

JEAI-25-18

BEWA: A Bayesian Epistemology-Weighted Artificial Intelligence Framework for Scientific Inference

Craig S Wright*

Department of Computer Science, University of Exeter Ltd, Exeter, UK

*Corresponding author: Craig S Wright, Department of Computer Science, University of Exeter Ltd, Exeter, UK, E-mail: craig@rcjbr.org

Received date: June 19, 2025; Accepted date: June 27, 2025; Published date: July 19, 2025

Citation: Wright CS (2025) BEWA: A Bayesian Epistemology-Weighted Artificial Intelligence Framework for Scientific Inference. J Eng Artif Intell Vol.1 No.2: 18.

Abstract

The proliferation of scientific literature and the accelerating complexity of epistemic discourse have outpaced the evaluative capacities of both human scholars and conventional artificial intelligence systems. In response, we propose Bayesian Epistemology with Weighted Authority (BEWA), a computational architecture for truth-oriented knowledge modelling. BEWA formalises belief as a probabilistic relation over structured claims, indexed to authors, contexts, and replication history, and updated via evidence-driven Bayesian mechanisms. Integrating canonical authorial identification, dynamic belief networks, replication-weighted citation metrics, and epistemic decay protocols, the system constructs an evolving belief state that prioritises truth utility while resisting social and citation-based distortions. By anchoring every propositional unit in structured metadata and linking updates to semantic replication and contradiction analysis, BEWA enables automated, principled reasoning across a corpus of scientific knowledge. This work advances the theoretical foundations and practical frameworks necessary for autonomous epistemic agents to assess, revise, and propagate beliefs in dynamic scientific environments.

Keywords: Bayesian epistemology; Belief update; Autonomous reasoning; Replication weighting; Scientific AI; Structured knowledge; Truth utility; Author credibility modelling; Epistemic integrity; Probabilistic knowledge representation

Introduction

The crisis of epistemic overload in modern scientific inquiry has exposed a critical deficiency in how truth claims are assessed, validated, and integrated across time and domain. The exponential growth in peer-reviewed publications, accompanied by inconsistent replication rates, entrenched citation biases, and the sociological entanglements of scientific authorship, has rendered traditional mechanisms of epistemic filtering increasingly obsolete. Simultaneously, artificial intelligence while having demonstrated capacity in data correlation and language generation remains fundamentally ill-equipped to perform rigorous epistemic reasoning. This gap is not merely technical but conceptual: current AI systems lack any principled framework for evaluating the truth-promoting value of claims, discerning authoritative sources, or understanding belief as a structured probabilistic relation between agents, claims, and contexts.

The present work introduces a formal architecture Bayesian Epistemology with Weighted Authority (BEWA)

which systematically encodes the logic of belief formation, update, and decay, guided by the core axioms of Bayesian rationality, tempered by structural mechanisms for authority weighting, replication scoring, and temporal reassessment.

BEWA does not seek to supplant human reasoning, but to embed within computational systems the formal machinery required to model, test, and evaluate propositional knowledge in scientific domains. Where contemporary models optimise for coherence or linguistic plausibility, BEWA operationalises scientific epistemology as a computable framework: every claim must be anchored, weighted, and situated within an evolving belief graph responsive to empirical reinforcement, contradiction, and decay.

At its core, this work constitutes a reimagining of what it means for AI to “know”. Rather than train a system to generate plausible continuations of language, BEWA instantiates an agent that reasons over propositions, scrutinises claims for

cross-referenced coherence, and reassigns belief weightings in response to changing empirical and social contexts. In doing so, it lays the foundation for autonomous epistemic agents not only capable of ingesting scientific knowledge, but of judging its weight, revising their stance, and constructing rational belief networks that remain resilient, transparent, and optimally truth-directed. This architecture forms not merely a tool for scientific indexing, but the cornerstone of a broader philosophical project: to make scientific reasoning computationally tractable without abandoning the rigour, fallibilism, and probabilistic humility that science demands.

Motivation and context

The production of scientific literature has expanded exponentially in the last two decades, with repositories such as PubMed, arXiv, and Scopus indexing tens of millions of publications across diverse domains. While this proliferation reflects a thriving global research enterprise, it has introduced a fundamental tension: the epistemic cost of abundance. Human cognitive capacity to read, assimilate, verify, and cross-correlate claims has not scaled proportionally, leaving both researchers and knowledge systems vulnerable to the pitfalls of unreplicated findings, citation cascades, and reputational distortions. As Ioannidis et al., and subsequent meta-analyses have shown, a significant fraction of published results even in high-impact journals fail to replicate, and yet are continuously cited, shaping the direction of scientific discourse in potentially misleading ways [1].

Conventional artificial intelligence systems, while effective at information retrieval and semantic clustering, remain epistemically shallow. Their models lack the formal apparatus to discriminate between verified knowledge, unsupported assertions, or strategic citations, and thus perpetuate the same problems of over-generalisation and under-qualification found in human discourse. The incorporation of Large Language Models (LLMs) into scientific workflows exacerbates this vulnerability, given their high fluency but epistemic indifference. What is needed is not a language generator, but a system for knowledge discrimination: a machine that believes rationally.

BEWA is proposed as a principled framework to address this lacuna. Drawing on foundational work in formal epistemology and probabilistic reasoning, BEWA offers a system that does not merely store scientific claims but assigns them structured belief values conditional on their provenance, corroboration, authorial record, and epistemic utility. By aligning computational epistemology with the standards of scientific methodology replication, citation credibility, authorial

accountability BEWA is positioned to serve as a formal counterpart to scientific reasoning in *silico*.

The problem of scientific epistemology in AI

The central problem of scientific epistemology in artificial intelligence is not the extraction of propositions, nor even their contextualisation, but the absence of principled belief management. Contemporary AI systems, particularly those driven by deep learning and language modelling architectures, have achieved notable success in identifying textual similarity, performing question answering, and generating plausible discursive outputs. However, these models lack any internalisation of epistemic constraints, such as the replicability of claims, the entrenchment of belief under sustained verification, or the significance of retraction and reputational decay. Without a formal epistemic scaffold, such systems merely reflect the statistical regularities of their training data, rather than critically adjudicating between levels of epistemic warrant.

The failure to incorporate epistemic stratification has deep consequences. Systems trained indiscriminately on corpora that blend rigorous research with fringe science, preprints with peer-reviewed material, or established consensus with contested speculation, conflate reliability with frequency. They lack the means to discount anomalous or outdated claims, to resolve conflicts between competing hypotheses, or to track the evolution of knowledge claims across time and replication. As such, they are incapable of supporting decision systems scientific, legal, medical, or policy-based that depend on epistemic stability and justified belief rather than surface-level coherence.

Moreover, epistemological naivete in AI raises a meta-scientific hazard: The uncritical automation of flawed scientific inference. The citation of unreplicated or retracted work, the omission of contrary findings, or the aggregation of selectively reported results has long been documented as a contributor to scientific dysfunction [1,2]. When AI systems ingest and reproduce these distortions without corrective filtering, they entrench misinformation under the guise of precision. A scientific AI must not merely infer, but adjudicate assigning belief in a manner consistent with rational constraints, probabilistic coherence, and methodological rigour. This is the foundational imperative that BEWA addresses.

Objectives of BEWA

The Bayesian Epistemology Weighting Architecture (BEWA) is developed to resolve the epistemic lacunae in contemporary AI systems by formalising a structured, axiomatic approach to scientific belief management. Its core objective is the construction of a truth-promoting knowledge architecture

in which claims are not simply stored or retrieved, but dynamically assessed, weighted, and propagated in accordance with their epistemic warrant. To this end, BEWA integrates a full Bayesian inferential layer wherein prior beliefs are established from canonical sources, updated through reproducible evidence, and decayed with temporal distance or contradiction. The resulting framework is not static but temporally responsive, adjusting belief states in proportion to the evolving consensus and integrity of underlying claims.

BEWA's second principal objective is to embed authorial credibility and citation structure as explicit variables in belief formation. Unlike models that treat all sources equivalently, BEWA distinguishes canonical authors from peripheral contributors through reputation modelling, peer review engagement, replication track record, and retraction history. Each claim inherits a weighted score derived not only from its own properties, but from the broader epistemic ecosystem of author, venue, and reception. In this way, BEWA promotes methodological accountability: Authors whose work has endured scrutiny and replication enhance the evidential force of their claims; those whose work is anomalous, inconsistent, or retracted attenuate the propagation of falsehoods.

Thirdly, BEWA is designed to operationalise these epistemic mechanisms across complex, multi-claim knowledge graphs. Through a formal propositional structure, semantic linkage, and probabilistic belief propagation, BEWA enables the dynamic evolution of belief networks wherein conflict, contradiction, and refinement are explicitly modelled. It does not seek finality, but stability under continuous critical evaluation. The architecture aims not merely to aggregate knowledge, but to support scientific reasoning through structured epistemic decay, replication-triggered resets, and contradiction-responsive attenuation. BEWA thereby elevates AI from passive aggregator to active epistemic agent, consistent with the axioms of Bayesian rationality and the imperatives of contemporary philosophy of science.

Contributions and novelty

This work introduces the Bayesian Epistemology Weighting Architecture (BEWA), a mathematically rigorous, axiomatically grounded system for autonomous epistemic processing in scientific domains. Unlike prior models that conflate data aggregation with epistemic judgment, BEWA offers a formal mechanism for belief formation, revision, and decay rooted in Bayesian inference and justified by philosophical and computational standards of rationality. Its novelty lies in the explicit coupling of structured propositional claims with a belief updating framework that is temporally sensitive, source-critical,

and capable of representing both support and contradiction across a dynamically evolving knowledge space.

BEWA's primary contribution is the integration of epistemic virtues replicability, citation strength, author credibility, and peer review participation into a formally specified weighting mechanism. These properties are not ad hoc features but are embedded as first-class terms within the probabilistic architecture. The system models belief not as a monolithic scalar but as a multidimensional function of authorial history, domain coherence, and cross-claim dependency. Moreover, it introduces the notion of epistemic decay and rejuvenation protocols: Aged or isolated claims lose credence over time unless renewed through replication or citation; conversely, successful replications trigger weight propagation throughout associated belief clusters.

A further contribution is the development of a semantic infrastructure that enables BEWA to interpret, map, and reconcile overlapping claims across multiple disciplines and terminological systems. By constructing belief networks with semantic equivalence mappings and contradiction matrices, BEWA supports cross-disciplinary synthesis without erasing contextual nuance. Finally, the architecture's modularity, cryptographic provenance tracking, and audit interfaces ensure reproducibility and transparency, making it suitable not only for scientific archiving but for regulatory, forensic, and policy-critical applications. In aggregating formal epistemology, machine reasoning, and information security, BEWA constitutes a foundational advancement in the design of truth-sensitive artificial intelligence.

Overview of System Architecture

This section provides a structural and conceptual overview of the Bayesian Epistemology-Weighted AI (BEWA) system. It delineates the foundational principles that govern the architectural framework, mapping the philosophical underpinnings of the system into concrete computational form. The system is not merely an exercise in engineering but a principled implementation of formal epistemology, designed to handle dynamic and conflicting scientific information within a rational, belief-updating paradigm. Accordingly, the architecture is built to ensure that knowledge claims are not merely stored or retrieved, but are actively evaluated, weighted, and revised through rigorously formalised evidentiary structures. BEWA is therefore conceived as a continuously self-correcting epistemic engine, capable of assessing propositional validity with statistical robustness, while maintaining scepticism until replicability thresholds are met.

The structure of the system rests on three critical pillars. First, the philosophical grounding in Bayesian epistemology ensures that all claims are embedded within a probabilistic inferential model, allowing beliefs to evolve incrementally and cautiously. Second, the architectural design is modular and hierarchical, integrating ingest pipelines, structured claim representation, belief propagation, and dynamic decay mechanisms, while ensuring coherence across each component. Third, the system embodies a core commitment to epistemic integrity: It prioritises replicability, resists premature belief inflation, and penalises epistemically hollow citations or uncorroborated popularity. This section thus introduces the high-level logic of BEWA's operation, setting the stage for the detailed mechanics explored in subsequent sections.

Philosophical basis: Bayesian epistemology

Axiom 1 (rational belief as probability): Any rational agent's degree of belief in a proposition H must be representable by a real-valued probability $P(H)$ within the closed interval $(0,1)$, such that the agent's belief system adheres to the Kolmogorov axioms of probability.

This axiom forms the foundational commitment of BEWA's inferential framework: epistemic states are mapped to probability distributions, and rational updates to these states are governed by Bayes' Theorem. This view is justified on both normative and decision-theoretic grounds. As shown in Ramsey et al., and formalised by de Finetti et al., the use of probabilities to represent belief is both behaviourally and logically coherent when an agent aims to avoid Dutch Book incoherence [3,4]. That is, if one's belief assignments violate the probability axioms, a set of bets can be constructed that guarantees a loss, revealing the irrationality of those beliefs.

Definition 1 (Bayesian agent): An agent A is said to be Bayesian if for any proposition H and evidence E , its belief update adheres to:

$$P(H | E) = \frac{P(E | H) \cdot P(H)}{P(E)}, \text{ Provided } P(E) > 0.$$

Theorem 1 (Bayesian coherence criterion): If an agent's beliefs are updated via Bayes' Theorem, and their priors obey the probability axioms, then the agent is immune to Dutch Book constructions.

Proof. See van Fraassen et al., and Joyce et al., where the coherence of Bayesian updating is proved using decision-theoretic formalism [5,6]. The key idea is that any deviation from Bayes' Rule allows a clever adversary to exploit inconsistencies via bets that yield guaranteed losses.

BEWA adopts the Bayesian framework not as a mere computational convenience, but as a foundational epistemological commitment. Within the philosophy of science, Bayesianism offers a natural model for understanding confirmation, falsification, and the accumulation of scientific knowledge over time. Classical falsificationism fails to account for degrees of belief and the nuanced role of partial evidence [7]. In contrast, Bayesian epistemology allows for iterative refinement of confidence levels, maintaining probabilistic caution in light of ambiguous or conflicting data a property critical for any AI epistemic engine operating over scientific domains.

Proposition 1 (gradual confirmation): Let $\{E_1, E_2, \dots, E_n\}$ be an increasing sequence of independent pieces of evidence favouring H . Then $P(H | E_1, \dots, E_n)$ converges to 1 as $n \rightarrow \infty$ if $P(E_i | H) > P(E_i | \neg H)$ for all i .

Proof. Follows from iterative application of Bayes' Theorem and the law of large numbers; see Howson et al., for formal derivation under the assumption of conditional independence [8].

From an architectural perspective, BEWA instantiates this philosophical position by encoding every scientific proposition as a probabilistically weighted claim, updated through conditionalisation as new evidence is introduced. This results in a dynamically evolving epistemic state space where each belief's trajectory reflects its empirical support history. By constructing belief as an evolving posterior under a well-founded prior, the system resists epistemic stasis and prevents premature convergence on erroneous claims an outcome observed in systems lacking formal uncertainty management [9].

Thus, Bayesian epistemology not only provides a rational framework for belief representation and update, but also satisfies the core design requirements of BEWA: dynamism, self-correction, resistance to noise, and formal auditability.

System-level design principles

The design of BEWA as an epistemic AI system is governed by a set of logically necessary architectural invariants, derived from the formal requirements of probabilistic reasoning, knowledge provenance, and inferential consistency. Each design principle serves a dual function: it guarantees operational coherence within the system, and it enforces alignment with normative principles of rational belief updating and scientific justification.

Principle 1 (compositional modularity): Every component of the system must admit independent verification and recomposability without epistemic leakage. This principle follows from foundational modularity theorems in software

verification and distributed system logic and ensures that any subsystem e.g., belief update, citation graph traversal, claim parsing can be tested and audited in isolation [10]. BEWA employs this to enable fault isolation and logical traceability.

Principle 2 (evidential locality): All belief updates must be a function only of local evidence and semantically adjacent claims. This constraint avoids logical omniscience problems, preserving decidability and computational tractability while ensuring that belief revision does not occur through non-causal or disconnected assertions [11]. The networked structure of BEWA propagates epistemic changes via bounded dependency paths.

Principle 3 (non-monotonic reversibility): No belief in the system is irrevocable; each posterior must remain subject to revision upon presentation of new evidence. This is a direct consequence of probabilistic logic and a formal rejection of monotonic reasoning frameworks that dominate traditional symbolic AI [12]. In BEWA, the posterior distribution over claims is non-monotonic, enabling revision under contradiction or new replication results.

Principle 4 (temporal sensitivity): All epistemic weights must be functions of both evidential strength and temporal distance from the current system state. Following formal models of information decay and memory-limited inference, BEWA introduces time-dependent weighting functions to encode the epistemic perishability of unreplicated or outdated claims [13].

Principle 5 (proof-carrying claims): Every claim object must carry forward the formal trace of its derivation and belief trajectory. Inspired by proof-carrying code models in formal verification, this principle ensures that every change to the system's epistemic state is accompanied by a verifiable, reproducible, and human-readable justification chain [14]. This enables not only internal consistency but external auditability and trust.

The collective enforcement of these principles ensures that BEWA is not merely probabilistic, but epistemologically principled. It maintains rigorous boundaries on the scope of inference, prevents epistemic drift through ungrounded propagation, and provides the infrastructure necessary for transparent and accountable scientific reasoning.

Epistemic integrity and truth-promoting utility

BEWA enforces epistemic integrity as a formal constraint on system-level belief formation and propagation. Epistemic integrity, within this architecture, is defined as the adherence of belief trajectories to a coherence-preserving inferential structure that prioritises evidentiary robustness over frequency,

visibility, or institutional bias. The motivation for this principle is grounded in both the philosophy of science and information theory: Scientific knowledge production must resist epistemic drift, bandwagon effects, and citation cascades that artificially inflate the credibility of unverified claims [2]. Accordingly, BEWA operationalises a utility function over claims that reflects not popularity or downstream use, but the claim's contribution to the discovery, confirmation, or rectification of scientific truth.

Definition 2 (truth-promoting utility function): Let C denote the set of structured claims, and let $U: C \rightarrow \mathbb{R}$ be a function such that for each $c \in C$,

$U(c) = \lambda_1 R(c) + \lambda_2 D(c) + \lambda_3 V(c) - \lambda_4 B(c)$, where $R(c)$ is the replication score of c ; $D(c)$ is the epistemic distinctiveness or novelty; $V(c)$ is the verified downstream influence (e.g., in confirmed applications); $B(c)$ is the belief inflation penalty due to network echo effects; $\lambda_i \in \mathbb{R}_{\geq 0}$ are domain-tunable weights.

This function is constructed to prioritise claims that not only survive empirical testing, but also contribute epistemically non-redundant insight. The component $D(c)$ penalises claims that merely replicate known results without methodological refinement or contextual extension. The penalty term $B(c)$ reflects the phenomenon where claims propagate in citation networks without independent validation a structure first analysed in Chu et al., who demonstrated the role of preferential attachment in distorting perceived scientific consensus [15].

Axiom 2 (integrity-first propagation): No belief update may propagate through the system unless the associated claim passes a minimum threshold of truth utility as evaluated by $U(c)$.

This axiom restricts the automatic diffusion of belief across the epistemic graph, thereby minimising the risk of structural bias or error reinforcement. It reflects a departure from naive Bayesian belief networks by introducing a truth-oriented constraint beyond conditional probability updates. This is aligned with recent findings in epistemic network theory, which show that long-range propagation of low-fidelity signals leads to error cascades [16]. BEWA's truth utility acts as a circuit breaker, ensuring only epistemically responsible claims participate in long-range influence operations.

Proposition 2 (utility-constrained epistemic stability): For a fixed claim set C and bounded belief update rates, the imposition of a lower bound $\delta > 0$ on $U(c)$ for propagation ensures bounded volatility in belief trajectories over time.

Proof. See the utility-stabilised propagation model in Banerjee et al., adapted with epistemic cost functions [17]. The

result follows from limiting the propagation of low-certainty, high-volatility nodes that dominate in unconstrained models.

Thus, epistemic integrity in BEWA is not a vague normative aim but a computationally enforceable constraint, instantiated through a formally defined utility function, bounded propagation logic, and volatility-dampening mechanisms. The result is an inferential architecture that aligns scientific computation with the normative demands of justification and truth-tracking.

Data Ingestion and Canonical Normalisation

This section outlines the foundational mechanisms by which BEWA acquires, processes, and standardises scientific information from disparate sources into a coherent and ontologically stable framework. Unlike conventional AI pipelines that indiscriminately ingest unstructured data, BEWA enforces stringent epistemic gatekeeping at the point of entry. All input material ranging from peer-reviewed publications and replication studies to technical reports and retraction notices is passed through a multi-layered filtration system designed to identify provenance, authenticate origin, and normalise both linguistic and structural content into canonicalised, author-bound claims. The ingestion pipeline is not only syntactic; it is epistemologically motivated, ensuring that the eventual reasoning engine operates on stable, authorial propositions rather than transient or ambiguous textual artefacts.

Central to this process is the resolution of authorship and the stabilisation of claims into persistent identifiers that preserve the semantic integrity of original assertions across time and format. This canonicalisation is not merely a bibliographic convenience; it serves as the scaffolding upon which belief updating, citation tracing, and contradiction mapping are built. Claims are disambiguated and indexed according to domain context, temporality, and authorial authority, with metadata extracted, verified, and cryptographically anchored. By securing both semantic and epistemic consistency at the ingestion stage, the system guarantees that downstream analytical processes are not corrupted by noise, duplication, or misattribution. The subsections that follow detail the protocols governing authoritative source selection, the logic of canonical ID formation for authors and claims, and the metadata integrity regime that underpins the entire structure.

Authoritative source domains

The integrity of any epistemic reasoning system is inextricably linked to the reliability of its inputs. In BEWA, the designation of authoritative source domains serves as an

axiomatic filtration criterion: Only sources that satisfy a minimal condition of epistemic legitimacy are permitted to influence belief formation. This restriction is not arbitrary but follows from formal constraints on information provenance, epistemic justification, and noise minimisation in inferential systems. To ensure that downstream probabilistic belief updates are not corrupted by unreliable or spurious assertions, all ingested sources are required to satisfy criteria derived from formal literature on information-theoretic trust and epistemic reliability models [18,19].

Axiom 3 (source legitimacy constraint): A document d may be admitted into the system's evidentiary graph \mathcal{E} if and only if it belongs to a domain D_{auth} satisfying: $\forall d \in \mathcal{E}, \text{source}(d) \in D_{auth} \Leftrightarrow \text{Verifiable}(d) \wedge \text{Indexed}(d) \wedge \text{PeerReviewed}(d)$

Here, $\text{Verifiable}(d)$ implies the document is publicly accessible and persistent (e.g., DOI-registered); $\text{Indexed}(d)$ indicates that the source exists within trusted citation databases (e.g., PubMed, Scopus, Web of Science); and $\text{PeerReviewed}(d)$ signifies that the document has passed through a documented, non-anonymous review process with editorial oversight.

Definition 3 (source domain tuple): Each domain D is a tuple (R, T, C) , where: R is a reputational index (e.g., impact factor, field-weighted citation impact); T is the review transparency metric (e.g., registered reports, open review); C is the corpus compliance rate fraction of submissions adhering to replication, data availability, or statistical preregistration norms.

Proposition 3 (reliability bound): Let $P_{err}(D)$ denote the probability that a claim extracted from domain D is later retracted or falsified. Then: If $C > 0.8$ and $T > 0.6$, $P_{err}(D) < \epsilon$, for some $\epsilon \in \mathbb{R}^+$ bounded above by 0.05 in empirical studies [1,20].

This probabilistic bound is critical: It establishes that the inclusion of sources from domains satisfying high C and T values e.g. those enforcing mandatory data sharing or pre-registration is mathematically associated with lowered epistemic risk. Conversely, domains lacking these features are systematically excluded to minimise noise infiltration.

Implementation: BEWA operationalises D_{auth} as a whitelist defined over registry-linked sources (e.g., Crossref DOIs), supported by cryptographic signatures where available (e.g., ORCID-verified authorship, PubPeer-linked review commentary). Grey literature, blog posts, and unverifiable claims are explicitly excluded, and their attempted inclusion triggers a rejection trace logged to the audit ledger.

Hence, authoritative source domains in BEWA are not designated by subjective prestige or ad hoc authority, but through mathematically defensible, verifiably structured compliance with epistemic reliability constraints. This enables

the inferential engine to maintain input integrity and safeguards the downstream logic against polluted or non-reproducible evidence structures.

Canonical author and claim identification

To ensure referential stability and semantic precision within BEWA, each scientific assertion must be linked to a canonical representation of both its author and its propositional content. This section formalises the identification of claims and authors as foundational invariants in the system's epistemic model. The goal is to eliminate ambiguity, resolve homonymy and synonymy across publications, and generate a stable ontology of attributions from which downstream belief calculations can proceed.

Definition 4 (canonical author identifier): Let A be the set of all authorial agents. For any author $a \in A$, define a Canonical Author Identifier $CAID(a)$ such that: $CAID(a) = \text{Hash}(\text{ORCID}(a) \parallel \text{DisambiguatedName}(a) \parallel \text{Affiliation}(a))$, where Hash is a collision-resistant cryptographic function (e.g. SHA-256), and DisambiguatedName is derived *via* a resolution algorithm over publication metadata [21].

This ensures that each author, even across variant naming conventions, contributes to the belief network under a persistent identity. Where ORCID is unavailable, disambiguation defaults to supervised learning over co-authorship graphs and venue clustering [22].

Definition 5 (canonical claim signature): Let C denote the set of claims extracted from the corpus. For any claim $c \in C$, its canonical form is given by: $CCS(c) = \text{Hash}(\text{Normalise}(c_{\text{text}}) \parallel CAID(a) \parallel t)$, where c_{text} is the syntactic surface form of the claim, a is the asserting author, and t is a temporal stamp (e.g. publication date). The Normalise function maps text to a logical-form expression or semantic vector representation, ensuring that logically identical claims receive the same hash signature [23].

Axiom 4 (claim identity stability): For all $c_1, c_2 \in C$, if $CCS(c_1) = CCS(c_2)$, then $\forall \phi, \text{Meaning}(c_1, \phi) \iff \text{Meaning}(c_2, \phi)$.

This axiom enforces that canonicalised claims are not merely syntactically similar but semantically equivalent in all model-theoretic interpretations of ϕ within BEWA's logical grammar.

Proposition 4 (disambiguation completeness under bounded ambiguity): Let N be the number of authors in the input corpus and k the average number of name-variants per author. Then for finite k and sufficient metadata (co-author vectors, ORCID coverage > 0.85), the CAID disambiguation

algorithm achieves convergence with high probability in $O(N \log N)$ time.

Proof. See Kim et al., for convergence properties of blocking-based author disambiguation, coupled with unique identifier overlays (e.g. DOIs and ORCIDs) [24]. Error rates empirically fall below 1.5% in corpora exceeding 10^6 entries.

Implementation-wise, each CAID is mapped to an evolving author profile, which records not only publication history but also replication success rate, citation diffusion, retraction record, and peer review activity. Each CCS is mapped to a version-controlled, context-enriched node in the epistemic graph, allowing BEWA to track revisions, contradictions, and semantic drift over time.

In sum, canonical author and claim identification provides the ontological substrate for epistemic accountability, belief traceability, and computational scalability. Without such formal anchoring, probabilistic inference over claims would be corrupted by aliasing, duplication, and incoherent attribution.

Metadata extraction and integrity validation

For BEWA to maintain epistemic reliability at scale, it must guarantee that all inferential operations are grounded in metadata that is both complete and verified. Metadata, in this context, refers to all structured information necessary to compute the credibility, provenance, and context of a scientific claim including, but not limited to, author identity, publication date, venue, citation relationships, funding declarations, methodological tags, and replication indicators. Incomplete or corrupted metadata threatens the stability of probabilistic reasoning and can result in spurious belief updates or the propagation of structurally invalid claims. This subsection formalises the extraction process as a mapping from source artefacts to structured tuples, and establishes axioms ensuring cryptographic integrity, schema completeness, and referential traceability.

Definition 6 (metadata record M): Let d be a document ingested into BEWA's system. Then the metadata record $M(d)$ is defined as: $M(d) := (CAID, CCS, DOI, t, V, F, R, S)$

where: CAID is the canonical author identifier; CCS is the canonical claim signature; DOI is the persistent digital object identifier; t is the publication timestamp; V is the venue vector (journal, impact metrics, editorial schema); F is funding disclosure; R is the replication status tag; S is the structural completeness flag (conformance to the metadata schema).

Axiom 5 (schema completeness constraint): A metadata record $M(d)$ is admissible if and only if:

$$\text{Complete}(\mathcal{M}) := \bigwedge_{f \in \{\text{CAID}, \text{CCS}, \text{DOI}, t, v, s\}} (f \neq \emptyset)$$

This ensures that no record lacking essential referential anchors is introduced into the belief network.

Axiom 6 (cryptographic integrity constraint): Let h_d be the hash of document d and σ_d the digital signature of its originator. Then: $\text{Verify}(d) = \text{SigCheck}(h_d, \sigma_d, \text{PubKey}_{\text{source}(d)}) = \text{True}$.

That is, every metadata record must be verifiably anchored to its source via cryptographic hash chains and digitally signed attestations from trusted identity providers (e.g., ORCID, Crossref, arXiv).

Proposition 5 (provenance immutability): If metadata $M(d)$ satisfies Axioms 5 and 6 at time t_0 , and if the document hash h_d is committed to a publicly verifiable ledger (e.g., IPFS, blockchain), then the provenance of d is immutable under the assumption of collision-resistance and ledger availability.

Proof. Direct application of cryptographic binding principles, where any tampering with the content of d will alter h_d , causing mismatch with the signed record [25]. Provided the public ledger is append-only and consensus-secured, this yields auditability and tamper-evidence.

BEWA's ingestion layer enforces these constraints using automated schema validation (e.g., JSON Schema against Crossref and PubMed records), digital signature verification tools (e.g., ORCID's signed assertions), and duplicate detection via simhash and minhash locality-sensitive hashing [26]. Failure to meet any condition results in rejection and logging of a metadata integrity fault.

Thus, the architecture guarantees that every belief within the system can be epistemologically anchored to a validated, complete, and immutable metadata structure ensuring computational soundness, inferential transparency, and audit integrity across the entire epistemic graph.

Claim Representation and Propositional Structure

This section defines the epistemologically disciplined schema by which scientific assertions are represented within the BEWA framework. The system is predicated on the understanding that language in scientific literature is often imprecise, context-dependent, and semantically overloaded. To render such data computationally tractable and logically analysable, BEWA translates these inputs into structured propositional claims abstract, semantically precise units that function as the minimal bearers of epistemic weight. Each claim is isolated from surrounding discursive noise, abstracted into a

canonical form, and endowed with logical integrity sufficient for independent evaluation and networked inference. These propositional units serve as the epistemic currency of the system: all weighting, belief revision, and cross-claim dependency is defined in terms of them.

Beyond structural parsing, each claim is embedded with a dense context map to capture domain relevance, experimental conditionals, and field-specific nuances. This contextualisation is critical: It guards against the epistemic error of false generalisation and enables the system to limit the scope of a proposition's influence based on declared or inferred domain boundaries. Claims are not timeless BEWA incorporates a versioning system that tracks the evolutionary history of any assertion, noting when and how its structure or interpretation has shifted. Temporal anchoring ensures that the system respects the chronological development of scientific knowledge, aligning belief weightings with the epistemi

Structured propositional claims

At the core of BEWA's epistemic model lies the formalisation of scientific assertions as structured propositional claims. Unlike natural language sentences, which are often ambiguous, context-sensitive, and syntactically irregular, structured propositional claims enable machine-tractable reasoning by encoding assertions in a logically well-formed, semantically normalised format. This structure facilitates belief assignment, dependency resolution, contradiction detection, and inferential propagation. The purpose of this subsection is to define the syntactic and semantic criteria for claim admissibility, to establish a compositional grammar for representing claims, and to demonstrate the formal soundness of the system's propositional encoding.

Definition 7 (structured propositional claim): A Structured Propositional Claim (SPC) is a tuple: $\text{SPC} = \langle \phi, \tau, \gamma \rangle$. Where ϕ is a well-formed formula in a domain-specific language L grounded in first-order logic; τ is a temporal index denoting the time of assertion or observation; γ is the contextual signature (ontology, experimental modality, statistical frame-work).

Axiom 7 (well-formedness of ϕ): The formula ϕ must satisfy the syntactic production rules of L : $\phi ::= P(t_1, \dots, t_n) \mid \neg \phi \mid \phi \wedge \phi \mid \phi \rightarrow \phi \mid \forall x \phi \mid \exists x \phi$

where P is a predicate symbol and t_i are typed terms over a domain ontology O .

Definition 8 (claim normalisation function): Let s be a natural language assertion extracted from a scientific document. Define $N(s)$ as the function that returns: $N(s) = \text{SPC} = \langle \phi, \tau, \gamma \rangle$,

where ϕ is derived via semantic parsing, entity disambiguation, and relation extraction [27].

Proposition 6 (injectivity of N on disambiguated inputs):

If s_1, s_2 are distinct scientific assertions with disjoint semantic parses under ontology O , then $N(s_1) \neq N(s_2)$.

Proof. Let s_1, s_2 be natural language sentences mapped to logical forms ϕ_1, ϕ_2 respectively via a pipeline that includes named entity recognition, syntactic parsing, and semantic role labelling. By assumption, s_1 and s_2 denote different referents in O . Then $\phi_1 \neq \phi_2$, and thus $N(s_1) \neq N(s_2)$ by construction. See Beltagy et al., and Liang et al., for similar injectivity guarantees under disambiguation assumptions [23,28].

The function N is implemented using a hybrid symbolic-neural parsing stack, where transformer-based models (e.g., SciBERT) generate candidate interpretations, which are then validated against ontological constraints and claim schemas using typed lambda calculus and ontology alignment [29]. Logical forms are grounded to probabilistic database schemas, enabling downstream inference over scientific claims as structured data.

Axiom 8 (epistemic decidability of ϕ): For any SPC $\langle \phi, \tau, \gamma \rangle$, there must exist a decision procedure D such that: $D(\phi) \in \{\text{verifiable, refutable, undecidable}\}$.

This guarantees that every claim entering the system is classifiable under a tractable epistemic status, supporting non-monotonic reasoning and dynamic belief assignment.

Structured propositional claims therefore function as the atomic elements of BEWA's reasoning calculus. By imposing logical formality, contextual grounding, and semantic tractability, the system ensures that belief manipulation operates over well-defined, auditable units mitigating ambiguity, enhancing comparability, and enabling precise epistemic operations across domains.

Contextual tagging and domain indexing

Scientific claims do not exist in isolation but are embedded within intricate domain-specific contexts that determine their scope, generalisability, and evidentiary strength. In BEWA, contextual tagging and domain indexing serve as orthogonal dimensions of claim normalisation and semantic disambiguation. Without contextualisation, semantically identical surface forms may yield radically different epistemic weights depending on underlying assumptions, methodological paradigms, or disciplinary boundaries. This subsection formalises the contextual tagging mechanism and defines a topological indexing structure over domain ontologies, ensuring that every structured propositional claim is correctly situated within its appropriate epistemic subspace.

Definition 9 (contextual tag set Γ): Let $\text{SPC} = \langle \phi, \tau, \gamma \rangle$ be a structured propositional claim. The contextual tag γ is an element of Γ , where: $\Gamma := C \times M \times S$, where C is the scientific concept ontology (e.g., MeSH, UMLS, ACM CCS); M denotes methodological descriptors (e.g., RCT, observational, meta-analysis); S represents statistical framing (frequentist, Bayesian, non-parametric, etc.).

Each claim thus carries a tripartite tag structure identifying its conceptual anchor, methodological provenance, and inferential semantics. This resolves cases where two identical predicates (e.g., "X increases Y") may differ in meaning if one arises from a randomised controlled trial and another from correlational modelling.

Axiom 9 (semantic stratification constraint): Two claims $c_1 = \langle \phi_1, \tau_1, \gamma_1 \rangle$ and $c_2 = \langle \phi_2, \tau_2, \gamma_2 \rangle$ may only be treated as epistemic equivalents if and only if: $\phi_1 \equiv \phi_2$ and $\gamma_1 = \gamma_2$.

This ensures that contextual differences are preserved at the semantic level, avoiding illicit aggregation or conflation of non-commensurable results.

Definition 10 (domain indexing function δ): Let D be the space of disciplinary domains (e.g., neuroscience, econometrics, bioinformatics). Define: $\delta : \Gamma \rightarrow 2^D$

Such that $\delta(\gamma)$ returns the minimal closed domain set in which the claim is epistemically coherent. This mapping is constructed using domain ontologies and citation graph embeddings [30].

Proposition 7 (transitive coherence via domain overlap): Let c_1, c_2 be two claims such that $\delta(\gamma_1) \cap \delta(\gamma_2) \neq \emptyset$. Then a coherence-preserving belief transformation is possible between c_1 and c_2 .

Proof. Follows from overlap in domain-indexed epistemic subgraphs. See Zeng et al., for graph embeddings over citation ontologies that support semantic transfer and influence modelling across neighbouring disciplines [31].

BEWA implements Γ via a multi-level tagging pipeline. Concepts are assigned using curated term matchers (e.g. MeSH taggers), methodological types are inferred from structured abstracts using neural classifiers trained on annotated corpora and statistical paradigms are extracted via pattern-matching over model descriptors and inference statements [32]. The domain function δ is implemented using vector-space projection over pretrained knowledge graph embeddings (e.g. SPECTER, SciGraph).

In summary, contextual tagging and domain indexing ensure that BEWA's belief updates respect the epistemic scope, semantic constraints, and methodological heterogeneity of

scientific claims preserving fidelity in reasoning and enabling high-resolution control over cross-domain inferential transfer.

Versioning and temporal anchoring of claims

Scientific claims are not static artefacts but dynamic assertions whose semantic content, evidentiary support, and inferential implications may evolve over time. BEWA addresses this epistemic mutability through a formal mechanism of versioning and temporal anchoring, which allows each structured propositional claim to exist as a temporally indexed sequence of revisions. This design not only enables historical traceability but ensures that belief updates are grounded in temporally coherent inference resisting anachronistic reasoning and preserving the causal integrity of scientific progression.

Definition 11 (claim version chain $V\phi$): Let ϕ be the logical core of a structured propositional claim. Then:

$$V_\phi := \left\{ \langle \phi^t, \gamma^t, \tau^t \rangle \right\}_{t_0 \leq t \leq T}$$

where each ϕ^t represents the version of the claim valid at time τ^t , with associated context tag γ^t . This sequence is totally ordered by timestamp: $\tau^{t_i} < \tau^{t_j} \Leftrightarrow t_i < t_j$

Axiom 10 (temporal monotonicity): For any claim ϕ and its version chain $V\phi$, no semantic regression is permitted: $\forall t_i < t_j$, if $\phi^{t_j} \vdash \phi^{t_i}$, then $\phi^{t_j} \equiv \phi^{t_i}$

This prohibits weakening or reversal of propositional commitment without explicit contradiction annotation, preventing silent epistemic erosion.

Definition 12 (temporal anchor map α): Let C be the set of all canonical claims. Define the anchor function: $\alpha : C \rightarrow T \times T$. Where $\alpha(c) = (\tau_{\text{start}}, \tau_{\text{end}})$ is the interval during which the claim version ϕ^t is active in the system's inferential graph.

This temporal anchoring enables BEWA to compute time-sensitive belief networks, where evidence, citations, or replications are only considered admissible if they fall within the claim's active interval.

Proposition 8 (temporal coherence in belief propagation): Let c_1, c_2 be two claims such that c_1 supports c_2 . Then temporal coherence requires: $\alpha(c_1). \tau_{\text{end}} \geq \alpha(c_2). \tau_{\text{start}}$

Proof. If c_1 's influence terminates before c_2 emerges, then c_2 cannot be justifiably inferred from c_1 without violating causality. This is a constraint on belief propagation scheduling, implemented in temporal DAG logic [33].

BEWA implements claim versioning through cryptographically chained hashes: each version ϕ^t is hashed together with its predecessor ϕ^{t-1} and signed by the system to

form an immutable update ledger. Temporal anchors are encoded using RFC 3339 timestamps, and update intervals are synchronised with publication, retraction, or amendment records extracted via Crossref, Retraction Watch, and publisher APIs.

Implementation note: All belief calculations on a given version of a claim are conducted with respect to the temporal interval in which that version is valid. Belief shifts caused by replications, contradictions, or epistemic reclassifications are constrained by temporal logic enforced at the graph layer. This prevents retrospective contamination of earlier belief states and ensures soundness in longitudinal inference.

Thus, versioning and temporal anchoring ensure that BEWA's epistemic graph not only represents what is believed and how strongly, but also when and in what form establishing a foundation for diachronic reasoning, evidence lifecycle modelling, and historiographic auditability.

Bayesian Weighting and Belief Updating

This section sets out the core inferential machinery of BEWA, wherein each scientific claim is evaluated through a principled Bayesian framework. At its heart, the system operationalises belief not as a binary metric of truth or falsehood, but as a dynamically updated probability that reflects the current weight of evidence. Every claim enters the system with an initial prior contextualised by its authorship, venue of publication, domain-specific norms, and the historical performance of associated entities. This prior is not arbitrarily assigned, but computed *via* a hierarchically conditioned model incorporating credibility, venue reliability, and baseline epistemic plausibility. From this baseline, each subsequent piece of evidence whether replication, contradiction, citation, or decay is treated as an updating factor in accordance with Bayesian conditionalisation.

Belief updating in BEWA is both modular and recursive. It incorporates replication success, citation influence (modulated by temporal decay and domain saturation), and epistemic counterweighting in the face of contradictions. The system does not naively reward frequency or visibility, but applies a critical filter to distinguish between epistemic endorsement and discursive noise. In addition, BEWA includes a decay mechanism that progressively reduces confidence in isolated or unreplicated claims over time, while simultaneously increasing sensitivity to emerging corroborations. The framework avoids premature convergence by maintaining uncertainty in the face of partial or ambiguous evidence and only reinforces belief where clear, replicated, and semantically consistent data exists.

The following subsections articulate the procedures for establishing priors, executing belief updates, processing contradictions, and enforcing temporal reassessment through probabilistic decay.

Initial prior formulation

In Bayesian epistemology, belief begins with a prior: a quantified estimate of the plausibility of a proposition in the absence of specific observational evidence. BEWA must assign initial prior probabilities to structured propositional claims upon ingestion, before any replication, contradiction, or citation-based update occurs. This prior is not drawn from subjective estimation but is computed as a function of measurable structural and epistemic properties of the claim's origin, including authorial credibility, publication venue, methodological design, and domain frequency statistics. This subsection formalises the prior function and proves its compliance with Bayesian consistency and epistemic calibration constraints.

Definition 13 (prior function π): Let $c = \langle \phi, \tau, \gamma \rangle$ be a structured propositional claim. The prior probability assigned to c is defined as: $\pi(c) := P(\phi \mid \gamma, \theta) = f(A(c), V(c), M(c), D(\phi))$, where: $A(c)$ is the authorial trust score (see §5.1); $V(c)$ is the venue credibility index; $M(c)$ is the methodological rigour metric (e.g., presence of preregistration, sample size adequacy); $D(\phi)$ is the historical base rate for similar claims in domain $\delta(\gamma)$.

Axiom 11 (probability axiom compliance): The prior function π satisfies: $\forall c, 0 \leq \pi(c) \leq 1$. Furthermore, for any mutually exclusive claims c_1, \dots, c_n within a disjoint claim partition $C' \subset C$:

$$\sum_{i=1}^n \pi(c_i) \leq 1$$

Definition 14 (authorial trust score $A(c)$): $A(c) := \sigma(\alpha_1 R_a + \alpha_2(1 - \rho_a) + \alpha_3 \log(1 + v_a))$, where: R_a is the replication success rate of author a ; ρ_a is the retraction frequency; v_a is the citation-normalised publication count; σ is the logistic squashing function to bound scores in $(0,1)$.

Proposition 9 (bounded variance of π under controlled inputs): Assume bounded author trust $A(c) \in (\epsilon, 1 - \epsilon)$ for $\epsilon > 0$ and fixed $M(c), V(c)$. Then the variance of $\pi(c)$ across similar claims is bounded.

Proof. Follows from the convexity of σ under bounded Lipschitz conditions over the input space, combined with bounded entropy in domain frequency distributions $D(\phi)$. Empirical results in Graves et al., demonstrate prior stability in similar hierarchical Bayesian settings [34].

Implementation note: In BEWA, each prior is computed during claim ingestion using a calibrated scoring function tuned on a corpus of validated claims from high-reliability domains (e.g., Cochrane Database, Nature Human Behaviour, NeurIPS reproducibility track). Venue credibility scores are computed from impact-adjusted replication rates, while $D(\phi)$ is estimated from smoothed claim frequency distributions indexed by contextual tags γ [35].

Axiom 12 (neutrality under ignorance): If $A(c), V(c), M(c), D(\phi)$ are all absent or undefined, then:

$$\pi(c) := \frac{1}{2}$$

This enforces a non-informative prior consistent with Laplacean indifference, ensuring epistemic neutrality in the absence of structural asymmetries.

By grounding priors in empirical author and venue data, methodological metadata, and domain frequency statistics, BEWA establishes a coherent epistemic starting point for all claims. The prior function π is not merely an artefact of convenience it is a mathematically principled device that constrains downstream belief trajectories to respect both epistemic caution and contextual informativeness.

Evidence-based posterior updating

In accordance with the Bayesian framework adopted by BEWA, posterior beliefs over scientific claims are updated through the application of Bayes' Theorem as new evidence is ingested. Evidence may take the form of citations, replications, contradictions, or derivations, each of which carries a quantifiable influence on the belief assigned to a structured propositional claim. Posterior updating must not only conform to the laws of probability, but also preserve inferential consistency, causal ordering, and network-level epistemic coherence.

Definition 15 (posterior belief): Let $c = \langle \phi, \tau, \gamma \rangle$ be a claim with prior $\pi(c) = P(\phi)$ and observed evidence e . The posterior belief $P(\phi \mid e)$ is given by:

$$P(\phi \mid e) = \frac{P(e \mid \phi) \cdot P(\phi)}{P(e)} \quad \text{provided } P(e) > 0$$

Axiom 13 (conditional independence of evidence streams): Let e_1, \dots, e_n be evidence events such that:

$$\forall i \neq j, \quad P(e_i \mid \phi, e_j) = P(e_i \mid \phi)$$

Then the joint likelihood is:

$$P(e_1, \dots, e_n \mid \phi) = \prod_{i=1}^n P(e_i \mid \phi)$$

This enables incremental belief updating over independent evidence observations.

Definition 16 (evidence likelihood function L): Each evidence unit e is associated with a type $\epsilon \in \{\text{replication, citation, contradiction, endorsement}\}$. Define the likelihood contribution of e to claim c as:

$$\mathcal{L}_e(e, c) := \begin{cases} \lambda_+ \cdot \sigma(q(e)), & \epsilon \in \{\text{replication, endorsement}\} \\ \lambda_- \cdot \sigma(-q(e)), & \epsilon = \text{contradiction}, \\ \eta \cdot \sigma(q(e)), & \epsilon = \text{citation}, \end{cases}$$

where $q(e)$ is the quality score of the evidence (e.g., journal credibility, sample size), σ is the sigmoid function, and λ_+ , λ_- , η are hyperparameters calibrated per domain.

Proposition 10 (monotonicity of posterior updates): Let e be a piece of evidence supporting ϕ with positive likelihood ratio:

$$\frac{P(e \mid \phi)}{P(e \mid \neg\phi)} > 1$$

Then $P(\phi \mid e) > P(\phi)$.

Proof. Follows directly from Bayes' Theorem:

$$P(\phi \mid e) = \frac{P(e \mid \phi) \cdot P(\phi)}{P(e \mid \phi) \cdot P(\phi) + P(e \mid \neg\phi) \cdot (1 - P(\phi))},$$

and the assumption implies that the numerator grows faster than the denominator.

Definition 17 (cumulative posterior update): Given a sequence of n evidence items $E = \{e_1, \dots, e_n\}$, define:

$$P(\phi \mid E) = \frac{\prod_{i=1}^n \mathcal{L}(e_i, \phi) \cdot \pi(c)}{\prod_{i=1}^n \mathcal{L}(e_i, \phi) \cdot \pi(c) + \prod_{i=1}^n \mathcal{L}(e_i, \neg\phi) \cdot (1 - \pi(c))}$$

Implementation note: BEWA's evidence ingestion pipeline scores each e using domain-specific evaluation functions (e.g., replication power, p-value correction, venue trust factor). These are normalised and aggregated into L using a neural calibration layer that ensures well-calibrated uncertainty estimates, thereby maintaining the interpretability of posterior shifts [36].

Axiom 14 (epistemic regularisation): Posterior updates are bounded within a temporal smoothing window: $|P(\phi \mid E_t) - P(\phi \mid E_{t-1})| \leq \delta$, for a fixed $\delta > 0$, to avoid belief volatility induced by low-confidence or adversarially injected evidence.

Evidence-based posterior updating in BEWA thus adheres to a strict formal framework: conditionally independent evidence is multiplicatively incorporated, update magnitude is modulated by evidence quality and type, and posterior trajectories are smoothed to reflect rational epistemic commitment. This ensures that beliefs are neither overfitted to noisy evidence nor underreactive to strong, reproducible support.

Contradiction handling and counter-evidence processing

A principled epistemic system must account not only for confirmatory evidence but also for disconfirmatory inputs namely, contradictions and counter-evidence. In Bayesian terms, this corresponds to negative likelihood ratios that reduce posterior belief in a proposition. Within BEWA, contradiction handling is formalised as the systematic down-weighting of claims upon presentation of semantically aligned but empirically conflicting evidence. The system distinguishes between strict logical contradiction and probabilistic disconfirmation, each with defined epistemic and computational implications.

Definition 18 (contradictory claim): Let $c_1 = \langle \phi, \tau_1, \gamma_1 \rangle$ and $c_2 = \langle \psi, \tau_2, \gamma_2 \rangle$ be claims. Then c_2 is a contradiction of c_1 if:

$$\text{Contradicts}(c_1, c_2) := \phi \models \neg\psi \quad \text{or} \quad \mathbb{P}(\phi \mid \psi) < \theta_c$$

for some threshold $\theta_c \in (0, 0.5)$ under a calibrated semantic contradiction detector.

Definition 19 (counter-evidence tuple): A unit of counter-evidence e^- against claim c is a tuple: $e^- := \langle \phi^-, \kappa, \rho \rangle$, where: ϕ^- is a contradicting statement (in logical or empirical form); κ is the replication consistency of the contradiction; ρ is the domain relevance alignment between c and e^- .

Axiom 15 (asymmetry of contradiction influence): Let e^+ and e^- be evidence in favour and against ϕ , respectively. Then:

$$\text{If } \kappa^- > \kappa^+ \text{ and } \rho^- \approx \rho^+, \text{ then } |\Delta P^-(\phi)| > |\Delta P^+(\phi)|$$

This reflects the epistemic principle that strong, reproducible contradictions should carry more weight than isolated confirmations a formalisation of Popperian falsifiability adapted into probabilistic calculus.

Definition 20 (contradiction likelihood function L_-):

$$\mathcal{L}_-(e^-, c) = \lambda_- \cdot \sigma(\kappa \cdot \rho \cdot q(e^-))$$

where $q(e^-)$ is the internal quality score of the counter-evidence and λ_- is a domain-calibrated influence constant.

Proposition 11 (posterior downdate under contradiction): Let $P(\phi \mid E)$ be the posterior belief in claim ϕ given evidence set E , and e^- a new contradiction. Then the updated posterior is:

$$P'(\phi \mid E \cup \{e^-\}) = \frac{\mathcal{L}_-(e^-, \phi) \cdot P(\phi \mid E)}{\mathcal{L}_-(e^-, \phi) \cdot P(\phi \mid E) + (1 - \mathcal{L}_-(e^-, \phi)) \cdot (1 - P(\phi \mid E))}$$

Proof. Follows directly from Bayes' Theorem with updated likelihood ratio reflecting counter-evidential influence.

Axiom 16 (non-retroactivity of future contradictions): If a contradiction e^- is timestamped at $\tau > \tau_c$, then it must not alter belief states computed at any $t < \tau$. This enforces the temporal integrity of belief trajectories and prevents epistemic contamination of past inferences.

Implementation note: Contradiction detection in BEWA is implemented via transformer-based entailment models (e.g., SciBERT fine-tuned on SNLI/SciNLI) with ontology-constrained semantic alignment. Empirical inconsistencies are flagged using statistical test comparisons (e.g., incompatible effect sizes or reversed directionality at $\alpha=0.05$), and their influence scaled based on methodological rigour, replication consistency, and domain proximity. Conflicting claims are not deleted but annotated with contradiction tags, preserving interpretability and enabling downstream audit.

Thus, contradiction handling in BEWA is not merely a passive attenuation of belief but a structured, principled recalibration mechanism. It enforces epistemic accountability, guards against dogmatic entrenchment, and aligns the system with rational principles of falsifiability and evidence-sensitive reasoning.

Bayesian decay and temporal reassessment

Belief in a proposition, in the absence of continual evidential reinforcement, should gradually diminish to reflect epistemic uncertainty introduced by the passage of time. BEWA encodes this principle through a formal decay mechanism, grounded in Bayesian logic and information theory, which reduces posterior certainty over time unless claims are explicitly reaffirmed through replication or re-evaluation. This temporal reassessment guards against unwarranted epistemic inertia and enforces a dynamic equilibrium in the system's belief network.

Definition 21 (decay function δ_t): Let $P_t(\phi)$ denote the posterior probability of a claim ϕ at time t , and $\Delta t = t - t_0$ the elapsed time since the last reinforcement (e.g., replication, high-confidence citation). Then: $P_t(\phi) := \delta_t(P_{t_0}(\phi)) = P_{t_0}(\phi) \cdot \exp(-\lambda \cdot \Delta t)$, where λ is the decay rate parameter, domain- and context-specific.

Axiom 17 (exponential temporal decay): In the absence of new evidence, belief in ϕ degrades according to:

$$\lim_{\Delta t \rightarrow \infty} P_t(\phi) = 0$$

This reflects the principle that unreplicated or unaudited claims should asymptotically approach epistemic nullity.

Definition 22 (reinforcement event R): A reinforcement event is any evidence unit $e \in E$ such that: $L(e, \phi) > \theta_r$, where θ_r is a domain-specific reinforcement threshold. When such an event occurs at time t_r , the decay counter resets:

$$P_{t_r}(\phi) := \frac{P_{t_r^-}(\phi) \cdot \mathcal{L}(e, \phi)}{P_{t_r^-}(\phi) \cdot \mathcal{L}(e, \phi) + (1 - P_{t_r^-}(\phi)) \cdot \mathcal{L}(e, \neg\phi)}$$

Proposition 12 (information-theoretic justification for decay): Let $H_t(\phi)$ denote the entropy of belief at time t . Then:

$$\frac{dH_t(\phi)}{dt} > 0 \quad \text{if no new evidence is introduced.}$$

Proof. Entropy $H(p) = -p \log p - (1 - p) \log(1 - p)$ increases as $p \rightarrow 0.5$, which is the attractor of the decay function δ_t in the absence of updates. Hence belief becomes increasingly uncertain over time.

Axiom 18 (monotonic temporal entropy): Let ϕ be a claim with no contradictory or reinforcing evidence in Δt . Then: $\forall t_1 < t_2, H_{t_1}(\phi) \leq H_{t_2}(\phi)$.

This enforces that epistemic uncertainty never decreases unless justified by evidentiary intervention.

Implementation note: BEWA computes decay continuously as a background process indexed by claim age and domain decay constants. For example, fast-moving empirical domains (e.g., oncology, machine learning) apply higher λ , while theoretical domains (e.g., mathematics, logic) decay at negligible rates. Claims with long decay intervals may still trigger reassessment flags if their influence pervades downstream belief chains.

In practice, all belief propagation routines are time-aware. Each inference step verifies whether a claim's current posterior reflects its decay-adjusted status, ensuring time-consistent epistemic integrity. Claims falling below a minimum trust threshold ϵ are marked stale and excluded from active inference unless revalidated.

Corollary (temporal reversibility via replication): Decay is not logically irreversible: Any high-confidence replication e resets Δt and re-establishes $P_t(\phi)$ without information loss.

Bayesian decay and reassessment thus serve as epistemic entropy regulators within BEWA, preventing the ossification of

stale claims and maintaining the responsiveness of belief states to the evolving structure of scientific knowledge.

Authorial Credibility and Impact Modelling

This section introduces the mechanisms by which BEWA evaluates and incorporates the epistemic weight of individual authors within its inferential framework. Unlike conventional systems that treat authorship as a bibliographic footnote or merely a source of provenance, BEWA treats authorial identity as a dynamic epistemic signal. Each author contributes not only specific claims but an accumulated pattern of credibility shaped by their historical accuracy, involvement in replication studies, susceptibility to retraction, and degree of engagement with rigorous peer review. These patterns are quantitatively scored and directly modulate the priors of future claims issued under their canonical identity. In this way, authorship becomes both a source of risk and trust: a vector through which epistemic integrity is either strengthened or diluted.

The architecture of BEWA's credibility engine recognises that impact must be distinguished from popularity. Authors who generate high-citation work may still score poorly if that work fails to replicate or accumulates contradictions. Conversely, less visible scholars who consistently produce durable, well-supported research are assigned epistemic weight in excess of their surface-level prominence. The system penalises epistemic malpractice such as irreproducibility, selective reporting, or frequent retraction while rewarding transparent and rigorous engagement with the scientific community. Authorial metrics are not static; they evolve with each new publication, review, or correction, and they propagate their influence into the weighting of every claim that author touches. The subsections below formalise the computation of author scores, the long-term impact of retraction history, and the significance of peer review participation within this credibility calculus.

Author score calculation

In BEWA, the epistemic weight assigned to a claim is partially inherited from its author's historical record. This record, formalised as an author score, quantifies the credibility of individual researchers based on their cumulative publication history, replication rate, retraction profile, and peer recognition. The author score serves as an input to the prior function π (Definition 13), acting as a probabilistic proxy for the reliability of future assertions. This subsection formalises the author scoring function, proves its boundedness and monotonicity, and ensures that it is resistant to gaming through citation inflation or strategic publishing.

Definition 23 (author score function A): Let $a \in A$ be a canonical author identifier. Then the author score $A(a) \in (0, 1)$ is given by: $A(a) := \sigma(\beta_1 \cdot r_a + \beta_2 \cdot \log(1 + v_a) - \beta_3 \cdot \rho_a + \beta_4 \cdot \mu_a)$, where: r_a is the verified replication rate of a 's claims; v_a is the field-normalised citation count; ρ_a is the retraction rate (fraction of retracted outputs); μ_a is a peer review engagement score (editorial or verified reviewer roles); $\sigma(x) = \frac{1}{1+e^{-x}}$ is the sigmoid function for bounding.

Axiom 19 (credibility monotonicity): For all $a \in A$, $\partial A(a)/\partial r_a > 0$, $\partial A(a)/\partial \mu_a > 0$, $\partial A(a)/\partial \rho_a < 0$. This ensures that credibility increases with empirical replication and peer engagement, and decreases with retraction incidence.

Proposition 13 (boundedness and stability): For bounded input parameters and fixed β coefficients, the function $A(a)$ is Lipschitz-continuous and satisfies: $0 < A(a) < 1$, $\forall a$ such that $v_a < \infty$.

Proof. The sigmoid function maps \mathbb{R} to $(0, 1)$ and is Lipschitz with constant $1/4$ over its full domain. Each feature term is finite and bounded (logarithmic or proportional), and thus their linear combination is bounded, implying the result.

Definition 24 (replication rate r_a):

$$r_a := \frac{|\text{Claims}_a^{\text{replicated}}|}{|\text{Claims}_a^{\text{testable}}|}$$

where only testable (empirical, non-theoretical) claims are included in the denominator. Replication is counted only when validated through high-confidence reproductions (cf. §6.2).

Definition 25 (retraction rate ρ_a):

$$\rho_a := \frac{|\text{Retracted}_a|}{|\text{TotalPublications}_a|}, \quad \text{with } 0 \leq \rho_a \leq 1$$

Definition 26 (peer review engagement μ_a):

$$\mu_a := \frac{|\text{VerifiedReviews}_a| + \omega \cdot |\text{EditorialRoles}_a|}{\log(1 + |\text{YearsActive}_a|)}$$

where ω adjusts editorial impact relative to review count and the denominator normalises by experience to prevent bias toward seniority.

Implementation note: BEWA integrates with ORCID, Publons, and Crossref to obtain data for v_a , μ_a , and ρ_a , and uses replication registries (e.g. Curate Science, ReplicationWiki) to compute r_a . All score components are versioned and updated upon the appearance of new publications, reviews, or corrections.

Axiom 20 (anti-gaming constraint): The marginal gain in $A(a)$ from increasing v_a is logarithmic and asymptotically flat:

$$\lim_{\nu_a \rightarrow \infty} \frac{\partial \mathcal{A}(a)}{\partial \nu_a} \rightarrow 0$$

This ensures resistance to citation inflation and self-citation abuse.

In sum, the author score $\mathcal{A}(a)$ is a formal, bounded, multidimensional estimator of scholarly reliability. It integrates empirical performance, reputational standing, and community participation to inform prior belief assignment in a manner that is mathematically stable, auditable, and epistemically justified.

Track record and retraction influence

The epistemic credibility of an author is not static but accumulates and may deteriorate over time through their published scientific output. In BEWA, an author's track record is modelled as a longitudinal sequence of claim events, from which performance metrics such as replication consistency, correction frequency, and retraction density are computed. This allows BEWA to condition belief formation not only on the present claim but also on the statistical integrity of an author's prior assertions. Of particular importance is the influence of retractions, which function as high-penalty negative evidence with persistent downstream impact on author trust and claim priors.

Definition 27 (author track record \mathcal{T}_a): For author $a \in A$, define the track record as:

$$\mathcal{T}_a := \{ \langle \phi_i, t_i, r_i, \text{status}_i \rangle \}_{i=1}^{N_a}$$

where each tuple represents a claim ϕ_i made at time t_i , with replication result $r_i \in \{0, 1\}$ and status $\in \{\text{active, corrected, retracted}\}$.

Axiom 21 (monotonic penalty of retraction): Let a have a claim $c = \langle \phi, \cdot, \cdot \rangle$ retracted at t_r . Then for all $t > t_r$, the author's credibility score $\mathcal{A}(a)$ must decrease or remain unchanged:

$$\mathcal{A}_t(a) \leq \mathcal{A}_{t_r}(a)$$

Definition 28 (retraction penalty function \mathcal{R}_a): Let $|\text{Retracted}_a|$ be the number of retractions and $|\text{Total}_a|$ the total number of publications. Define:

$$\mathcal{R}_a := \gamma \cdot \frac{|\text{Retracted}_a|}{1 + \log(1 + |\text{Total}_a|)}$$

with $\gamma > 0$ controlling the steepness of reputational decay. This sublinear normalization ensures early-career authors are not disproportionately penalised.

Proposition 14 (replicability-weighted recovery): Let a 's retraction-adjusted score be: $\mathcal{A}_r(a) := \mathcal{A}(a) - \mathcal{R}_a + \eta \cdot r_a$, where r_a is the cumulative replication ratio and η a tunable recovery factor. Then:

$$\text{If } r_a \rightarrow 1, \quad \lim_{|\text{Retracted}_a| \ll |\text{Total}_a|} \mathcal{A}_r(a) \rightarrow \mathcal{A}(a).$$

Proof. As $r_a \rightarrow 1$, the recovery term $\eta \cdot r_a$ counteracts the penalty \mathcal{R}_a , particularly when the retraction count is negligible. Boundedness of \mathcal{R}_a under the logarithmic denominator ensures $\mathcal{A}_r(a)$ remains finite and recoverable.

Definition 29 (retraction propagation suppression): Let ϕ be a claim authored by a and retracted at t_r . Then:

$$\forall c_i \text{ such that } \phi \rightarrow c_i, \quad \text{BeliefWeight}(c_i) \stackrel{!>t_r}{\leftarrow} \min(\text{BeliefWeight}(c_i), \epsilon)$$

where ϵ is a lower bound threshold. This ensures no claim retains high confidence if it descends from a retracted parent without independent support.

Implementation note: BEWA integrates with Retraction Watch and Crossref's Crossmark metadata feeds to detect retraction events. Correction notices are distinguished from retractions, with lower penalty weight. Author track records are maintained as immutable audit logs, and the penalty function is evaluated recursively to determine cascading credibility loss in co-authored publications.

Axiom 22 (persistent negative weight of retraction): Retractions are never forgotten but may be contextually overcome. Formally,

$$\inf_{t > t_r} \mathcal{A}_t(a) < 1 \quad \text{for any } t_r \text{ such that } \text{Retracted}_a \neq \emptyset.$$

This ensures that retractions impose an irreversible epistemic cost, even if reputational recovery is possible through sustained accuracy and replication. BEWA thereby upholds the principle that scientific trust must be earned continuously and is susceptible to justified revision when the record warrants it.

Peer review engagement metrics

While publication record and replication history provide empirical evidence for the quality of an author's output, BEWA supplements these with metrics derived from an author's participation in the scientific community as a reviewer, editor, or contributor to structured quality assurance processes.

These peer review engagement metrics serve as auxiliary indicators of epistemic diligence, institutional trust, and expertise recognition. By incorporating these into the author scoring function (cf. §5.1), BEWA ensures that scholarly contributions beyond authorship are also reflected in credibility assignments.

Definition 30 (peer review engagement score μ_a): Let $a \in A$ be an author. Then the engagement score is defined as:

$$\mu_a := \frac{\theta_1 \cdot |\mathcal{R}_a| + \theta_2 \cdot |\mathcal{E}_a|}{1 + \log(1 + Y_a)}$$

where $|\mathcal{R}_a|$ is the count of verified, completed peer reviews attributed to a ; $|\mathcal{S}_a|$ is the number of documented editorial roles held (e.g., editor-in-chief, associate editor); Y_a is the number of active years since first publication; $\theta_1, \theta_2 \in \mathbb{R}_{\geq 0}$ are tunable weights (typically $\theta_2 > \theta_1$ to reflect greater epistemic responsibility).

Axiom 23 (normalised seniority adjustment): The denominator term prevents inflation of μ_a in long-career authors who engage minimally with peer review over time. Thus:

$$\mu_a \xrightarrow{|\mathcal{R}_a|, |\mathcal{E}_a| \rightarrow 0} 0 \quad \text{even as } Y_a \rightarrow \infty$$

Definition 31 (review quality adjustment μ_a^*): Let each review $r \in \mathcal{R}_a$ have a quality score $q_r \in (0, 1)$ based on editor-provided or community-evaluated metrics (e.g., depth, punctuality, constructiveness). Then:

$$\mu_a^* := \frac{\theta_1 \cdot \sum_{r \in \mathcal{R}_a} q_r + \theta_2 \cdot |\mathcal{E}_a|}{1 + \log(1 + Y_a)}.$$

Proposition 15 (stability of μ_a^* under incomplete review data): If q_r is missing for a subset $\mathcal{R}'_a \subseteq \mathcal{R}_a$, and $q_r = 0.5$ is imputed (maximum entropy prior), then:

$$|\mu_a^* - \mu_a| \leq \theta_1 \cdot \frac{|\mathcal{R}'_a|}{1 + \log(1 + Y_a)} \cdot 0.5$$

Proof. Imputed reviews contribute at most 0.5 per unit to the numerator, and the score remains bounded under the normalised scaling. See Shah et al., for review quality scoring under partial observability [37].

Definition 32 (institutional trust overlay): Let $\tau_a \in [0, 1]$ denote an institutional trust index for a based on confirmed service in high-trust journals or conferences (e.g., Nature, NeurIPS, The Lancet). Then define the final engagement term as:

$$\mu_a^\dagger := \tau_a \cdot \mu_a^*,$$

which up-weights peer review influence when validated by trusted venues.

Implementation note: Review metadata is drawn from services such as Publons, ORCID peer review activity feeds, and journal editorial APIs. Editor-confirmed reviews are weighted more heavily than self-reported entries. To avoid gaming, only verified contributions with timestamps and linked manuscript identifiers are admitted. Missing data is flagged but not penalised beyond zero contribution.

Axiom 24 (bounded review influence): Peer review engagement cannot exceed a fixed maximum contribution $\mu_a^\dagger < \mu_{max}$ to prevent inflation from prolific low-impact reviewing.

BEWA thus treats peer review as a first-class epistemic signal not sufficient alone to establish credibility, but essential to contextualise an author's role in maintaining scientific integrity. Through weighted aggregation, temporal normalisation, and venue stratification, peer engagement becomes a structured and auditable source of epistemic trust.

Citation and Replication Framework

This section formalises the twin pillars of epistemic reinforcement within the BEWA system: Citation analysis and replication scoring. In conventional bibliometric systems, citations are often treated as indicators of prestige or influence; BEWA rejects this conflation. Instead, citations are deconstructed as weighted epistemic endorsements context-sensitive signals that must be interpreted within temporal, semantic, and disciplinary boundaries. The system assesses not merely the quantity of citations but their origin, relevance, decay, and semantic congruity with the original claim. This allows for the discrimination between superficial citations and those that represent genuine affirmations of an assertion's scientific merit. Each citation contributes to the belief calculus only to the extent that it reflects rigorous uptake and considered engagement.

Replication, by contrast, is treated as the epistemic gold standard. Unlike citation, which may reflect popularity or discursive inertia, replication constitutes a direct test of claim durability under independent conditions. BEWA encodes each replication event as a distinct evidentiary input, scored according to methodological fidelity, effect size congruence, and semantic alignment with the original claim. Successful replications raise a claim's posterior significantly, while failed replications especially when consistent across studies precipitate rapid belief retraction. Additionally, the system monitors contradictions, mapping their propagation across the epistemic network and adjusting neighbouring claims as appropriate. Contradictions are not treated as binary negations but as

probabilistic disruptors with network-wide consequences. The following subsections explicate how citation decay is modelled, how replications are evaluated for semantic and statistical integrity, and how the contradiction graph dynamically alters local and global belief states.

Citation weighting and decay functions

Citations in scientific literature represent not only informational dependency but also communal endorsement. However, raw citation counts often distort epistemic relevance due to temporal inflation, field-specific density, and citation cascades. BEWA introduces a calibrated citation weighting system grounded in time-discounted Bayesian import, where each citation contributes to the belief in a claim proportionally to its contextual and temporal provenance.

Definition 33 (citation influence function C_i): Let ϕ_i be a claim cited n times by documents $\{d_1, \dots, d_n\}$, each at time t_j , with each citing document d_j assigned credibility $A(d_j)$. Then:

$$C_i := \sum_{j=1}^n \delta(t_j) \cdot A(d_j)$$

where $\delta(t_j)$ is a temporal decay function normalised to $(0, 1)$.

Definition 34 (decay function δ): Let t_j be the timestamp of citation j , and T the current system time. Then: $\delta(t_j) := \exp(-\lambda(T - t_j))$, where $\lambda > 0$ is the decay constant. A higher λ accelerates obsolescence, reflecting fast-moving fields.

Axiom 25 (monotonic citation decay): The influence of an unreinforced citation must decrease over time: $\forall t_1 < t_2, \delta(t_1) > \delta(t_2)$.

Proposition 16 (stability under citation saturation): Let ϕ_i be cited $n \rightarrow \infty$ times from low-authority documents. Then:

$$\lim_{n \rightarrow \infty} C_i < \infty \quad \text{if} \quad \sup_j A(d_j) < M$$

Proof. Since $\delta(t_j) \cdot A(d_j) < M$ for all j , the sum converges if $\delta(t_j)$ decays faster than $1/j$, which is true for $\lambda > 0$ under exponential decay.

Definition 35 (Contextual Citation Modifier χ_j): For each citation j , define:

$$\chi_j := \begin{cases} 1 & \text{if citation is supportive} \\ -1 & \text{if citation is critical or refuting} \\ 0 & \text{if neutral or incidental} \end{cases}$$

This is derived from citation intent classification *via* NLP techniques. The final citation contribution becomes:

$$C_i := \sum_{j=1}^n \chi_j \cdot \delta(t_j) \cdot A(d_j)$$

Definition 36 (citation entropy penalty σ_i): To prevent belief inflation through redundant citation clusters, define:

$$\sigma_i := - \sum_k p_k \log p_k$$

where p_k is the proportion of citations from venue or author cluster k . A low entropy implies citation redundancy; thus: $C_i^* := C_i \cdot (1 - \epsilon \cdot (1 - \sigma_i / \log K))$, with $\epsilon \in (0, 1)$ and K the number of clusters.

Implementation note: BEWA integrates citation metadata from Crossref, Dimensions, and Semantic Scholar APIs. Citation intent is classified using transformer-based models fine-tuned on datasets like SciCite and ACL-ARC. Author clusterings and venue normalisation address bias from prolific low-impact publication venues.

Axiom 26 (non-linearity of citation weight): The influence of citations is not additive but modulated by both credibility and decay; hence

$$\frac{dC_i}{dn} \neq \text{constant}.$$

In conclusion, citations are not raw votes but weighted endorsements. BEWA's formalisation ensures they are evaluated dynamically, critically, and contextually ensuring that citation inflation does not substitute for replicable epistemic merit.

Replication scoring and semantic equivalence

In scientific inquiry, replication serves as a keystone of epistemic reliability. A claim that has been successfully replicated across independent studies accrues higher confidence than one supported by a single observation. However, direct replication is rare; more frequently, validation occurs via semantically equivalent or derivatively confirmatory studies. Hence, BEWA implements a dual-layer scoring function that quantifies both replication frequency and semantic proximity.

Definition 37 (replication instance ρ_{ij}): Let ϕ_i be a primary claim and d_j a study. Then $\rho_{ij} = 1$ if d_j contains a successful replication of ϕ_i under predefined methodological fidelity and independence constraints. Otherwise, $\rho_{ij} = 0$.

Definition 38 (replication score \mathcal{R}_i):

$$\mathcal{R}_i := \sum_{j=1}^m \rho_{ij} \cdot \omega_j,$$

where m is the number of replication attempts recorded; $\omega_j \in (0, 1)$ is the credibility weight of the replicating study, derived from §5–6 metrics.

Axiom 27 (diminishing marginal replication): After a sufficient number of high-quality replications, the marginal gain in \mathcal{R}_i should diminish. Thus, define:

$$\mathcal{R}_i^* := \log(1 + \mathcal{R}_i).$$

Definition 39 (semantic equivalence function ς): Given two claims ϕ_i and ϕ_k , define semantic equivalence as: $\varsigma(\phi_i, \phi_k) := \text{sim}(\text{vec}(\phi_i), \text{vec}(\phi_k))$,

where $\text{vec}(\cdot)$ denotes a structured embedding (e.g., SPECTER, Sentence-BERT) and sim a cosine similarity function. A threshold τ is set such that $\varsigma > \tau$ implies candidate equivalence.

Definition 40 (replication *via* equivalence): Let $S_i = \{\phi_k : \varsigma(\phi_i, \phi_k) > \tau\}$ be the set of semantically equivalent claims. Then the extended replication score is:

$$\tilde{\mathcal{R}}_i := \mathcal{R}_i^* + \eta \cdot \sum_{\phi_k \in S_i} \alpha_k \cdot \mathcal{R}_k^*,$$

where $\eta \in (0, 1)$ is a semantic discount factor, and $\alpha_k = \varsigma(\phi_i, \phi_k)$.

Proposition 17 (semantic replication transitivity bound): For three claims ϕ_i, ϕ_j, ϕ_k such that $\varsigma(\phi_i, \phi_j), \varsigma(\phi_j, \phi_k) > \tau$, the triangle inequality in embedding space yields: $\varsigma(\phi_i, \phi_k) \geq \varsigma(\phi_i, \phi_j) + \varsigma(\phi_j, \phi_k) - 1$.

Proof. Follows from cosine similarity properties under triangle inequality in unit norm vector space.

Implementation note: BEWA incorporates NLP-based equivalence scoring with fine-tuned SPECTER embeddings and indexes replication records from curated sources such as the Center for Open Science, ReplicationWiki, and the Cochrane Library. Methodological fidelity is validated using structured experiment metadata and controlled vocabulary (e.g., MeSH, CRediT taxonomy).

Axiom 28 (independence of replication): Only replications with no overlapping authors, funding sources, or institutions are included in the high-confidence tier of \mathcal{R}_i .

Definition 41 (contradictory replication penalty): Let $\rho_{ij} = -1$ denote a failed replication. Then:

$$\mathcal{R}_i := \sum_j \rho_{ij} \cdot \omega_j, \quad \rho_{ij} \in \{-1, 0, 1\}.$$

BEWA's mechanism captures not merely quantity but quality, novelty, and semantic alignment of replication constructing a resilient, logically grounded foundation for belief revision. The weight of a claim becomes a composite function of both repeated validation and structural equivalence to a corpus of convergent assertions.

Contradiction mapping and network response

The formal integrity of an epistemic architecture such as BEWA depends critically on its capacity to map, diagnose, and structurally respond to contradictions in its claim network. Contradictions arise when two or more propositions are logically or empirically incompatible, yet simultaneously persist within the network under nonnegligible posterior probabilities. This subsection formalises contradiction detection through graph-theoretic structures and probabilistic thresholds, and outlines the system's response protocols for epistemic rebalancing.

Definition 42 (contradictory claim pair): Let $\phi_i, \phi_j \in \Phi$ be distinct claims. We define the contradiction predicate $\chi(\phi_i, \phi_j) = 1$ if and only if: $\text{Entails}(\phi_i \rightarrow \neg\phi_j) \vee \text{Entails}(\phi_j \rightarrow \neg\phi_i)$. Entailment is evaluated *via* a formal logical entailment engine or, in empirical domains, high-confidence semantic contradiction scores: $\chi_{\text{sem}}(\phi_i, \phi_j) = \text{antisim}(\vec{\phi}_i, \vec{\phi}_j)$, where antisim is a model trained to capture negation and contradiction, such as those developed in the ANLI corpus or NLI benchmarks.

Definition 43 (contradiction graph \mathcal{C}): Construct a contradiction graph $\mathcal{C} = (V, E)$ where: $V = \{\phi_i \in \Phi\}$; $E = \{(\phi_i, \phi_j) \mid \chi(\phi_i, \phi_j) = 1\}$.

Axiom 29 (no persistent high-probability contradiction): There exists no (ϕ_i, ϕ_j) such that $\chi(\phi_i, \phi_j) = 1$ and $P(\phi_i), P(\phi_j) > \theta_c$, where θ_c is a contradiction coherence threshold.

Protocol 11 (network response to contradictions): When $P(\phi_i), P(\phi_j) > \theta_c$ and $\chi(\phi_i, \phi_j) = 1$:

- Evaluate global evidence distributions: $\mathcal{S}\phi$ and $\mathcal{S}\phi$.
- Compute resolution function:

$$\Delta_{ij} := \log \left(\frac{\sum_{e \in \mathcal{E}_{\phi_i}} \mathcal{L}(e | \phi_i)}{\sum_{e \in \mathcal{E}_{\phi_j}} \mathcal{L}(e | \phi_j)} \right)$$

- Adjust beliefs: $P'(\phi_i) := P(\phi_i) \cdot \sigma(\Delta_{ij})$, $P'(\phi_j) := P(\phi_j) \cdot \sigma(-\Delta_{ij})$, where σ is the logistic function.
- Flag both claims with instability tags and reduce propagation radius.

Proposition 18 (cycle detection and contradiction clustering): If C contains a cycle of size $n > 2$, then at least one subcluster is inconsistent under Axiom 29 and must be quarantined.

Proof. By contradiction: Assume all nodes in a cycle have posterior $> \theta_c$ and mutually contradict. This violates Axiom 29. Therefore, network coherence mandates demotion or reappraisal of one or more claims.

Definition 44 (epistemic quarantine set Q): Let $Q \subset \Phi$ be the minimal subset of C for which:

$$\sum_{\phi_i \in Q} P(\phi_i) \leq \min_{\text{partition of } C} \left(\sum_{\phi_i \in S} P(\phi_i) \right)$$

subject to $\chi(\phi_i, \phi_j) = 1$ for all $\phi_i, \phi_j \in Q$.

Implementation note: Contradiction detection is continuous and hierarchical. Local contradictions trigger lightweight Bayesian rebalancing, while global topological instabilities in C invoke structural pruning and epistemic quarantine. Contradiction edges are stored in a conflict adjacency matrix, versioned and cryptographically anchored to support full auditability (see §12).

BEWA's contradiction management maintains belief coherence without resorting to naive conflict resolution. Claims are demoted or discounted in proportion to their evidentiary support relative to conflicting counterparts, enforcing rational consistency throughout the system's propositional lattice.

Cross-Claim Belief Networks

This section describes the architecture and function of the cross-claim belief network that undergirds the epistemic reasoning of BEWA. Scientific knowledge is not atomised; claims are interdependent, nested within broader conceptual lineages, and frequently contingent on the veracity of adjacent assertions. BEWA encodes these interdependencies within a dynamic belief graph, wherein each node represents a structured propositional claim and edges encode semantic similarity, logical entailment, or evidentiary correlation. This allows the system to perform not only claim-specific inference, but also to propagate belief states across domains, responding in a principled manner to reinforcement or disruption within the wider network. Such propagation is not indiscriminate: it is

governed by weighted edge relationships, epistemic thresholds, and domain-informed constraints to prevent spurious or disproportionate influence.

The networked design permits both local sensitivity and global coherence. A robustly supported claim may stabilise a fragile neighbouring cluster, while a replicated contradiction in one domain may reverberate through others, lowering confidence in structurally or semantically linked propositions. At the same time, the architecture is designed to tolerate ambiguity and local conflict without catastrophic failure. BEWA's graph model incorporates mechanisms for instability detection, epistemic damping, and edge decay, allowing it to gracefully manage evolving scientific landscapes. The belief network is thus neither rigid nor anarchic: It is an adaptive inferential structure capable of encoding nuance, resolving tension, and integrating new information with epistemic discipline. The subsections to follow detail the methods by which claims are linked, how belief values propagate, and how conflicts are identified and managed across clusters.

Semantic and logical linkage of claims

Within BEWA, claims are not evaluated in isolation but are situated within a rich web of inferential dependencies. These relationships semantic, logical, and evidential form the structural backbone of the belief graph, enabling propagation, inference, contradiction resolution, and coherence maintenance. This subsection formalises the mapping of such linkages and defines the criteria under which claims are unified into coherent epistemic structures.

Definition 51 (claim linkage function λ): Let $\phi_i, \phi_j \in \Phi$ be distinct claims. The linkage function $\lambda : \Phi \times \Phi \rightarrow \{0, 1\}^3$ decomposes as: $\lambda(\phi_i, \phi_j) := (\lambda_{\text{sem}}, \lambda_{\text{log}}, \lambda_{\text{evd}})$, where: $\lambda_{\text{sem}} = 1$ iff the semantic distance $d_{\text{sem}}(\vec{\phi}_i, \vec{\phi}_j) < \epsilon_s$; $\lambda_{\text{log}} = 1$ iff $\phi_i \vdash \phi_j$ or $\phi_j \vdash \phi_i$ in formal deductive logic; $\lambda_{\text{evd}} = 1$ iff $\exists e$ such that e supports both ϕ_i and ϕ_j with high likelihood.

Axiom 31 (triangulated linkage validity): A composite claim network must maintain closure under transitive semantic and logical relations: $\lambda_{\text{sem}}(\phi_i, \phi_j) = \lambda_{\text{sem}}(\phi_i, \phi_k) = 1 \Rightarrow \lambda_{\text{sem}}(\phi_j, \phi_k) = 1$; $\phi_i \vdash \phi_j, \phi_j \vdash \phi_k \Rightarrow \phi_i \vdash \phi_k$.

Definition 52 (linkage graph L): Construct the semantic-logical-evidential linkage graph $L = (V, E)$ where: $V := \{\phi_i \in \Phi\}$, $E := \{(\phi_i, \phi_j) \mid \lambda(\phi_i, \phi_j) \neq (0, 0, 0)\}$.

Edges in L are labelled with the type(s) of linkage present and weighted according to: $w(\phi_i, \phi_j) := \alpha_s \cdot \text{sim}(\vec{\phi}_i, \vec{\phi}_j) + \alpha_l \cdot \mathbb{I}_{\phi_i \vdash \phi_j} + \alpha_e \cdot \text{shared evidence}(\phi_i, \phi_j)$, with $\alpha_s + \alpha_l + \alpha_e = 1$ and domain-specific calibration.

Proposition 20 (semantic-evidential coherence): If $\lambda_{\text{sem}} = 1$ and $\lambda_{\text{evd}} = 1$ but $\lambda_{\text{log}} = 0$, then the system infers potential

unformalised inference. A soft-logic implication is queued for automated induction.

Definition 53 (claim cluster C_k): A claim cluster is a maximal subgraph of L with dense interlinkage: $\forall \phi_i, \phi_j \in C_k$, $w(\phi_i, \phi_j) > \tau$, for a coherence threshold τ .

Implementation note: Semantic embeddings use domain-specific transformers trained on curated corpora (e.g., SPECTER, SciBERT for scientific domains), while logical entailments are computed *via* automated theorem provers or higher-order logic systems.

Evidential overlap is computed through co-citation matrices and joint support probabilities.

Linkage graphs are updated incrementally upon ingestion of new claims or refinement of existing nodes. The graph L serves as the substrate for belief propagation, contradiction alerting, and claim contextualisation.

This structure enables BEWA to go beyond flat, isolated assertions transforming its knowledge base into an epistemic manifold of interconnected, evolving propositions, recursively interlinked by logic, meaning, and evidence.

Graph structures and belief propagation

To enable rigorous updating, evaluation, and querying of interconnected epistemic content, BEWA employs a formal belief graph over structured claims. This graph encodes both propositional assertions and their inferential, semantic, and evidentiary relationships, facilitating structured belief propagation *via* a generalised Bayesian network architecture augmented with non-monotonic belief revision rules.

Definition 54 (belief graph G): Let $G := (V, E, \pi)$ be a directed graph where: $V := \{\phi_i \in \Phi\}$ is the set of structured propositional claims; $E := \{(\phi_i, \phi_j, t_{ij})\}$ encodes directed influences (e.g., inferential, evidential) labelled with type $t_{ij} \in \mathcal{T}$; $\pi : V \rightarrow (0, 1)$ assigns each node a marginal belief value $\pi(\phi_i) = P(\phi_i \mid \mathcal{E})$ conditional on the current epistemic state \mathcal{E} .

Axiom 32 (typed edge semantics): Edges in G must reflect the dominant source of dependency: $t_{ij} \in \{\text{deductive, evidential, semantic, contrapositive, replicative}\}$.

Each t_{ij} triggers a distinct propagation operator $P_{t_{ij}}$ defined in §7.3.

Definition 55 (local belief propagation): For node ϕ_i with parents $\{\phi_j\}$, let: $\pi(\phi_i) := F\phi_j(\{\pi(\phi_j)\}, \{w_{ij}\})$, where F is a weighted aggregator function (typically log-linear or Noisy-OR) determined by edge types and propagation policies.

Proposition 21 (stability under sparse updates): Let G be a belief graph with sparse, bounded-degree topology. Then under

localised updates $\pi'(\phi_k) = \pi(\phi_k) + \delta$ for a single node ϕ_k , the update propagation is bounded in depth by the edge-type decay parameters $\{\gamma_{t_{ij}}\}$:

$$\forall \phi_m \notin \mathcal{N}_k^{(d)}, \quad |\pi'(\phi_m) - \pi(\phi_m)| < \epsilon, \quad \text{where } d \geq \frac{\log(\epsilon/\delta)}{\log(\min \gamma)}.$$

Proof. See Pearl, et al., for convergence properties in belief networks with decaying influence factors [12].

Definition 56 (global belief fixpoint): BEWA maintains a consistent belief state by iterative updates until:

$$\forall \phi_i \in V, \quad \left| \pi^{(t+1)}(\phi_i) - \pi^{(t)}(\phi_i) \right| < \delta,$$

where δ is a convergence threshold. The system uses asynchronous propagation, prioritising high-impact nodes (e.g., highly cited or recently updated claims).

Axiom 33 (non-monotonic revision): When high-weight counterevidence enters the graph (cf. §4.3), previously stable beliefs $\pi(\phi_i)$ may be revised downward. This violates monotonicity and requires a belief revision operator R that preserves network coherence: $\pi' := R(\pi, \phi_{\text{new}}, \pi(\phi_{\text{new}}))$.

Implementation note: Belief propagation is implemented atop a high-performance probabilistic graph database (e.g., DGraph, TigerGraph) with runtime support for incremental updates, batch re-evaluation, and fast approximate querying. Differentiable propagation layers are also available for neural fine-tuning in high-noise domains.

By mapping structured claims onto a belief graph, BEWA transforms raw assertion data into a dynamic, self-revising network of truth valuations. These valuations evolve over time through evidence ingestion, retraction, contradiction, and human input mirroring the dynamic structure of knowledge in high-integrity epistemic systems.

Handling conflicts and cluster instability

In any epistemic framework that aggregates heterogeneous sources, contradictions are not an anomaly but a structural inevitability. The BEWA system explicitly incorporates conflict detection and instability analysis within the belief network to both surface unresolved disputes and prevent epistemic contagion *i.e.*, the unbounded propagation of uncertainty or error through connected claims.

Definition 57 (conflict edge χ): Given claims $\phi_i, \phi_j \in \Phi$, a directed conflict edge $\chi(\phi_i, \phi_j)$ is instantiated if: Contradict $(\phi_i, \phi_j) = 1$, and $\pi(\phi_i), \pi(\phi_j) > \delta_c$, where δ_c is the threshold above which conflicting claims are considered epistemically relevant.

Axiom 34 (conflict non-coexistence constraint): For any claim pair (ϕ_i, ϕ_j) with a mutual contradiction edge: $\pi(\phi_i) +$

$\pi(\phi_i) \leq 1 + \epsilon$, with $\epsilon \rightarrow 0$ in the absence of new ambiguity-resolving evidence.

Definition 58 (epistemic cluster C): An epistemic cluster is a maximal Strongly Connected Component (SCC) of the belief graph G wherein: $\forall \phi_i, \phi_j \in C, \exists$ a directed path $\phi_i \rightarrow \phi_j$.

Definition 59 (instability score $\iota(C)$): Let C be an epistemic cluster. Its instability is:

$$\iota(C) := \frac{\sum_{\chi(\phi_i, \phi_j) \in C} \pi(\phi_i) \cdot \pi(\phi_j)}{|C|^2},$$

reflecting the density and severity of contradiction among highly believed claims.

Proposition 22 (instability containment bound): Let C_k and C_ℓ be distinct clusters with $\iota(C_k) > \tau$. Then $\forall \phi_i \in C_k, \phi_j \in C_\ell$, the system enforces attenuation:

$$\left. \frac{d\pi(\phi_j)}{dt} \right|_{\text{link } \phi_i \rightarrow \phi_j} \leq \alpha \cdot (1 - \iota(C_k)),$$

where α is a normal propagation coefficient.

Axiom 35 (conflict partitioning protocol): When $\iota(C) > \theta$, the system forks C into subclusters $\{C_1, C_2, \dots\}$ via minimum-cut algorithms that minimise inter-cluster contradiction flow:

$$\min_{\text{partition}} \sum_{\substack{\phi_i \in C_a \\ \phi_j \in C_b \\ a \neq b}} \pi(\phi_i) \cdot \pi(\phi_j) \cdot \neg \chi(\phi_i, \phi_j) = 1.$$

Implementation Note: Conflict edges are derived from formal inconsistency detection (via propositional and predicate logic), semantic negation models, and contradiction mining using NLP (models like DeBERTa for natural contradiction detection). Instability scores are recomputed after each batch update, and clustering is performed using approximate spectral methods for scalability.

This mechanism ensures that the BEWA system does not just assimilate information passively, but actively surveils the topology of belief for signs of epistemic fracture. High-instability clusters trigger alerts and conditional discounting, enabling the system to contain the epistemic contagion until new evidence preferably from independent replication resolves the underlying tension.

Truth Utility and System Optimisation

This section introduces the metarational layer of BEWA: The evaluation of claims not solely by probabilistic confidence, but by their contribution to the epistemic aim of truth-promotion. While Bayesian posterior probabilities quantify the likelihood of a claim being correct, they do not discriminate between claims that are trivial, inconsequential, or epistemically sterile. BEWA addresses this limitation by introducing a secondary axis of evaluation the truth utility function which modulates the visibility, prioritisation, and application of claims based on their potential to enhance scientific understanding, inform future inquiry, or rectify systemic error. This is not a measure of popularity or impact in the sociological sense, but of epistemic consequence: a calibrated utility score that integrates replication fidelity, methodological depth, inferential reach, and systemic relevance.

Optimisation within BEWA is governed by this layered assessment of epistemic value. High-certainty but low-utility claims are not suppressed, but they are deemphasised in recommendations, while claims with high truth-promotion potential even if currently epistemically marginal are flagged for further scrutiny and resource allocation. This truth utility paradigm enables the system to function not merely as a passive aggregator of belief, but as an active epistemic agent, capable of prioritising lines of inquiry, identifying underexamined anomalies, and suggesting evidentiary targets for the scientific community. The system also accounts for risk: epistemic weighting incorporates not just certainty, but the cost of error in downstream reasoning. The following subsections define how utility scores are constructed, how epistemic risk is embedded in weighting protocols, and how these factors drive decision-making at the application level.

Truth promotion score construction

To operationalise the epistemic objective of truth-conduciveness within the Bayesian Epistemic Weighted Architecture (BEWA), we introduce the Truth Promotion Score (TPS), a formal scalar quantity designed to evaluate the net effect of a claim, author, or domain on the system's ability to approximate true propositions. TPS is not merely a re-expression of posterior belief; rather, it measures the expected long-run contribution to truth discovery, subject to epistemic dynamics, counterfactual perturbation, and information flows.

Definition 60 (truth promotion score τ): Let $\phi \in \Phi$ be a structured claim. Then the Truth Promotion Score $\tau(\phi)$ is defined as: $\tau(\phi) := E(\Delta_{\text{TP}}(\psi) \mid \text{Inclusion of } \phi)$, where ψ ranges over all influenced downstream claims, $\pi(\cdot)$ is the system's belief

function, and Δ_T denotes the marginal contribution towards true claims $T \subset \Phi$ under the truth- model defined in §1.

Axiom 36 (causal relevance constraint): A claim ϕ has non-zero $\tau(\phi)$ only if its causal graph descendants ψ intersect with claims whose verified truth status is established via: Replicated experimental outcomes; axiomatic derivations; authoritative peer-consensus convergence; or semantically equivalent high-truth claims $\phi' \in \Phi$ with $\pi(\phi') \approx 1$.

Definition 61 (weighted TPS for composite nodes): For composite structures $X \in \{\text{author, domain, corpus}\}$, let:

$$\tau(X) := \sum_{\phi_i \in X} \tau(\phi_i) \cdot w_i,$$

where w_i encodes context-adjusted weights (e.g., citation impact, belief volatility, down- stream propagation depth).

Proposition 23 (TPS decomposition): The TPS admits an additive decomposition:

$$\tau(\phi) = \sum_{\psi \in \text{Desc}(\phi)} \underbrace{\kappa(\phi, \psi)}_{\text{causal influence}} \cdot \underbrace{\pi(\psi) \cdot \mathbb{K}_{\psi \in \mathbb{T}}}_{\text{downstream truth fide}}$$

where κ is a directional influence coefficient derived from the belief graph's adjacency tensor and belief propagation kernel.

Definition 62 (anti-truth penalty): A claim ϕ that consistently promotes falsehood reduces system epistemic integrity. Define: $\tau^-(\phi) := -E[\Delta_{\text{Fit}}(\psi)]$, with $F := \Phi \setminus T$, such that $\tau(\phi) = \tau^+(\phi) + \tau^-(\phi)$.

Axiom 37 (TPS normalisation for belief-based querying): All truth promotion scores are min-max normalised within each belief update cycle:

$$\tau'(\phi) = \frac{\tau(\phi) - \min_{\phi} \tau(\phi)}{\max_{\phi} \tau(\phi) - \min_{\phi} \tau(\phi)},$$

to ensure stable comparability across domains and preserve bounded impact in probabilistic inference.

Implementation note: TPS is used to prioritise which claims are (a) surfaced to users in summary or audit views; (b) subjected to targeted re-evaluation or replication; and (c) given epistemic preference in domain-specific queries. It also serves as the foundational metric in risk-aware reasoning (see §9.2) and time-ordered epistemic triage (§10).

By centring its computational evaluation on truth promotion, BEWA reorients the logic of ranking and reasoning

away from popularity or novelty and toward epistemic utility a formalisation of the normative goals espoused by Bayesian and Popperian traditions alike.

Risk-aware epistemic weighting

In constructing a system whose outputs influence epistemic reasoning and potentially policy or scientific action, it is insufficient to treat all uncertainty as epistemically symmetric. The Bayesian Epistemic Weighted Architecture (BEWA) integrates a formal mechanism of risk-aware weighting, adjusting posterior belief propagation not merely by epistemic probability but by the potential epistemic loss associated with errors in belief attribution. This parallels and extends decision-theoretic Bayesian frameworks, particularly in high-stakes inferential settings.

Definition 63 (epistemic loss function $L(\phi)$): Let $\phi \in \Phi$ be a claim, and define a loss function $L : \Phi \times \{0, 1\} \rightarrow \mathbb{R}_{\geq 0}$ such that:

$$\mathcal{L}(\phi, t) = \begin{cases} \lambda_{\text{FN}}(\phi) & \text{if } \pi(\phi) < \theta \wedge t = 1, \\ \lambda_{\text{FP}}(\phi) & \text{if } \pi(\phi) \geq \theta \wedge t = 0, \\ 0 & \text{otherwise,} \end{cases}$$

where t is the truth state of ϕ , $\pi(\phi)$ the belief assignment, and θ the classification threshold.

Definition 64 (weighted belief utility $U(\phi)$): $U(\phi) := \pi(\phi) \cdot u_{\text{TP}}(\phi) - (1 - \pi(\phi)) \cdot \lambda_{\text{FP}}(\phi)$, for utility term u_{TP} conditioned on true belief propagation, and penalty λ_{FP} for false positives.

Axiom 38 (risk dominance bias): In contexts with asymmetric error costs (i.e., $\lambda_{\text{FP}} \gg \lambda_{\text{FN}}$), the system shall bias belief updating conservatively unless replication or corroboration reduces L to below a system-defined ϵ .

Definition 65 (domain risk profile ρ_d): Each domain $d \in D$ is assigned a contextual risk scalar $\rho_d \in (0, 1)$ derived from:

$$\rho_d := \frac{\sum_{\phi_i \in d} \lambda_{\text{FP}}(\phi_i) + \lambda_{\text{FN}}(\phi_i)}{|\phi_i \in d|}.$$

Proposition 24 (risk-adjusted belief propagation): For any edge $(\phi_i \rightarrow \phi_j)$ in G , belief propagation is governed by:

$$\frac{d\pi(\phi_j)}{dt} \propto \pi(\phi_i) \cdot (1 - \rho_{d_j}) \cdot U(\phi_i),$$

where d_j is the domain of ϕ_j .

Axiom 39 (minimum risk integrity constraint): For all $\phi \in \Phi$ with $\rho_d(\phi) > \gamma$ (e.g., medical, engineering, safety domains),

any claim admitted must be: From an authoritative source, and independently replicated or derived from axioms.

Implementation note: The risk-aware structure is embedded into both forward belief propagation and posterior adjustments after contradiction analysis (S6). It enables the system to weight beliefs not solely on epistemic strength but on the potential harm of incorrect belief, aligning with real-world use cases where epistemic robustness must account for asymmetric downstream costs.

Risk-aware epistemic weighting thus elevates BEWA beyond naive belief ranking systems, binding probabilistic reasoning to a utility-sensitive infrastructure that aligns with the pragmatic imperatives of scientific integrity and responsible automation.

Application-level prioritisation strategies

The value of an epistemic inference system such as BEWA emerges most prominently when its operations are purpose-aligned: That is, when the abstract belief network is filtered, ordered, and evaluated in terms of the decision-utility of downstream applications. This section formalises how BEWA distinguishes, weights, and channels information depending on the application-layer context in which it is deployed be it clinical inference, financial policy, safety-critical system validation, or scientific knowledge curation.

Definition 66 (application class A): Let $A = \{A_1, A_2, \dots, A_k\}$ denote the set of distinct application domains. Each A_i is characterised by a tuple: $A_i := (\Gamma_i, \Lambda_i, O_i)$, where Γ_i defines input claim relevance criteria, Λ_i defines acceptable epistemic risk thresholds, and O_i encodes output functionals (e.g., ranking, filtering, recommendation).

Axiom 40 (contextual utility projection): For each claim $\phi \in \Phi$ and application A_i , define: $\pi_{A_i}(\phi) := \pi(\phi) \cdot U_{A_i}(\phi)$, where $U_{A_i}(\phi)$ is a domain-specific utility transformation as defined by Λ_i and O_i .

Definition 67 (domain-prioritised subgraph G_{A_i}): Let $G = (\Phi, E)$ be the global belief graph. Then the application-filtered subgraph is: $G_{A_i} := (\Phi_{A_i}, E_{A_i})$, with $\Phi_{A_i} := \{\phi \in \Phi \mid \Gamma_i(\phi) = 1\}$.

Proposition 25 (maximal belief flow selection): Let $\phi_{\max}^{A_i} := \arg \max_{\phi \in \Phi_{A_i}} \pi_{A_i}(\phi)$. Then the system exposes:

$$\text{Top-}K_{A_i} := K\text{-argmax}_{\phi \in \Phi_{A_i}} \pi_{A_i}(\phi),$$

with optional re-ranking under O_i such as trust-aware PageRank, belief-aware influence spread, or task-specific thresholds.

Definition 68 (query-adaptive prioritisation function Ψ): Let Q be a structured query from an application. Then: $\Psi_Q(\phi) := f(\pi(\phi), \text{sim}(\phi, Q), \rho_d(\phi), U_{A_i}(\phi))$, where $\text{sim}(\phi, Q)$ measures

semantic distance and f is a monotonic ranking function with attenuation on risky but low-utility claims.

Axiom 41 (saturation control and freshness bias): For applications with high volatility or innovation rate (e.g., real-time systems, scientific frontier models), prioritisation shall favour claims ϕ such that: $\text{Age}(\phi) < \Delta_t$, and replication status(ϕ) = pending, with optional quarantine if $\pi(\phi)$ exhibits extreme volatility.

Implementation note: Application-level prioritisation is handled post-belief-update and pre-output stage. Each domain invokes Ψ over its subgraph, applies domain-specific constraints from Λ_i , and resolves a sorted response set through O_i (e.g., recommendation, diagnostic inference, triage). In critical systems, safety predicates additionally filter ϕ for compliance with predefined constraints, such as ISO-26262 or FDA standards.

This framework ensures that BEWA's epistemic outputs are not only consistent and truthful within the graph but dynamically useful within contextual deployments, harmonising epistemic strength with the practical realities of truth utility across disciplines.

Temporal Dynamics and Critical Delay Protocol

This section delineates the temporal logic that governs the evolution of epistemic confidence within the BEWA framework. Scientific knowledge is not static; the reliability of claims fluctuates over time based on patterns of usage, replication, contradiction, and neglect. BEWA internalises this diachronic instability through a set of principled temporal mechanisms that modulate belief in accordance with the claim's evidentiary lifespan. Key to this architecture is the understanding that claims, however initially robust, may lose epistemic standing as their replication stales, as methodologies evolve, or as new evidence supersedes old paradigms. The system integrates decay protocols that algorithmically diminish belief in claims that fail to remain epistemically active or substantiated. This prevents epistemic inertia the retention of unverified assertions due to their historical prominence and reinforces a living standard of proof.

To counterbalance decay, BEWA implements a critical delay protocol that governs the assimilation of new claims into the belief network. Early-stage assertions are treated with caution, regardless of venue or authorship, and are assigned epistemic probation until they accrue evidence commensurate with their proposed influence. This guards the system against transient trends, premature consensus, or the diffusion of findings that

have not undergone adequate methodological vetting. At the same time, replication events function as belief resets, halting or reversing decay trajectories and reactivating aged claims with renewed force. The system thus exhibits both epistemic scepticism and responsiveness able to resist informational volatility while dynamically re-evaluating beliefs in light of cumulative and time-sensitive evidence. The subsections to follow formalise the decay algorithms, describe the mechanics of replicative reinforcement, and specify the temporal constraints imposed on the admission and weighting of nascent scientific claims.

Decay protocols for aged or isolated claims

To sustain epistemic relevance and prevent the long-term inflation of unsupported or obsolete information, the BEWA system introduces mathematically defined decay protocols that attenuate belief in claims based on age, isolation, and lack of supporting interaction. Such decay is not uniform but governed by temporal, structural, and evidence-based metrics, ensuring that persistent claims remain weighted in proportion to their sustained relevance and interaction.

Definition 70 (claim age $\Delta_t(\phi)$): Let t_0 be the time of first ingestion of claim ϕ . Then: $\Delta_t(\phi) := t_{\text{current}} - t_0$.

Definition 71 (isolation score $\iota(\phi)$): Let $k_{\text{in}}(\phi)$ and $k_{\text{out}}(\phi)$ be the in-degree and out-degree of ϕ in the belief network G . Then:

$$\iota(\phi) := \frac{1}{1 + k_{\text{in}}(\phi) + k_{\text{out}}(\phi)},$$

reflecting topological disconnection.

Axiom 43 (time-based decay law): Each belief score $\pi(\phi)$ is updated continuously via a decay differential equation:

$$\frac{d\pi(\phi)}{dt} = -\lambda_t(\phi) \cdot \pi(\phi),$$

with $\lambda_t(\phi)$ defined as: $\lambda_t(\phi) = \alpha_t \cdot \log(1 + \Delta_t(\phi))$, where α_t is the system decay constant calibrated per domain.

Definition 72 (compound decay rate $\Lambda(\phi)$): $\Lambda(\phi) := \lambda_t(\phi) + \lambda_i(\phi)$, where, $\lambda_i(\phi) = \beta \cdot \iota(\phi)$, and β is a domain-specific attenuation constant penalising isolation.

Proposition 26 (belief half-life $T_{1/2}(\phi)$): Assuming $\Lambda(\phi)$ is constant, the belief score halves in:

$$T_{1/2}(\phi) = \frac{\ln(2)}{\Lambda(\phi)}.$$

Axiom 44 (decay immunity clause): Claims ϕ with verified replication count $r(\phi) \geq \rho_{\text{min}}$ and citation count $c(\phi) \geq \kappa$ are decay-immune: $\Lambda(\phi) \rightarrow 0$ iff $r(\phi) \geq \rho_{\text{min}} \wedge c(\phi) \geq \kappa$.

Definition 73 (decay checkpoint audit D_ϕ): Let $D_\phi := \{t_i\}_{i=1}^N$ denote a series of decay checkpoints for claim ϕ . At each t_i :

$$\text{if } \left. \frac{d\pi(\phi)}{dt} \right|_{t_i} < -\epsilon \quad \text{and } \nexists \text{ new citation or support,} \quad \pi(\phi) \leftarrow \pi(\phi) \cdot \delta,$$

where $0 < \delta < 1$ is a discrete decay multiplier.

Implementation note: Temporal decay is implemented as a background process operating on a delta queue. Isolation is recomputed via dynamic degree tracking and page-rank metrics. Claims with only a single inbound edge from low-weight sources are flagged for manual or peer-audit. These decay dynamics ensure that claims only retain epistemic weight proportional to their integration and relevance in the knowledge substrate.

The decay framework thus institutionalises the notion that belief is not merely acquired—it must be maintained through ongoing relevance, validation, and use. In epistemic systems designed for truth promotion, neglect is synonymous with obsolescence.

Replicative reset mechanisms

In order to prevent epistemic inertia wherein outdated or contextually misaligned claims retain undue influence the BEWA framework incorporates a class of mechanisms termed replicative resets. These mechanisms allow for the systematic re-evaluation and belief recalibration of claims following independent replication events, especially when these events introduce novel experimental conditions, methodological advancements, or domain shifts. The replicative reset ensures that belief is not only responsive to new evidence but also dynamically restructured based on the depth and robustness of repeated validation.

Definition 74 (replication event $R_k(\phi)$): A replication event $R_k(\phi)$ is a semantically equivalent experimental or analytical claim confirming ϕ under modified or novel boundary conditions Θ_k . That is: $R_k(\phi) := \phi'$ such that $\text{SemEq}(\phi, \phi') = 1$, and $\Theta_k \not\subseteq \Theta(\phi)$.

Axiom 45 (reset threshold criterion): Let ϕ be a claim with current belief score $\pi(\phi)$. If $|R(\phi)| \geq \gamma_r$, and:

$$\forall \phi'_i \in \mathcal{R}(\phi) : \pi(\phi'_i) > \delta_r,$$

then a replicative reset is triggered, updating $\pi(\phi)$ as follows:

$$\pi(\phi) \leftarrow \frac{1}{Z} \sum_{i=1}^{|\mathcal{R}(\phi)|} w_i \cdot \pi(\phi'_i),$$

where w_i are confidence weights based on the source credibility and context of ϕ'_i , and Z is a normalising constant.

Definition 75 (reset modifier $\mu(\phi)$): If ϕ has previously decayed via standard time-based or isolation decay, the reset factor $\mu(\phi)$ defines the proportion of belief reinstated: $\mu(\phi) := \min(1, \log(1 + |\mathcal{R}(\phi)|) \cdot \eta)$, with η as a calibration constant set per epistemic domain.

Proposition 27 (decay reversal bound): Let $\pi_d(\phi)$ be the belief after decay and $\pi_r(\phi)$ after reset. Then: $\pi_r(\phi) \leq \mu(\phi) \cdot \pi^*(\phi)$, where $\pi^*(\phi)$ is the pre-decay score stored in the claim's historical register.

Definition 76 (domain-specific reset filter FQ): Replications from incompatible domains (e.g., biology \leftrightarrow computer science) are discarded unless cross-domain transfer is validated. Formally:

$$\phi'_i \in \mathcal{R}(\phi) \Rightarrow \Omega(\phi'_i) = \Omega(\phi) \quad \text{or} \quad \text{CrossMap}(\Omega(\phi'_i), \Omega(\phi)) = 1.$$

Axiom 46 (reset immunity for terminally discredited claims): If ϕ has an accumulated contradiction weight $\kappa(\phi) > \kappa_{\max}$, it is flagged as terminally discredited. No reset shall occur: $\mathcal{R}(\phi) \neq \emptyset \Rightarrow \pi(\phi) = 0$ if $\kappa(\phi) > \kappa_{\max}$.

Implementation note: Replication equivalence is verified using semantic matching across claim ontologies and domain-specific syntactic cores, leveraging transformer-based embeddings with domain adaptation layers (e.g., SciBERT, BioBERT). The reset process is executed asynchronously in the update queue, triggering belief cascade updates in downstream dependent claims.

By providing a structured yet adaptive mechanism for resetting epistemic weight in response to robust replication, the BEWA system mirrors the scientific imperative of falsifiability and cumulative verification, thereby maintaining alignment with the principles of progressive epistemic refinement.

Probationary periods for new claims

In the lifecycle of a scientific or factual assertion within the Bayesian Epistemic Weighting Architecture (BEWA), a newly introduced claim must not immediately inherit epistemic authority or be granted high belief priors. This subsection formalises the concept of a probationary period as a transitional epistemic state where new claims ϕ are initially treated with guarded scepticism until their stability, replicability, and citation potential are sufficiently demonstrated.

Definition 81 (probationary claim ϕ^\dagger): A claim ϕ enters probationary status upon first ingestion. It is denoted as ϕ^\dagger and is subject to restricted propagation and reduced network influence.

$$\phi^\dagger := \{\phi \mid t_0 = \tau(\phi), \quad \pi(\phi, t_0) = \pi_0^\dagger, \quad \pi_0^\dagger \ll \pi_{\min}\}$$

where π_0^\dagger is the designated low initial belief (e.g., $\pi_0^\dagger = 0.05$), and $\tau(\phi)$ denotes the claim's timestamp.

Axiom 48 (Probationary Isolation Constraint): During the interval $(t_0, t_0 + \Delta_\dagger)$, where Δ_\dagger is the probation duration, ϕ^\dagger is not permitted to influence any of the following: Belief propagation networks; dependent claim weightings, or citation-enhanced weight transference.

$$\text{Formally: } \forall t \in [t_0, t_0 + \Delta_\dagger], \quad \nexists \psi : \phi^\dagger \rightarrow \psi$$

Definition 82 (maturation function $M(\phi^\dagger, t)$): Let M be a non-decreasing function mapping observed external validations (citations, peer endorsements, replications) to belief score updates. For all $t > t_0$: $\pi(\phi, t) = \pi_0^\dagger + \lambda_c \cdot C(\phi, t) + \lambda_r \cdot R(\phi, t)$, where $C(\phi, t)$ and $R(\phi, t)$ are the cumulative citation and replication scores up to time t , and λ_c, λ_r are calibration weights empirically derived per domain.

Definition 83 (promotion criterion Π^\dagger): A claim ϕ^\dagger exits probation and becomes canonical (ϕ) if:

$$\mathcal{M}(\phi^\dagger, t) \geq \pi_{\min}, \quad \text{and} \quad \Delta_\dagger^\phi \leq \Delta_{\max}$$

where π_{\min} is the minimum threshold for epistemic impact and Δ_{\max} is the maximum probation duration permitted (after which the claim expires unless promoted).

Axiom 49 (decay upon probation failure): If ϕ^\dagger has not met the promotion criterion by Δ_{\max} , it undergoes exponential belief decay: $\pi(\phi, t) \leftarrow \pi_0^\dagger \cdot e^{-\lambda_d(t - \Delta_{\max})}, \forall t > t_0 + \Delta_{\max}$

where λ_d is a domain-specific decay constant. Claims failing probation are demoted to the archival register and disconnected from belief propagation chains.

This probationary structure preserves epistemic integrity by enforcing a cooling-off period for novel assertions, allowing belief to emerge proportionally to community uptake, replication stability, and authorial credibility.

Interfaces and Applications

This section articulates the outward-facing components of BEWA how its epistemic machinery is rendered accessible, interpretable, and actionable to end-users. The system is not designed to remain a closed inferential engine, but to serve as a transparent, intelligible platform for researchers, auditors, and

knowledge institutions. Its interface architecture reflects a dual imperative: first, to expose the reasoning process with precision and granularity, enabling full auditability of belief formation and claim evolution; second, to present this complexity without compromising usability or interpretative clarity. At every stage, the goal is to make epistemic justification navigable to allow users not only to know what BEWA believes, but to understand precisely why those beliefs are held and how they change over time.

The system supports a layered interaction model. Through structured query interfaces and programmatic APIs, external agents can interrogate belief states, trace evidentiary pathways, and simulate hypothetical updates. The user interface is designed to reveal the provenance of each claim, its temporal evolution, its position within the belief network, and the interplay of citations, contradictions, and replications that define its epistemic status. Beyond raw data access, BEWA offers visualisation modules that map the shifting topography of scientific belief across time and domain, enabling both high-level trend analysis and fine-grained epistemic inspection. The following subsections explore these capabilities in detail, covering the design of the query and audit interface, considerations in human-computer epistemic interaction, and the tools developed for visualising dynamic belief landscapes.

Query and audit API

The Query and Audit API (QAAPI) serves as the external interface layer through which users and systems interact with the Bayesian Epistemic Weighting Architecture (BEWA). It enables both epistemic interrogation and forensic traceability, bridging machine-readable knowledge representations with external agents seeking to evaluate, interrogate, or reproduce claims and their underlying justifications.

Definition 101 (query interface Q): Let $Q: \Sigma \rightarrow B$ be the mapping from well-formed structured queries Σ to belief-annotated outputs B such that: $Q(\sigma_i) = \langle \phi_i, \pi(\phi_i, t), \Gamma_i, \Delta_i \rangle, \forall \sigma_i \in \Sigma$, where ϕ_i is a matched claim, $\pi(\phi_i, t)$ is its current belief score at time t , Γ_i is the set of all epistemic justifications (citations, replications, metadata), and Δ_i is the set of contradictions or counter-evidence.

The semantics of σ_i support: Propositional logic-based queries (e.g., conjunctions, disjunctions); temporal filters (e.g., belief at time t'); authorial lineage traces (e.g., claims originating from author α); domain-scoped filtering (e.g., within neuroscience).

Definition 102 (audit interface A): The audit function $A: \phi \rightarrow \mathcal{T}$ maps any canonical or probationary claim ϕ to its historical trajectory \mathcal{T} , which includes:

$$\mathcal{T} = \{(t_k, \pi_k, E_k, M_k)\}_{k=0}^n$$

where each tuple denotes the belief state at time t_k , posterior π_k , contributing evidence E_k , and modifying events M_k (e.g., retractions, contradictory citations, revisions).

Axiom 59 (verifiability axiom): For every output from Q and A , the system must return: A cryptographic hash of the current belief graph state (anchor provenance); a path to all causal justifications leading to $\pi(\phi_i, t)$; signatures of any verified peer-reviewed replication or retraction events.

This guarantees that belief evolution is not opaque and that all justifications underlying current weights are auditable, immutable, and externally reconstructable.

Definition 103 (immutable query anchor HQ): Each response $Q(\sigma)$ is accompanied by:

$$\mathcal{H}_Q(\sigma, t) = \text{Hash}(\text{DAG}_t^{\phi_j} \cup \Gamma_j \cup \Delta_j)$$

where $\text{DAG}_t^{\phi_i}$ is the directed acyclic graph of dependent and influencing claims for ϕ_i at time t .

Together, the QAAPI ensures that users may perform both forward inference over the BEWA system and backward audit to challenge, verify, or dispute belief assignments. It enforces methodological transparency and underpins all mechanisms of public accountability and reproducibility.

User interface design considerations

While the Bayesian Epistemic Weighting Architecture (BEWA) is mathematically and structurally rigorous, its accessibility and epistemic transparency depend crucially on a User Interface (UI) that exposes complex inferential structures in a cognitively tractable form. The system must bridge deductive provenance chains, probabilistic evolution, and domain-specific claim networks without obfuscation, distortion, or interpretive bias.

Axiom 64 (cognitive parsimony axiom): The UI must minimise cognitive load while preserving formal accuracy. Let Π_t be the posterior distribution over all claims $\{\phi_i\}$ at time t , and let L be the user-visible instantiation of this distribution. Then:

$$\forall u \in \mathcal{U}, \quad \mathcal{C}(\mathcal{I}_u(\Pi_t)) < \theta$$

where \mathcal{U} is the user population, \mathcal{C} is the cognitive complexity function measured via empirical usability trials, and θ is the maximum tolerable complexity threshold.

Definition 104 (visual traceability layer V): A directed acyclic belief graph $G_t = (V, E)$ must be interactively explorable such that each node $v_i \in V$ (representing ϕ_i) displays: The belief

trajectory $\pi(\phi_i, t)$; sources of corroboration (Γ_i); contradictions (Δ_i); temporal provenance $\tau(\phi_i)$ and update events.

The user must be able to traverse causal and evidential pathways with $O(\log n)$ interaction depth for any ϕ_i in the graph, maintaining constant-time visibility of justification clusters.

Axiom 65 (truth-promoting layout axiom): The interface shall not encode salience purely via popularity, citation volume, or recency. Instead, layout heuristics L must prioritise: $L(\phi_i)$ a $U_i(\phi_i)$

where $L_i(\phi_i)$ is the truth promotion score as defined in section 9, ensuring that epistemically rigorous but under-cited claims are not visually buried.

Definition 105 (claim card component C_i): Each node ϕ_i is rendered via a modular interface component C_i that includes: Author and institution metadata; current posterior π with time-evolution sparkline; trust lineage: most influential corroborators and contradictors; epistemic role: Foundational, auxiliary, speculative, or deprecated.

Finally, the UI design must adhere to cryptographic anchoring of claim state snapshots, with visible indicators for state hash validation, consistency verification, and audit trail access. This ensures that what the user sees is not only comprehensible but also verifiable and immutable under tamper-sensitive conditions.

Visualisation of belief evolution

Belief evolution within the Bayesian Epistemic Weighting Architecture (BEWA) represents a dynamic topological transformation of knowledge, requiring visual encodings that preserve semantic fidelity, temporal coherence, and interpretive clarity. The representation of $\pi(\phi_i, t)$ across $t \in \mathbb{R}^+$ for each claim ϕ_i constitutes a non-Markovian process, influenced by latent structural couplings, propagative updates, and external contradiction events. The visualisation layer must expose these updates without distorting epistemic meaning or privileging transient statistical artefacts.

Definition 106 (belief evolution function $B_i(t)$): Let $B_i(t) : \mathbb{R}^+ \rightarrow (0, 1)$ denote the belief trajectory of claim ϕ_i . For all $t_k < t_{k+1}$, the update relation is given by: $B_i(t_{k+1}) = \eta \cdot B_i(t_k) + (1 - \eta) \cdot \Delta_k$.

where Δ_k is the net belief change due to evidence accrued at t_{k+1} , and $\eta \in (0, 1)$ is the epistemic inertia coefficient.

Definition 107 (epistemic flux diagram): A directed streamgraph $S = \{(\phi_i, B_i(t))\}$ visualises time-series belief trajectories as variable-width timelines. The vertical width at time t encodes $B_i(t)$, while colour denotes domain. Cross-claim

correlations are indicated via Bezier-curve arcs linking events of synchronous update or belief reversal. Such correlations derive from underlying contradiction networks Δ_i and citation couplings Γ_i .

Temporal visualisation must adopt a non-linear time axis $\tau : \mathbb{R}^+ \rightarrow \mathbb{R}$ satisfying:

$$\tau'(t) = \kappa \cdot \left(1 + \sum_{i=1}^N |\Delta \mathcal{B}_i(t)| \right)$$

where κ is a normalisation constant. This ensures high-resolution renderings of epistemic shocks (e.g., landmark replications, retractions) while compressing periods of stability.

Definition 108 (belief cascade overlay): A claim cluster $\Phi = \{\phi_1, \dots, \phi_n\}$ subject to interdependent updates is visualised as a multivariate event tensor $\mathcal{T}\Phi \in \mathbb{R}^{n \times m}$, where m is the number of discrete update epochs. An animated projection of $\mathcal{T}\Phi$ over time enables the user to perceive both diffusive and abrupt propagations through the semantic network, with stability zones encoded via spectral gradient consistency.

Axiom 72 (provenance-preserving visual transform): Every visual transformation $V: \{B_i(t)\} \rightarrow \mathbb{R}^2$ must admit an inverse V^{-1} such that: $x \in \text{Render}(V)$, $V^{-1}(x) \rightarrow (\phi_i, t_k, \Gamma_i, \Delta_i)$ ensuring that no graphical abstraction severs the audit trail or detaches from the cryptographically anchored knowledge graph.

Through these constructs, belief visualisation becomes not merely illustrative but epistemically rigorous enabling dynamic, transparent inspection of the system's rational evolution.

System Integrity and Provenance

This section addresses the foundational infrastructure required to ensure that BEWA's epistemic operations are not merely accurate, but also secure, auditable, and irreversibly anchored. In an environment where the integrity of scientific inference is paramount, it is not sufficient to generate belief updates algorithmically; those updates must themselves be immune to tampering, retrospectively traceable, and cryptographically verifiable. BEWA is thus built upon a provenance-first design, in which every assertion, transformation, and belief revision is recorded, hashed, and anchored to a persistent and immutable ledger. The system does not assume trust it enforces it through cryptographic architecture and rigorous historical traceability. Every claim object, every piece of metadata, and every belief state change is secured against both accidental corruption and deliberate subversion.

Beyond cryptographic anchoring, BEWA integrates a robust audit trail mechanism that provides complete transparency

across the system's decision processes. Each belief trajectory can be decomposed into a temporally ordered series of evidence applications, credibility adjustments, and network interactions allowing for post hoc forensic analysis and institutional scrutiny. This design ensures that no epistemic drift or structural manipulation can occur without immediate detection and full exposability. Furthermore, the system incorporates protocols for so.

Cryptographic anchoring of claims

Ensuring the immutability, auditability, and trust of claims within an epistemic network requires that each propositional unit ϕ_i be verifiably fixed at the time of its assertion and resistant to post-hoc modification. This is achieved through cryptographic anchoring: The process of associating each claim with a tamper-evident and chronologically provable cryptographic commitment embedded in a distributed ledger or equivalent verifiable data structure.

Definition 109 (claim commitment hash $H(\phi_i)$): Let ϕ_i denote a structured, normalised claim. Its cryptographic anchor is defined by a secure hash: $H(\phi_i) := H(\text{serialize}(\phi_i) \parallel \text{meta}(\phi_i))$, where H is a collision-resistant hash function (e.g., SHA-256), and $\text{meta}(\phi_i)$ encodes the claim's timestamp τ_i , author ID, source ID, and version number.

Axiom 81 (uniqueness of anchoring): $\forall \phi_i, \phi_j, \phi_i \neq \phi_j \Rightarrow H(\phi_i) \neq H(\phi_j)$, under the assumption of pre-image resistance and collision-resistance of H .

Definition 110 (Merkle root embedding): Claims ingested within the same temporal epoch ϵ_t are grouped into a Merkle tree T_t , with leaf nodes $\{H(\phi_i), \dots, H(\phi_j)\}$ and root R_t . The root R_t is then time-stamped and embedded into a tamper-evident ledger L : $L_t := L_{t-1} \parallel \text{Tx}(R_t, \tau_t, \sigma_{\text{admin}})$, where σ_{admin} is a digital signature by the anchoring authority or decentralised key quorum.

Proposition 33 (verifiability): Given any claim ϕ_i and a public ledger L containing root R_t , the membership of ϕ_i in epoch ϵ_t can be verified in $O(\log n)$ time via a Merkle proof path.

Definition 111 (anchored claim tuple $\Phi_i^\#$): The complete representation of a claim within the anchored system is: $\Phi_i^\# := (\phi_i, H(\phi_i), \tau_i, R_t, \pi_i, \rho_i)$, where π_i is the prior belief and ρ_i is the claim's reference in the belief graph.

Axiom 82 (cryptographic finality): Once anchored, ϕ_i is epistemically finalised: No retroactive modification is permitted. If revision is required, a superseding claim ϕ_j must be anchored independently with an explicit reference to ϕ_i as predecessor: $\phi_j > \phi_i \Rightarrow \text{meta}(\phi_j).\text{parent} := H(\phi_i)$.

Definition 112 (temporal anchor map A): Let $A: T \rightarrow R$ map times to Merkle roots such that: $\forall_t \in T, A(t) = R_t$. By anchoring claims cryptographically, the system ensures each knowledge assertion becomes a fixed point in the epistemic topology verifiable, immutable, and resistant to manipulation. This establishes a foundational substrate for all higher-order trust, weighting, and reasoning operations across the network.

Auditability and historical traceability

In any epistemically robust system where claims influence belief propagation and decision-making, the capacity to reconstruct the provenance, transformation, and historical interpretation of each claim is foundational. This subsection formalises the auditability requirements and defines protocols that render the history of claims and their associated belief states both immutable and interrogable.

Definition 113 (epistemic record $\mathcal{E}(\phi_i)$): For a given claim ϕ_i , the epistemic record is the complete chronological history of its instantiations, references, updates, and influence over the network. Formally,

$$\mathcal{E}(\phi_i) := \langle (\phi_i^{(t_0)}, \pi_0), (\phi_i^{(t_1)}, \pi_1), \dots, (\phi_i^{(t_n)}, \pi_n) \rangle$$

where each $\phi_i^{(t_k)}$ denotes the state of the claim at time t_k , and π_k the belief weight assigned at that point.

Axiom 91 (sequential traceability): For all $t_k < t_{k+1}$, $\phi_i^{(t_k)} \neq \phi_i^{(t_{k+1})} \Rightarrow \exists \text{update}_{k \rightarrow k+1} \in U$, where U is the set of permitted claim update operations (e.g., citation delta, replication event, contradiction handling). Each operation must be cryptographically recorded and time-stamped.

Definition 114 (immutable claim ledger $L\phi$): The ledger of a claim ϕ is an append-only sequence:

$$\mathcal{L}_\phi = \bigcup_{k=0}^n \text{Tx}_k(\phi_i, \text{op}_k, \pi_k, t_k, \sigma_k)$$

where Tx_k records the operation performed, the resulting belief score, timestamp t_k , and signature σ_k attesting to the legitimacy of the actor or process.

Definition 115 (verifiable trace chain $\mathcal{T}\phi$): A trace chain $\mathcal{T}\phi$ is a hash-linked sequence of ledger entries: $\mathcal{T}\phi := H(Tx_0) \rightarrow H(Tx_1) \rightarrow \dots \rightarrow H(Tx_n)$. This construction guarantees tamper-resistance and order-preservation.

Proposition 41 (audit consistency theorem): For any two auditors A and B with access to the same public ledger L ,

$$\mathcal{T}_\phi^A = \mathcal{T}_\phi^B$$

provided L is synchronised and hash functions are consistent. This guarantees epistemic transparency across trust domains.

Axiom 92 (fork-prevention constraint): A claim ϕ may only have one active trace at any time. Branching updates must occur via supersession: $\exists \phi_j : \phi_j > \phi_i \Rightarrow \mathcal{T}\phi_i \subsetneq \mathcal{T}\phi_j$.

Definition 116 (audit query protocol A): Given a claim ID and a time window (t_a, t_b) , the audit query returns: $A(\phi_i, t_a, t_b) = \{Tx_k \in L_\phi \mid t_a \leq t_k \leq t_b\}$, enabling forensic reconstruction of claim belief trajectories, author actions, and downstream epistemic impact.

This systematic commitment to traceability ensures that no claim can rise or fall in influence without public, cryptographically bound records, enabling accountable scientific discourse and formal rational auditability.

Security, sovereignty, and tamper-proofing

The epistemic system described herein operates in adversarial informational environments where claim injection, manipulation of weights, or surreptitious revisionism could compromise the truth-promoting objective function. This subsection formalises the architectural and algorithmic safeguards ensuring that claim security, authorial sovereignty, and global tamper-proofing are not merely assumed, but mathematically guaranteed.

Definition 117 (claim sovereign space Σ_ϕ): For each canonical claim ϕ , define a sovereign space Σ_ϕ as the bounded domain in which only authorised transformations and referenced interactions can occur. Formally: $\Sigma_\phi := \{Tx_i \mid \sigma_i \in S_\phi\}$, where S_ϕ is the authorisation set cryptographically defined for ϕ , including claim authors, authorised replicators, and verified peer reviewers.

Axiom 101 (non-repudiable authorial rights): A claim ϕ entered into the system at time t_0 under digital signature σ_α must retain immutable attribution to its originator: $\forall t > t_0$, $\text{author}(\phi) = \alpha \Leftrightarrow \text{Verify}(\phi, \sigma_\alpha, PK_\alpha) = \text{true}$.

Definition 118 (system-wide Merkle digest M_t): At time t , define a recursive hash summarisation over all claims ϕ_i and their trace chains $\mathcal{T}\phi_i$: $M_t := \text{MerkleRoot}(\{H(\mathcal{T}\phi_i) \mid \phi_i \in \Phi_t\})$.

This digest is committed to a distributed, auditable ledger and ensures global immutability.

Proposition 47 (global tamper-evident condition): If $\exists \phi_i$ such that $\mathcal{Y}_{\phi_i}^{(t)} \neq \mathcal{Y}_{\phi_i}^{(t+\delta)}$, then:

$$H(\mathcal{T}_{\phi_i}^{(t)}) \neq H(\mathcal{T}_{\phi_i}^{(t+\delta)}) \Rightarrow M_t \neq M_{t+\delta}$$

implying detection via periodic Merkle digest comparisons.

Definition 119 (tamper-proofness oracle Ω): Define $\Omega : \Phi_t \rightarrow \{\text{secure}, \text{compromised}\}$ as:

$$\Omega(\phi_i) := \begin{cases} \text{secure} & \text{if } \forall k, \text{Verify}(Tx_k, \sigma_k, PK_{\phi_i}) = \text{true} \\ \text{compromised} & \text{otherwise} \end{cases}$$

Axiom 102 (bounded external write principle): All operations Op_j on ϕ_i must satisfy:

$$Op_j \in \mathcal{U}_{\phi_i}, \quad \text{signed}(Op_j) \in \Sigma_{\phi_i}, \quad \text{and} \quad Tx_j \subset \mathcal{L}_{\phi_i}$$

ensuring only internally scoped, signed actions affect the epistemic status.

Definition 120 (cryptographic anchoring commitment C_t): Let $C_t = \text{Hash}(M_t \parallel \text{nonce}_t)$ be a hash committed to a permissioned timestamping network, anchoring claim integrity at global system time t .

This architecture ensures that not only are claims bound to their originators, but that the integrity of belief propagation is enforced through algorithmic sovereignty where authors maintain control, transformations are verifiable, and every form of epistemic manipulation is detectable, bounded, and accountable.

Conclusion and Future Work

The design and realisation of a truth-promoting Bayesian epistemological engine represents a comprehensive convergence of epistemic formalism, probabilistic reasoning, and algorithmic rigour. The architectural synthesis presented herein consolidates the theoretical infrastructure necessary for the intelligent parsing, structuring, and belief-weighting of scientific claims within a fully traceable and dynamically re-evaluative system. Each component from source ingestion and canonical authoring, to semantic belief propagation and decay modelling has been constructed to reflect strict axiomatic principles, reproducible scientific metrics, and scalable formal logic. The system promotes epistemic integrity by conditioning claim weightings on reproducibility, peer consensus, and historical provenance, thereby embedding within its architecture a form of active, computable rational scepticism.

Beyond the methodological backbone, the platform formalises epistemic utility by mathematically capturing the interaction between credibility, replication, and citation dynamics within a probabilistic framework. In doing so, it not only responds to claims as static entities but integrates them as participants in an evolving belief network governed by truth-conducive Bayesian optimisation. Mechanisms such as temporal decay, contradiction normalisation, and replicative reset protocols are not auxiliary features but essential

safeguards, mathematically enforcing resilience against epistemic drift and noise.

The architecture ultimately aims at converging toward stable, high-confidence propositional belief states with rigorous traceability to underlying data sources, ensuring no epistemic closure without sufficient probabilistic justification.

The scope for further elaboration and refinement is substantial. While the current schema has been constructed to reflect a high degree of modularity and extensibility, future efforts will expand upon computational expressiveness, autonomous knowledge generation, and integration with domain-specific ontologies. Additionally, enhancing interface interactivity, implementing belief-based query visualisation, and developing de-centralised protocols for sovereign integrity anchoring are all avenues requiring structured investigation. As a scientific engine, the system is conceived not merely to model belief but to contribute to the autonomous formalisation and critique of scientific knowledge, initiating the foundation for a rigorous machine-epistemology grounded in both logic and scientific realism.

Summary of architecture

The architecture of the system is predicated upon a layered epistemological stack grounded in Bayesian probability theory and formal logic, with a strict emphasis on traceability, source integrity, and truth-conducive inference. At the lowest layer, authoritative data ingestion mechanisms process structured and semi-structured texts from scientific repositories, enforcing canonicalisation procedures that disambiguate authorship, claim origin, and publication metadata. This is immediately followed by a propositional parsing layer wherein all epistemic content is transduced into formal claim structures propositions represented in first-order logical schema with attached contextual indices and temporal signatures.

Subsequent layers govern belief attribution and update mechanisms. Prior formation draws upon author credibility scores, domain trust calibrations, and known replication histories. Belief updates are performed through exact Bayesian conditioning, where posteriors integrate new evidence sources adjusted for citation latency, replication authority, and peer review metadata. Special attention is given to the maintenance of belief coherence under contradictory evidence: probabilistic contradiction models and decay functions modulate the influence of outdated, refuted, or anomalous data in accordance with probabilistic divergence and entropy thresholds.

At the system level, claims are not isolated. They are embedded in a directed acyclic graph of propositional

dependencies where belief propagation is mediated by semantic and logical linkages. This cross-claim network formalises the dependency topology of epistemic assertions and enables propagation of belief updates through both direct citation and inferred logical entailment. All computations are cryptographically anchored to ensure tamper-proof integrity, while interfaces are exposed through a query and audit

API, facilitating both real-time interrogation and longitudinal epistemic analysis. The architectural model is thus a closed-loop rational system, capable of iterated self-correction, alignment with scientific consensus dynamics, and computational scepticism grounded in formally defined axioms.

Limitations and challenges

Despite its foundational rigor and systematic coherence, the proposed architecture confronts several limitations that are intrinsic to both epistemic modelling and real-world data integration. Foremost is the ontological underdetermination of structured claim representation. While the use of first-order logic and propositional schemata enables formal representation, natural language propositions often remain context-sensitive, under-specified, or epistemically ambiguous. Efforts to canonicalise claims must therefore confront the Löwenheim-Skolem problem in formal semantics: there exist multiple non-isomorphic models for a given set of logical formulae, complicating efforts at stable interpretation and entailment propagation.

A second critical limitation arises in the weighting of belief updates in the presence of inconsistent or adversarial input. While Bayesian models offer robust priors and posteriors under idealised assumptions of data fidelity, real scientific ecosystems often include fraudulent publications, citation gaming, and unretracted but discredited material. The system attempts to mitigate this through decay functions and contradiction processing, but no purely statistical method can resolve epistemic conflict without presupposing ground truth or normativity a task that invites regress unless externally anchored. The architecture therefore incorporates penalisation structures and recursive epistemic networks, but these remain vulnerable to latent bias in structural priors and institutional trust assignments.

Moreover, computational tractability presents challenges of scale. The graph-theoretic representation of cross-claim dependencies and the dynamic re-weighting of belief states across time introduces NP-hard problems, particularly under belief propagation in cyclically approximated networks or where combinatorial explosion emerges from fine-grained semantic disambiguation. Approximations such as Markov blanket

reductions and belief threshold pruning are employed, but this inevitably introduces epistemic compression and potential loss of fidelity. Finally, while cryptographic anchoring ensures integrity and auditability, it does not guarantee correctness; false claims may be securely recorded, necessitating an ongoing philosophical distinction between provenance and truth.

Prospects for autonomous scientific reasoning

The architectural paradigm outlined in this work sets the foundational substrate for the emergence of autonomous scientific reasoning systems capable not merely of pattern detection, but of formal epistemic judgement constrained by logical validity, probabilistic coherence, and structured domain knowledge. At the intersection of formal epistemology, statistical learning, and automated reasoning lies the prospect of a system that engages with claims, not as isolated data points, but as dynamic epistemic propositions embedded within an evolving knowledge network. The application of Bayesian inferential structures across temporally anchored, semantically indexed, and authoritatively sourced claims enables a machine agent to enact a form of constrained rationality: an epistemic agent that not only updates beliefs, but scrutinises contradiction, weights replication, and assesses the downstream implications of belief revision.

Future extensions will involve incorporating more powerful inferential schemas that go beyond conditional probability, including modal logics of necessity and possibility, counterfactual reasoning (via Lewisian or Pearlian semantics), and algorithmic causal discovery. Such capacities will support not merely the assimilation of scientific findings, but their critical reconstruction in the light of new data, hypothesis testing, and even meta-analytic synthesis. The introduction of deductive-constructive mechanisms such as type-theoretic verification of scientific models, Coq-assisted proof structures for claim derivation, or category-theoretic mapping of ontological types will push the architecture from Bayesian updating to full scientific rational reconstruction.

Autonomous scientific reasoning also necessitates a procedural ethics of inquiry: Mechanisms for audit, challenge, and redress. Such a system must not merely assimilate peer-reviewed literature but evaluate the social structure of peer review itself, detecting collusive citation clusters, authorial dependence networks, or epistemic monopolisation. Through recursive evaluation of authority, dynamic reassessment of belief, and iterative challenge mechanisms, the proposed system paves the way for epistemically autonomous agents capable of engaging in structured dispute, reasoned consensus, and the principled rejection of unfounded claims. This constitutes not

only a technical advance but a shift in the philosophy of machine

References

1. JPA Ioannidis (2005) Why most published research findings are false. *PLoS Med* 2: e124. [Crossref], [Google Scholar], [Indexed]
2. SA Greenberg (2009) How citation distortions create unfounded authority: Analysis of a citation network. *BMJ* 339: b2680. [Crossref], [Google Scholar], [Indexed]
3. FP Ramsey (1926) *The Foundations of Mathematics and Other Logical Essays*. 156-198. [Crossref], [Google Scholar]
4. BD Finetti (1937) *La pr´evision: Ses lois logiques, ses sources subjectives*. 7: 1-68. [Google Scholar]
5. BCV Fraassen (1989) *Laws and Symmetry*. [Google Scholar]
6. JM Joyce (1998) A nonpragmatic vindication of probabilism. *Philos Sci* 65: 575-603. [Crossref], [Google Scholar]
7. Karl RP (1934) *Logik der Forschung*. Springer, Vienna.
8. C Howson, P Urbach (1989) *Scientific reasoning: The bayesian approach*. [Google Scholar]
9. A Doucet, ND Freitas, N Gordon (2001) *Sequential monte carlo methods in practice*. [Google Scholar],
10. JA Goguen, JW Thatcher, EG Wagner, JB Wright (1979) An introduction to algebraic semantics of abstract data types. *Inf Ctrl* 43: 152-179.
11. R Parikh (1997) Ronald Fagin, Joseph Y. Halpern, Yoram Moses, and Moshe Y. Vardi. *Reasoning About Knowledge*. MIT Press, Cambridge, Mass., and London 1995, xiii + 477 pp. *J Symb Log* 62: 1484-1487 [Crossref], [Google Scholar]
12. J Pearl (1988) *Probabilistic reasoning in intelligent systems: Networks of plausible inference*. [Google Scholar]
13. JY Halpern, R Pucella (2006) Using reasoning about knowledge to analyze security protocols. 8: 1-61.
14. GC Necula (1997) Proof-carrying code. 106-119. [Crossref], [Google Scholar]
15. CH Chu, JA Evans (2003) Modeling preferential attachment and knowledge accumulation in scientific citation networks. *Scientometrics* 66: 421-432.
16. KJS Zollman (2007) The communication structure of epistemic communities. *Phil Sci* 74: 574-587. [Crossref], [Google Scholar]
17. Abhijit V. Banerjee (1992) A simple model of herd behavior. *Q J Econ* 107: 797-817. [Crossref], [Google Scholar]
18. S Goldwasser (1998) New paradigms in theoretical computer science. *ACM Computing Surveys* 30 (4es):153.
19. JJ Ch (2003) Meyer and Wiebe van der Hoek. A theory of epistemic reliability. *Artif Intell* 149: 75-97.

20. MR Munafo, BA Nosek, DVM Bishop, KS Button, CD Chambers, et al. (2017) A manifesto for reproducible science. *Nat Hum Behav* 1: 0021. [Crossref], [Google Scholar], [Indexed]
21. AA Ferreira, A Veloso, WMeira Jr (2012) Brief author name disambiguation in medline. *J Am Med Inform Association* 19: 244-248.
22. IS Kang, SS Na (2009) Name disambiguation in author citations using a k-way spectral clustering method. 21-25.
23. I Beltagy, K Lo, A Cohan (2019) Scibert: A pretrained language model for scientific text. 3615-3620. [Crossref], [Google Scholar]
24. J Kim, J Diesner (2018) Author name disambiguation using a graph model with node splitting and merging based on bibliographic information. *Scientometrics* 115: 997-1019.
25. A Narayanan, J Bonneau, E Felten, A Miller, S Goldfeder (2016) Bitcoin and cryptocurrency technologies: A comprehensive introduction. [Google Scholar]
26. GS Manku, A Jain, AD Sarma (2007) Detecting near-duplicates for web crawling. 141-150. [Crossref], [Google Scholar]
27. CD Manning, M Surdeanu, J Bauer, J Finkel, S Bethard, et al. (2014) The stanford corenlp natural language processing toolkit. 55-60. [Google Scholar]
28. P Liang, J Berant, M Iyyer, J Krishnamurthy, WT Yih (2013) Learning dependency-based compositional semantics.
29. M Gardner, J Grus, M Neumann, O Tafjord, P Dasigi, et al. (2018). Allennlp: A deep semantic natural language processing platform. 1-6. [Crossref], [Google Scholar]
30. M Valenzuela, M Ha, O Etzioni (2015) Identifying meaningful citations. 244-250. [Google Scholar]
31. D Zeng, K Lo, W Ammar, S Kohlmeier, V Zayats, et al. (2020) A literature graph-based knowledge extraction system for scientific publications. arXiv: 2005.00596.
32. LL Wang, K Lo, Y Chandrasekhar, R Reas, J Yang, et al. (2020) Cord-19: The Covid-19 open research dataset. arXiv: 2004.10706. [Google Scholar], [Indexed]
33. A Berti, JV Benthem, E Pacuit (2015) A temporal logic framework for reasoning about belief dynamics. *J Philos Log* 44: 635-664.
34. A Graves (2016) Stochastic backpropagation and approximate inference in deep generative models.
35. A Altmejd, A Dreber, E Forsell, J Huber, T Imai, et al. (2019) Predicting the replicability of social science lab experiments. *PLoS One* 14: e0225826. [Crossref], [Google Scholar], [Indexed]
36. C Guo, G Pleiss, Y Sun, KQ Weinberger (2017) On calibration of modern neural networks. 1321-1330. [Google Scholar]
37. NB Shah, J Zou, K Shah, M J Wainwright (2019) Designing peer review systems to incentivize high quality reviews. 33: 5184-5191.