scoring consistency of large language models with faculty for formative assessments in medical education," *J. Gen. Intern. Med.*, vol. 40, no. 1, pp. 127–134, 2025. doi: 10.1007/s11606-024-09050-9.

[48] K. F. Stanger-Hall, "Multiple-choice exams: An obstacle for higher-level thinking in introductory science classes," *CBE—Life Sci. Educ.*, vol. 11, no. 3, pp. 294–306, 2012. doi: 10.1187/cbe.11-11-0100.

[49] J. Sweller, "The development of cognitive load theory: Replication crises and incorporation of other theories can lead to theory expansion," *Educ. Psychol. Rev.*, vol. 35, Art. 87, 2023. doi: 10.1007/s10648-023-09817-2.

[50] J. Sweller, P. Ayres, and S. Kalyuga, *Cognitive Load Theory*. New York, NY, USA: Springer, 2011. doi: 10.1007/978-1-4419-8126-4.

[51] L. Y. Tan et al., "Artificial intelligence-enabled adaptive learning platforms: A review," *Comput. Educ. Artif. Intell.*, vol. 9, Art. 100429, 2025. doi: 10.1016/j.caeai.2025.100429.

[52] M. Tarrant and J. Ware, "Impact of item-writing flaws in multiple-choice questions on student achievement in high-stakes nursing assessments," *Med. Educ.*, vol. 42, no. 2, pp. 198–206, 2008. doi: 10.1111/j.1365-2923.2007.02957.x.

[53] G. Tsaparlis, "Higher and lower-order thinking skills: The case of chemistry revisited," *J. Baltic Sci. Educ.*, vol. 19, no. 3, pp. 467–483, 2020. doi: 10.33225/jbse/20.19.467.

[54] B. Taylor, F. Kisby, and A. Reedy, "Rubrics in higher education: An exploration of undergraduate students' understanding and perspectives," *Assess. Eval. Higher Educ.*, vol. 49, no. 6, pp. 978–992, 2024. doi: 10.1080/02602938.2023.2299330.

[55] W. J. van der Linden and C. A. W. Glas, Eds., *Elements of Adaptive Testing*. New York, NY, USA: Springer, 2010. doi: 10.1007/978-0-387-85461-8.

[56] W. J. van der Linden and R. K. Hambleton, Eds., *Handbook of Modern Item Response Theory*. New York, NY, USA: Springer, 1997.

[57] H. Vermeiren et al., "Balancing stability and flexibility: Investigating a dynamic K value approach for the Elo rating system in adaptive learning environments," *User Model. User-Adapted Interact.*, vol. 36, Art. 4, 2026. doi: 10.1007/s11257-025-09439-z.

[58] H. Vermeiren et al., "Psychometrics of an Elo-based large-scale online learning system," *Comput. Educ. Artif. Intell.*, vol. 8, Art. 100376, 2025. doi: 10.1016/j.caeai.2025.100376.

[59] E. Virranmäki, K. Valta-Hulkkonen, and A. Pellikka, "Geography tests in the Finnish Matriculation Examination in Paper and digital forms: An analysis of questions based on revised Bloom's taxonomy," *Stud. Educ. Eval.*, vol. 66, Art. 100896, 2020. doi: 10.1016/j.stueduc.2020.100896.

[60] C. Voigt, L. Blömer, and U. Hoppe, "Analysing the cumulative hierarchy of the taxonomy of learning objectives in flipped classroom," *Int. J. Higher Educ. Pedagogies*, vol. 1, no. 1, pp. 1–12, 2020. doi: 10.33422/ijhep.v1i1.13.

[61] H. Xie et al., "Trends and development in technology-enhanced adaptive/personalized learning: A systematic review of journal publications from 2007 to 2017," *Comput. Educ.*, vol. 140, Art. 103599, 2019. doi: 10.1016/j.compedu.2019.103599.

[62] A. C. M. Yang, B. Flanagan, and H. Ogata, "Adaptive formative assessment system based on computerized adaptive testing and the learning memory cycle for personalized learning," *Comput. Educ. Artif. Intell.*, vol. 3, Art. 100104, 2022. doi: 10.1016/j.caeai.2022.100104.

[63] M. Yudelson et al., "Leveraging skill hierarchy for multi-level modeling with Elo rating system," in *Proc. 6th ACM Conf. Learning@ Scale (L@S '19)*, 2019, pp. 1–4. doi: 10.1145/3330430.3333645.

[64] N. L. B. Zaidi et al., "Pushing critical thinking skills with multiple-choice questions: Does Bloom's taxonomy work?" *Acad. Med.*, vol. 93, no. 6, pp. 856–859, 2018. doi: 10.1097/ACM.0000000000002087.

[65] Y. Zhang et al., "Development of a computerized adaptive testing for internet addiction," *Frontiers Psychol.*, vol. 10, Art. 1010, 2019. doi: 10.3389/fpsyg.2019.01010.

APPENDIX A. SUMMARY OF NOTATION

TABLE III. SUMMARIZES THE PRINCIPAL NOTATION USED THROUGHOUT THE FORMAL SPECIFICATION OF THE HBC-GLICKO ARCHITECTURE. SYMBOLS ARE GROUPED BY FUNCTIONAL CATEGORY. SUBSCRIPT CONVENTIONS FOLLOW STANDARD PSYCHOMETRIC AND RATING-SYSTEM USAGE

Category	Symbol	Definition
<i>General indices</i>	$n$	Learner index
	$t$	Rating period (session) index
	$b$	Cognitive band index, $b \in \{1, 2, 3\}$

	$\mathcal{B}$	Set of cognitive bands {Low, Mid, High}
	$\mathcal{K}$	Set of Bloom categories within band b
<b>Learner state</b>	$r_{n,b,t}$	Rating of learner n in band b at period t
	$RD_{n,b,t}$	Rating deviation (uncertainty) of learner n in band b at period t
	$\mathcal{S}_{n,t}$	Full learner measurement state at period t
<b>Item parameters</b>	$d_i$	Calibrated difficulty of item i on the unified scale
	$RD_{i,t}$	Rating deviation of item i at period t
<b>Response model</b>	$q$	Scaling constant, $q = \ln(10) / 400$
	$g(RD)$	Uncertainty weighting function
	$E_{n,i,t}$	Expected score of learner n on item i at period t
	$S_{n,i,t}$	Observed score, $S \in [0, 1]$
	$\delta_{n,i,t}$	Prediction error, $S - E$
	$\omega_{n,b,t}$	Glicko-style adaptive update weight
<b>Within-band routing</b>	$\pi_b(k)$	Target sampling distribution over categories in band b
	$c_t(k)$	Delivery counter for category k at step t
	$\Delta_t(k)$	Coverage deficit for category k at step t
	$p^*$	Target performance level for adaptive item selection
	$H_s$	Set of items already scored in session s
<b>Progression</b>	$A_{b+1}$	Entry anchor set for band b + 1
	$d_{b+1}^{ref}$	Mean reference difficulty of entry anchor items for band b + 1
	$p_{b \rightarrow b+1}$	Target transition probability from band b to band b + 1
	$T_{b \rightarrow b+1}$	Progression threshold from band b to band b + 1
<b>Decision rules</b>	$z_b$	Confidence multiplier for band b
	$LCB_{n,b}$	Lower confidence bound, $r - z \cdot RD$
	$UCB_{n,b}$	Upper confidence bound, $r + z \cdot RD$
	$N_{n,b}(k)$	Scored administrations for learner n in category k of band b
	$m_b$	Minimum evidence per category required for promotion
<b>Inactivity</b>	$t_{n,b}$	Sessions since last scored attempt by learner n in band b
	$c_b$	Inactivity-driven uncertainty growth rate for band b
	$RD_{max}$	Maximum admissible rating deviation
	$RD_{n,b}^{pre^*}$	Effective pre-update uncertainty after inactivity adjustment