

# Episteme: The Market for Scientific Truth

June 2025

Sasha Shilina, PhD  
sasha@episteme.ac

## Abstract

Contemporary science faces a paradox. Never before have we had such computational power, such abundant data, and such global interconnectivity, yet our systems for validating scientific knowledge remain plagued by opacity, delay, and misaligned incentives. Peer review, the linchpin of credibility, is increasingly gamed or bypassed; retractions surge, reproducibility declines, and artificial intelligence (AI) accelerates both promise and peril by operating on flawed training data. Episteme proposes an alternative: a decentralized, AI-driven prediction market infrastructure that allows scientific hypotheses to be traded as tokenized contracts, resolved by probabilistic AI oracles, and governed by human-verified fallback mechanisms. This paper outlines Episteme's design and philosophical grounding, positioning it as a protocol for incentivized validation and truth-discovery in the age of ontological acceleration.

# Table of contents

<b>1. Introduction.....</b>	<b>6</b>
<b>2. Background and Related Work.....</b>	<b>7</b>
2.1 Problem Statement: Challenges in the Modern Scientific Ecosystem.....	7
2.1.1 Lack of Real-Time Data Synchronization.....	8
2.1.2 The Funding Bottleneck.....	8
2.1.3 The Reproducibility Crisis.....	9
2.1.4 The Devaluation of Validation.....	9
2.1.5 Data Integrity and Trust in Science.....	10
2.1.6 Barriers to Public Engagement and the Opacity of Science.....	10
2.1.7 Delayed Impact.....	10
2.2 Emerging Solutions.....	11
2.2.1 Decentralized Technologies for Science.....	11
2.2.2 Prediction Markets for Science.....	12
2.2.3 AI-for-Science (AI4S).....	13
<b>3. Episteme: Mission, Vision, and a Theory of Market-Mediated Truth.....</b>	<b>15</b>
3.1 Vision and Mission.....	15
3.1.1 Vision.....	15
3.1.2 Mission.....	16
3.2 Mapping Challenges to Mechanisms.....	17
3.2.1 Overcoming Data Fragmentation: AI Oracles and Provenance.....	17
3.2.2 Democratizing Scientific Funding: Tokenized Claims and Markets.....	17
3.2.3 Incentivizing Validation and Replication: Rigor as First-Class Good.....	17
3.2.4 Safeguarding Data Integrity: Verifiable Provenance.....	18
3.2.5 Public Engagement Without Spectacle: Participatory Forecasts.....	18
3.2.6 Trust and Transparency: On-Chain Records.....	18
3.2.7 Accelerating Impact: Evidence-to-Funding Loops.....	19
3.3 Innovation and Positioning.....	20
3.3.1 Competitor Landscape.....	20
3.3.2 Episteme’s Distinctives.....	20
<b>3.4 Design Philosophy.....</b>	<b>22</b>
3.4.1 Ontology: From Papers to Claims.....	22
3.4.1.1 Hypothesis (H).....	22
3.4.1.2 Evidence (E).....	23
3.4.1.3 Belief $P(H E)$ .....	23
3.4.1.4 Resolution (R).....	23
3.4.2 First Principles.....	23
3.4.2.1 Auditability over authority.....	23
3.4.2.2 Incentive compatibility.....	23
3.4.2.3 Pluralism with convergence.....	24
3.4.2.4 Human accountability.....	24
3.4.2.5 Public-good primacy.....	24
3.4.3 Market Epistemology.....	24

3.4.4 Applied Doctrines.....	25
3.4.4.1 Epistemic Robustness.....	25
3.4.4.2 Value of Information as a Funding Primitive.....	25
3.4.4.3 Mechanism Falsifiability (Meta-Epistemics).....	25
3.4.4.4 Temporal Ontology of Truth.....	26
3.4.4.5 Epistemic Justice and Access.....	26
3.4.4.6 Mythotechnics of Legitimacy.....	26
<b>4. Architecture.....</b>	<b>27</b>
4.1 Design Goals.....	27
4.2 Roles.....	27
4.2.1 Authors.....	28
4.2.2 Forecasters.....	28
4.2.3 Verifiers.....	28
4.2.4 Replicators.....	29
4.2.5 Curators and Indexers.....	29
4.2.6 AI Agents.....	29
4.3 Layered Overview.....	29
4.3.1 Base Layer: Humanode L1.....	31
4.3.2 Knowledge Layer: Scientific Database and Provenance.....	31
4.3.3 Asset Layer: Tokenized Claims & Positions.....	33
4.3.4 Interpretation Layer: AI, Evidence, and Collective Belief.....	34
4.3.5 Execution Layer.....	35
4.4 Token Ecosystem.....	36
4.5 Hypothesis Markets.....	37
4.5.1 Core Components.....	38
4.5.2 Lifecycle of a Hypothesis Market.....	38
4.5.2.1 Creation.....	38
4.5.2.2 Activation.....	39
4.5.2.3 Monitoring.....	39
4.5.2.4 Trading.....	39
4.5.2.5 Resolution.....	39
4.5.3 Incentive Flows.....	40
4.5.3.1 Authors.....	40
4.5.3.2 Forecasters.....	40
4.5.3.3 Verifiers.....	41
4.5.3.4 Replicators.....	41
4.5.3.5 Curators & Indexers.....	41
4.5.3.6 The Incentive–Verification Loop.....	41
4.5.3.7 Operational Metrics & Governance Triggers.....	41
4.5.4 Anticipating Failure Modes.....	42
4.5.4 Epistemic Significance.....	44
<b>5. Human–AI Epistemics: Adjudication &amp; Contestability.....</b>	<b>45</b>
5.1 Process Architecture.....	45
5.2 AI oracles as first pass, not final say.....	46

5.3 Traceability by default.....	46
5.4 Designed contestability.....	46
5.4 Human accountability and layered review.....	47
5.5 Calibration and performance guarantees.....	47
5.6 AI-specific failure modes and mitigations.....	47
5.7 Normative boundary.....	47
5.8 Agent ecology.....	47
<b>6. Governance: Polycentric Oversight.....</b>	<b>48</b>
6.1 Decentralization: From Core Team to Polycentric Order.....	48
6.2 Protocol Layer: A Constitutional Baseline.....	49
6.3 Verifier DAO: Procedure as Living Case Law.....	49
6.4 Domain Panels: Specialized Doctrine.....	49
6.5 Hybrid Modalities: Identity, Stake, and Foresight.....	50
6.5.1 1h1v House via Humanode Vortex (Identity).....	51
6.5.2 Token House (Stake).....	51
6.5.3 Futarchy for Bounded Allocations (Foresight).....	51
6.6 Reflexivity and Adaptivity.....	51
6.6 Risks and Mitigations.....	52
6.6.1 Plutocracy and Capture.....	52
6.6.2 Verifiers Cartels.....	52
6.6.3 Over-Correction and Procedural Rigidity.....	52
6.6.4 Ethical Drift.....	52
6.7 Governance as Truth-Making.....	53
<b>7. Revenue Flow &amp; Treasury Policy.....</b>	<b>53</b>
7.1 Sources of Revenue.....	53
7.2 Movement & Accounting.....	54
7.3 Allocation & Accumulation.....	54
7.4 Distribution: Financing Reliability.....	55
7.5 Policy Guards.....	55
7.6 Design Ethos.....	55
<b>8. Roadmap and Long-Term Vision.....</b>	<b>55</b>
8.1 Phased Rollout.....	55
8.1.1 Chōra (Internal Genesis, Q1 2025).....	56
8.1.2 Eidōlon (Public Testnet, Q3 2025).....	56
8.1.3 Krino (Adversarial Testnet, Q1 2026).....	56
8.1.4 Aletheion (Mainnet Launch, 2026).....	56
8.2 Long-Term Vision: Episteme 2030+.....	57
<b>9. Discussion.....</b>	<b>58</b>
9.1 Theory of Change.....	58
9.2 Predictions and Evaluation.....	59
9.3 Fairness and Epistemic Justice.....	61
9.4 Scope and Anti-Goals.....	61
9.5 Limitations and Open Questions.....	61
9.5.1 Core design fragilities.....	61

9.5.1.1 Speculation Versus Rigor.....	61
9.5.1.2 Futarchy Under Adversarial Pressure.....	62
9.5.2 Ethical and epistemic dilemmas.....	62
9.5.2.1 Ethical Conflicts in Research Participation.....	62
9.5.2.2 Complexity of Scientific Prediction.....	62
9.5.2.3 Information Hazards and Dual Use.....	62
9.5.2.4 Insider Information and Market Abuse.....	63
9.5.3 Cultural and institutional adoption barriers.....	63
9.5.3.1 Cultural Adoption Barriers.....	63
9.5.3.2 Governance Fatigue and Participation Inequality.....	63
9.5.3.3 Legal and Institutional Constraints.....	63
9.5.4 Technical and infrastructural fragilities.....	63
9.5.4.1 Premature Automation in Validation Loops.....	63
9.5.4.2 Technical Fragilities.....	64
9.5.4.3 Biometric Equity and Proof of Personhood.....	64
10.13 Open Research Questions.....	64
<b>10. Conclusion.....</b>	<b>65</b>
<b>11. References.....</b>	<b>65</b>

# 1. Introduction

"Knowledge is not for knowing: knowledge is for cutting."

— Michel Foucault, *The Order of Things: An Archaeology of the Human Sciences*

Science has long been celebrated as the bedrock of technological progress and collective enlightenment. Yet, in the early decades of the twenty-first century, this foundation shows visible cracks. The so-called “reproducibility crisis,” first identified in psychology and biomedicine (Ioannidis, 2005; Camerer et al., 2016, 2018), has metastasized across disciplines, undermining confidence in the reliability of published research. Surveys and meta-studies reveal alarming trends: flawed methodologies, selective reporting, publication bias, and outright misconduct (Baker, 2016; Fang & Casadevall, 2016; Prinz et al., 2011; Powell, 2016). In 2024 alone, over 10,000 scientific articles were retracted, a record indictment of systemic dysfunction in peer review, editorial oversight, and research incentives (Else, 2024).

Long considered the guarantor of scientific legitimacy, peer review increasingly reveals its fragility. Studies document bias, opacity, and vulnerability to manipulation in review processes (Grieneisen & Zhang, 2012). As Kitcher (1993) warned, science’s epistemic authority cannot remain unquestioned if its institutional mechanisms fail to deliver reliability. The legitimacy of science is at stake not only in its findings but in the infrastructures through which those findings are validated.

At the same time, the very instruments of sensemaking are transforming. Artificial Intelligence (AI), deployed at nearly every stage of research, from hypothesis generation to literature review and statistical analysis, learns from a corpus already riddled with inconsistencies and errors. Generative models, while powerful, risk amplifying systemic flaws and producing authoritative-sounding but spurious knowledge (Birhane et al., 2021). Indeed, *Nature* has warned of an “AI-driven reproducibility crisis,” where machine learning (ML) workflows are themselves compromised by data leakage, overfitting, and hidden confounders (Ball, 2023). In such a landscape, the latency between discovery and verification widens, while public trust in expert systems wanes.

This crisis is not only academic. The reliability of knowledge underpins global public health, climate action, technological innovation, and democratic decision-making. When scientific infrastructures fail, the consequences ripple through society at large. Traditional remedies (incremental reform of peer review, stricter compliance frameworks, or increased funding alone) are insufficient. What is required is not simply reform but re-architecture: new infrastructures capable of keeping pace with the velocity, complexity, and scale of contemporary knowledge production.

Signs of systemic change are already visible. Since June 2025, *Nature* has published referee reports and author rebuttals by default, signaling a collective push toward transparency in scholarly communication (Nature Editorial, 2025). Parallel to these

reforms, technological revolutions are emerging. Decentralized Science (DeSci) harnesses distributed ledgers and tokenized incentives to reconfigure how research is funded, verified, and governed. AI-for-Science (AI4S) deploys autonomous systems to accelerate discovery and automate analysis. Each alone represents a profound shift; fused together, they suggest the possibility of a new paradigm. We call this convergence Decentralized Science + AI (DeScAI): a recursive, epistemically accountable system that unites provenance, cryptoeconomic incentives, AI oracles, and participatory governance into a coherent architecture of discovery (Shilina, 2025). In this whitepaper, we present Episteme as the first concrete instantiation of such an infrastructure. Episteme is a platform for decentralized, tokenized prediction markets on scientific hypotheses. In Episteme, claims are minted as tradable assets, evidence and attention update their market prices in real time, and AI oracles provide probabilistic resolution. Hypothesis creators receive rewards; forecasters allocate belief through positions that reflect evidence; verifiers ensure epistemic integrity through transparent dispute resolution. In this design, incentives align not with prestige or citation counts but with verifiable truth dynamics.

Episteme does not seek to replace science. It seeks to rewire its circuitry: to transform fragmented and slow processes into ones that programmable, profitable and participatory. By integrating markets, AI, and collective insight, it positions knowledge as a living, collectively priced asset, an epistemic commons open to both scientists and enthusiasts.

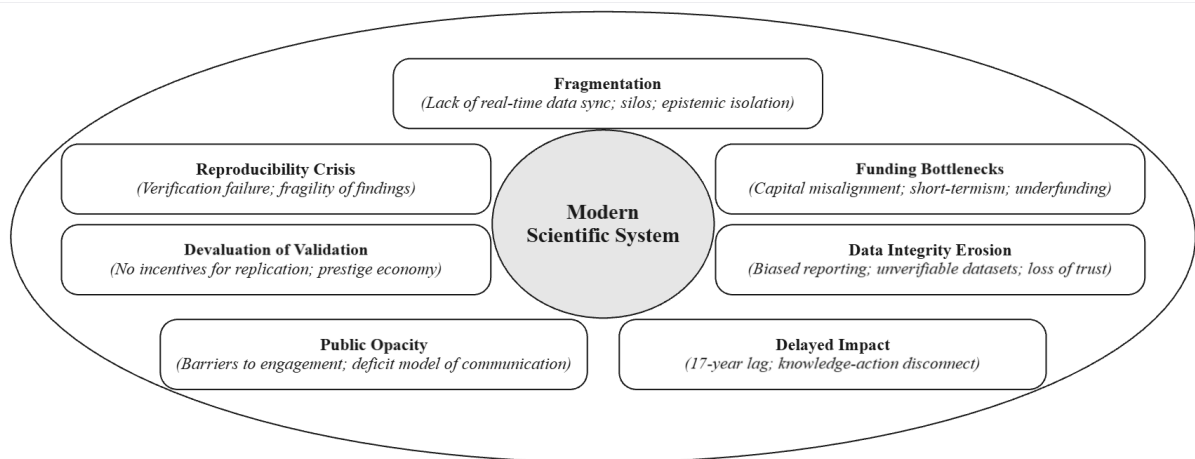
This whitepaper presents Episteme as both a technological platform and a philosophical intervention: a protocolal layer for truth production in an era defined by accelerating complexity.

## 2. Background and Related Work

*“Science must begin with myths, and with the criticism of myths.”*  
— Karl Popper, *Conjectures and Refutations*, 1963.

### 2.1 Problem Statement: Challenges in the Modern Scientific Ecosystem

Science has long been the bedrock of progress, yet today its institutions strain under inefficiency and exclusion. We group the systemic failures of the modern research system into seven interlocking challenges: fragmentation, funding bottlenecks, reproducibility, under-incentivized validation, data integrity, public opacity, and delayed impact.



**Figure 1. Systemic challenges in the modern scientific ecosystem.** *This diagram maps the seven interlocking dysfunctions that hinder contemporary science: fragmented data flows, misaligned funding, a reproducibility crisis, under-incentivized validation, compromised data integrity, limited public access, and chronic delays in real-world application.*

### 2.1.1 Lack of Real-Time Data Synchronization

A central weakness of contemporary science is its failure to achieve real-time data synchronization. What was once a linear, methodical process has become a globalized endeavor generating vast volumes of data, yet these remain scattered across institutions, locked in proprietary silos, and often inaccessible. In this fractured environment, knowledge is isolated, collaboration slowed, and the circulation of information, the lifeblood of discovery, constrained. The National Academies of Sciences, Engineering, and Medicine (2018) identify the inability to access and integrate real-time data as a core bottleneck for open science. Foucault reminds us that knowledge infrastructures are shaped by power, here expressed through institutional gatekeeping. Leonelli (2016) shows how fragmented data flows prevent unified scientific discourse, while Borgman (2015) underscores the epistemological costs of weak data infrastructures. Research data, in this light, resembles puzzle pieces, potentially transformative, yet disconnected and underused.

### 2.1.2 The Funding Bottleneck

A second structural weakness is the persistent disconnect between capital and scientific potential. Despite its ambition, discovery is constrained by slow, competitive, and politically influenced funding systems. Governments and philanthropies still provide most support, yet global R&D spending hovers at only 2–3% of GDP, with sharp regional disparities (UNESCO, 2021). Underfunded fields stagnate, while grant cycles and venture capital shape research toward narrow horizons.

This economic tethering narrows science’s imagination. As Wilsdon et al. (2015) observe, the dominance of competitive grant funding promotes short-termism, rewarding projects with fast returns while sidelining speculative or long-range inquiry. Karl Popper’s reminder that discovery demands risk and openness to the unknown stands in stark contrast to a system that discourages ventures into uncertainty. Van Noorden (2014) further illustrates how many researchers struggle to fund novel or high-risk work, casting doubt on the sustainability of current funding models.

When science is bound to present economic pressures, its capacity to explore the future is diminished.

### 2.1.3 The Reproducibility Crisis

Reproducibility, the principle that scientific results should be independently verifiable, has entered crisis. Large-scale initiatives in psychology, economics, and biomedicine report replication success rates as low as 40–60% (Open Science Collaboration, 2015; Camerer et al., 2016, 2018). Ioannidis (2005) warned early on that “most published findings are false,” due to statistical and methodological biases embedded in research systems.

This is not merely a matter of flawed experiments—it is a signal of deeper epistemic instability. Science, long viewed as a self-correcting process, increasingly resembles an attention economy where fragile findings go viral while rigorous replication is ignored. Baker (2016) and Fanelli (2018) document how institutional pressures (publish-or-perish culture, journal selectivity, and the demand for novelty) undermine reliability. Smaldino and McElreath (2016) call this an “evolutionary” process: academic systems select for visibility over veracity, leading to the proliferation of unreliable claims.

The reproducibility crisis is therefore not merely technical but epistemological. It exposes a tension between science as a pursuit of truth and science as a commodity governed by prestige, funding, and publication metrics. When validation is devalued, error correction stalls, and the credibility of science itself is compromised.

### 2.1.4 The Devaluation of Validation

Despite being a foundational element of the scientific method, replication is systematically devalued in current research ecosystems. Most scientists report difficulty reproducing peers’ results (Baker, 2016), yet few pursue replication work due to its lack of institutional reward. Studies that yield null or corrective outcomes struggle to get published (Franco, Malhotra, & Simonovits, 2014), and replication papers receive significantly fewer citations (Makel, Plucker, & Hegarty, 2012).

The funding and publishing systems are structurally biased in favor of innovation, not verification. As Munafò et al. (2017) observe, replication work is seen as low-prestige and low-impact, despite being essential to the integrity of scientific discourse. Smaldino and

McElreath (2016) show that such systems, left unchecked, select against epistemic robustness.

Until the architecture of incentives changes, the epistemic foundations of science will remain precarious. Validation must be not only encouraged but made viable, visible, and valued.

### 2.1.5 Data Integrity and Trust in Science

Scientific knowledge depends on the integrity of its data, yet this integrity is increasingly in doubt. In an environment of information abundance, data is often selectively reported, manipulated, or framed to fit desired outcomes. Ioannidis (2005) showed that systemic biases in data collection and analysis render a large share of published research unreliable. The result is an erosion of confidence not only in specific studies but in the credibility of science as a whole.

The problem is structural. When datasets are incomplete, unverifiable, or distorted by incentives, the claims built upon them inherit the same fragility. Accuracy, transparency, and reproducibility, the pillars of trust, are compromised, leaving researchers, policymakers, and the public unable to distinguish between findings that can be relied on and those that cannot. In this way, the crisis of data integrity is a crisis of trust itself.

### 2.1.6 Barriers to Public Engagement and the Opacity of Science

Science is often described as a public good, yet paradoxically much of it remains inaccessible. Complex language, gated publications, and opaque processes keep knowledge out of reach from the very publics it is meant to serve. The consequences are measurable: a Pew Research Center survey (2020) found that only 35% of people in developed economies feel confident in their understanding of scientific advancements. This disconnect fuels skepticism toward scientific consensus on urgent issues such as climate change, vaccination, and public health.

Opacity in scientific practice widens the gap between experts and lay publics. As Dunwoody (2014) notes, the perception of science as a closed, elite domain reinforces mistrust and distance. The traditional “deficit model” of communication, where experts disseminate knowledge to passive recipients, has proven inadequate (Gibbons, 1999). Instead, a participatory model is required: one in which publics engage as active interlocutors rather than passive audiences. This call echoes Habermas’ vision of the public sphere, where knowledge emerges through dialogue shaped by societal needs and values, not merely decreed by institutions.

### 2.1.7 Delayed Impact

Even when robust discoveries are made, their translation into application is slow. Classic studies estimate that the average time lag between biomedical research and its uptake in

clinical practice is around 17 years (Balas & Boren, 2000). Similar delays characterize other fields, where findings accumulate in journals but rarely feed back quickly into funding priorities, policy decisions, or technological development.

This latency is not neutral: during public health crises, climate emergencies, or technological races, the gap between evidence and implementation translates into lost lives, wasted resources, and missed opportunities. The disconnect between early signals and downstream action reveals a structural weakness in the scientific ecosystem, knowledge may be produced, but its impact is chronically deferred.

## 2.2 Emerging Solutions

### 2.2.1 Decentralized Technologies for Science

Decentralized Science (DeSci) has emerged as a movement to re-engineer the infrastructure of knowledge production. It leverages distributed ledger technologies (DLTs), smart contracts, decentralized autonomous organizations (DAOs), token-curated registries (TCRs), quadratic funding, and intellectual-property non-fungible tokens (IP-NFTs) to redesign how research is funded, owned, verified, and governed (Asgaonkar & Krishnamachari, 2018; Bamakan et al., 2022; Bischof et al., 2022; Buterin et al., 2019; DeFrancesco & Klevecz, 2022; Fantaccini et al., 2024; Weidener & Boltz, 2025; Weiss, 2022). In shifting epistemic legitimacy from journals, grant committees, and corporate funders toward programmable systems of provenance, incentive alignment, and transparent participation, DeSci represents not merely a technical innovation but a new social contract for science.

Observers have situated DeSci within a broader historical arc: following the alphabet, the printing press, and the World Wide Web, it is described as the fourth great decentralization of knowledge (Weidener & Lukács, 2025). In this frame, DeSci is less a niche experiment than a structural reconfiguration of how societies record, verify, and distribute truth. Early proposals for open peer-review protocols (Tenorio-Fornés et al., 2021) provided the conceptual groundwork, while recent scholarship by Weidener and Spreckelsen (2024) has formalized DeSci's guiding values of transparency, collaboration, and integrity, alongside principles of collective ownership and incentive design. Empirical surveys indicate that the ecosystem is surprisingly resilient: 96% of identified DeSci projects remain active, with partnership-driven DAOs and collective funding pools dominating the landscape (Díaz, Menchaca, & Weidener, 2025). Concrete case studies further illustrate its traction (Unfried, 2024; Ortlepp, 2022).

DeSci also draws heavily on mechanism design innovations. Quadratic funding (Buterin, Hitzig, & Weyl, 2019), enables broad-based communities to match individual contributions with collective pools, rewarding projects with wide support and addressing chronic funding bottlenecks. Token-curated registries (TCRs) extend this logic to curation, using staking and challenge mechanisms to maintain decentralized lists, applied to reviewer pools, data

repositories, or project registries. These primitives complement efforts in scholarly publishing to increase auditability and accountability via blockchain (Bartling & Friesike, 2022).

Finally, DeSci intersects with the broader Open Science movement. The FAIR principles (Findable, Accessible, Interoperable, Reusable) have become normative standards for data stewardship, while blockchain-anchored provenance adds tamper-resistant, verifiable histories of datasets, code, and publications. Identity systems such as ORCID provide unique researcher identifiers that can integrate seamlessly into decentralized infrastructures, closing the loop between attribution, provenance, and incentive flows.

In sum, DeSci represents a diverse but converging set of innovations aimed at restructuring the provenance, ownership, and funding rails of science. Where traditional institutions rely on opaque editorial and grant systems, DeSci experiments with transparent, programmable alternatives rooted in distributed infrastructure and incentive alignment.

Beyond technical novelty, DeSci remedies structural problems long debated in economic theory. Where Adam Smith and Karl Marx disagreed on whether scarcity was natural or imposed, DeSci reduces artificial scarcity by making research data and funding openly accessible. Where Friedrich Hayek argued that dispersed knowledge outperforms centralized authority, DAOs encode this insight into programmable governance. And following Elinor Ostrom's vision of the commons, DeSci treats knowledge not as proprietary resource but as a collectively stewarded public good. In this lineage, Episteme extends DeSci by tokenizing hypotheses themselves, embedding provenance, incentives, and governance directly into the epistemic unit of science.

### 2.2.2 Prediction Markets for Science

Prediction markets have long been studied as efficient mechanisms for aggregating dispersed information into forecasts. Wolfers and Zitzewitz (2004) demonstrated across multiple domains that well-structured markets routinely outperform expert judgment and conventional benchmarks, producing probability distributions that closely approximate Bayesian beliefs. Hanson (1995) extended this argument to science, proposing markets as an institutional alternative to peer review, where claims are continuously priced rather than passively archived. More recent theory reinforces this epistemic function: Osband (2025) conceptualizes markets as *rational learning systems*, where participants update in proportion to expected profit and variance, turning volatility into a signal of collective adaptation rather than noise.

Empirical evidence supports the feasibility of science-focused markets. Dreber et al. (2015) ran prediction markets on the replicability of 44 psychology studies and found that aggregated market probabilities tracked actual replication outcomes more closely than individual experts. Camerer et al. (2016, 2018) replicated this approach in economics and psychology, again showing predictive accuracy in the 40–60% range. A meta-analysis of replication forecasting (Gordon et al., 2021) confirmed these results across four large-scale

projects, demonstrating that markets and structured elicitation are viable tools for triaging which studies warrant replication. Holzmeister et al. (2024) advanced this further with “decision markets,” where traders priced which experiments should be prioritized for replication, and those selected indeed replicated at much higher rates. These results illustrate the epistemic promise of markets: low-cost, scalable, real-time assessments of scientific credibility.

At the platform level, however, participation and liquidity have remained persistent challenges. The Foresight Exchange (est. 1994) and New Zealand’s iPredict hosted markets on topics from climate change to new elements of the periodic table, but suffered from thin engagement, sometimes registering fewer than ten trades per contract (Thicke, 2016). SciCast, developed at George Mason University and sponsored by IARPA, was more ambitious: a combinatorial market on science and technology forecasting that showed crowd forecasts could outperform domain experts (Pfeiffer & Almenberg, 2015). Despite promising results, it too struggled to sustain user engagement. Recent crypto-native platforms have begun to address these constraints. Polymarket, with billions in trading volume, demonstrates that blockchain-based markets can achieve robust liquidity and diverse participation, while Augur’s difficulties with governance and market curation highlight the need for careful design.

Alongside participation, resolution remains a second major bottleneck. Markets require trusted, timely, and auditable settlement. Here, emerging work points toward hybrid models. Chakravorti et al. (2021) introduce “artificial prediction markets” for human–AI collaboration, where exogenous agents simulate bounded rationality and improve classification performance. Such hybrid designs suggest how algorithmic oracles and human forecasters might jointly produce calibrated epistemic signals. Ducreé et al. (2021) likewise emphasize how blockchain-based token economies (staking, community participation, prediction markets) can re-engineer collective intelligence for science. Murphy (2023) goes further, situating prediction markets within a speculative “meta-episteme,” a structural transformation of knowledge production comparable to the Scientific Revolution.

### 2.2.3 AI-for-Science (AI4S)

AI-for-Science (AI4S) is re-engineering both the laboratory bench and the research pipeline. It refers to the application of machine learning, generative models, and autonomous agents to augment or perform core scientific tasks: hypothesis generation, literature synthesis, experiment design, and data interpretation. Unlike traditional computational tools, AI4S systems learn from unstructured data, propose novel experimental paths, and iteratively refine knowledge through closed-loop automation.

Its reach is already visible across disciplines. In biomedicine, AI accelerates drug discovery and protein engineering (Dara et al., 2022; Mak et al., 2023; Blanco-González et al., 2023; Qureshi et al., 2023). In materials science, models predict new compounds with

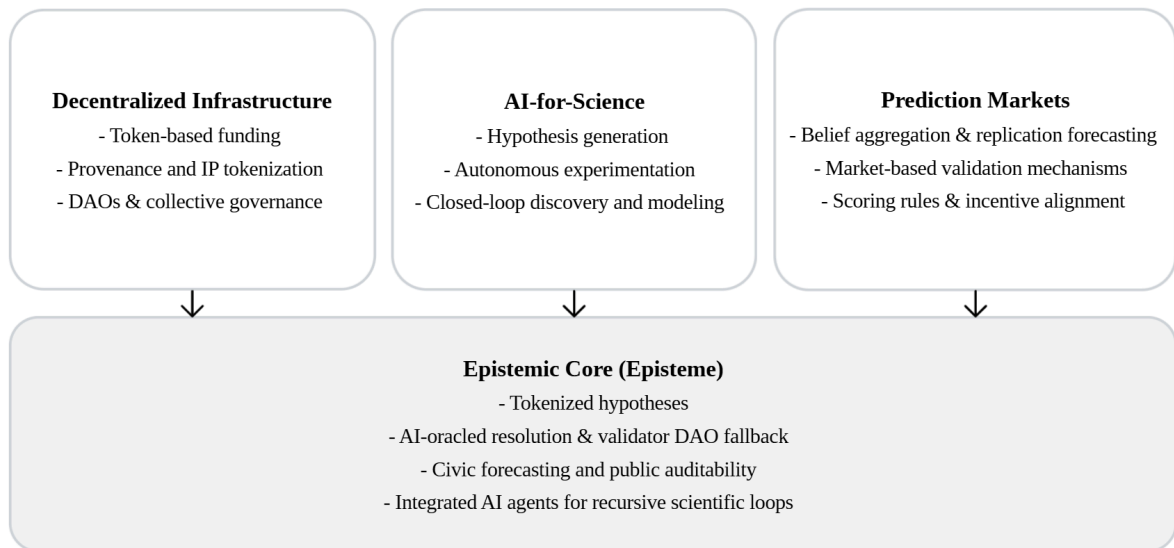
unprecedented accuracy (Sha et al., 2020; Li et al., 2020; Guo et al., 2021). Even in the social sciences, AI-driven literature synthesis and modeling are becoming common (Xu et al., 2024). Breakthroughs illustrate this shift: AlphaFold 2’s near-complete mapping of protein structures (Jumper et al., 2021); autonomous chemist robots capable of self-directed experimentation (Dai et al., 2024); and SAMPLE’s rapid protein engineering loops (Rapp et al., 2024). The “AI Scientist” (Lu et al., 2024) highlights the frontier of autonomous research agents, while recent surveys catalog a fast-growing ecosystem of LLM-based agent frameworks for planning, reasoning, and collaboration (Ferrag et al., 2025). At the institutional level, DARPA’s FoundSci program pursues an “autonomous scientist” capable of domain-general skeptical reasoning (DARPA, 2024). Early demonstrations are striking: Berkeley Lab’s A-Lab, guided by DeepMind’s GNoME, validated 41 new materials in 17 days with a 71% success rate (Biron, 2023).

While DeSci and AI4S have mostly evolved separately, a handful of frameworks now signal their convergence. Wei and Li (2025) propose ISEK, a six-phase architecture (Publish, Discover, Recruit, Execute, Settle, Feedback) that organizes human and AI agents into token-incentivized research loops with recursive participation. Ding et al. (2023) outline an AI4S  $\times$  DLT reference model, emphasizing challenges in provenance, replication, and agent coordination. Kaal (2025) introduces a Directed Acyclic Graph (DAG)-based system with DAO-governed validation pools to audit and adapt AI agent behavior. ETHOS (Chaffer et al., 2025) offers decentralized governance for AI agents, integrating soulbound tokens (SBTs), zero-knowledge proofs (ZKPs), and proportional dispute resolution to ensure ethical and accountable operation.

Taken together, these proposals suggest that integrating AI-native discovery engines with blockchain-native incentive layers could compress validation timelines from years to weeks, reduce gatekeeping bias, and open high-impact research to communities traditionally excluded from funded science. Yet the two domains still largely evolve in isolation: most DeSci protocols prioritize governance, funding, and provenance without integrating intelligent systems, while AI-for-science initiatives often rely on centralized compute, proprietary data, and black-box reasoning. This bifurcation represents a missed opportunity.

### 3. Episteme: Mission, Vision, and a Theory of Market-Mediated Truth

Emerging solutions in DeSci, prediction markets, and AI4S hint at a new architecture for knowledge, but they remain fragmented. Episteme steps into this gap as the synthesis, rewiring flows of knowledge, capital, and public insight into a unified protocol for scientific truth.



**Figure 2. Episteme as the convergence point of DeSci, AI4S, and prediction markets.** This diagram illustrates how three emerging domains (DeSci, prediction markets, and AI) converge into epistemic core.

The name Episteme (from the Greek ἐπιστήμη) means “knowledge” or “understanding.” The project starts from a simple premise: if scientific knowledge is a public good, its credibility should be produced by transparent, incentive-compatible processes, not opaque authority. We therefore treat truth as *programmable*: claims are precisely specified; evidence is registered with verifiable provenance; beliefs are continuously aggregated into prices; and resolutions are justified with recorded reasons. In this frame, markets are not spectacle but *epistemic machinery*, devices that reward accuracy, fund validation, and expose reasoning to audit.

#### 3.1 Vision and Mission

##### 3.1.1 Vision

Episteme aspires to weave together three domains that modern science has allowed to drift apart: *expert inquiry*, *capital allocation*, and *civic participation*. Its vision is a research commons where hypotheses are not only articulated but also *priced, traded, and resolved* under transparent and auditable rules; where capital flows continuously toward decision-relevant questions rather than being stalled in slow bureaucratic cycles; and where

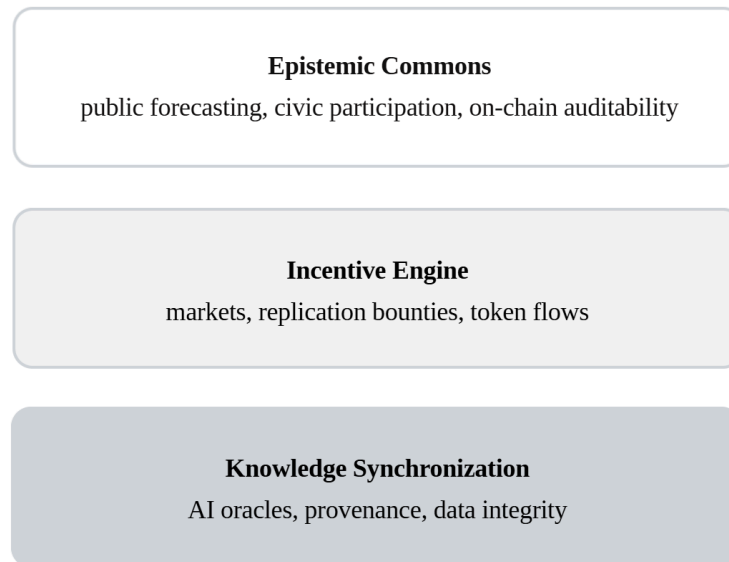
public forecasters, independent researchers, and institutional actors alike can contribute evidence and judgment without passing through the gates of legacy hierarchies.

### 3.1.2 Mission

Episteme’s mission is to build an open scientific economy where capital, computation, and collective intelligence converge to accelerate discovery by:

1. *Synchronizing knowledge* through continuously indexing AI oracles that aggregate and structure global research in real time.
2. *Democratizing funding* via tokenized, AI-resolved hypothesis markets that open access beyond traditional gatekeepers.
3. *Realigning incentives* so that replication, falsification, and validation become not peripheral but economically rewarding activities.
4. *Restoring legitimacy* through on-chain provenance—immutable, transparent, and auditable by the public.
5. *Engaging the public* via gamified participation with real epistemic stakes, fostering literacy and civic involvement in science.

DeFi, DeSci, and AI here are not ends in themselves but composable instruments of epistemic coordination, assembled to transform the production of knowledge into a process that is simultaneously rigorous, participatory, and trustworthy.



**Figure 3. The three functional strata of Episteme’s epistemic infrastructure.** *This diagram outlines the core architecture of Episteme’s design, structured across three vertically integrated strata:*

- *Knowledge Synchronization* forms the foundation, ensuring data integrity and verifiable provenance through AI oracles and distributed timestamping mechanisms.

- *Incentive Engine* enables cryptoeconomic coordination by rewarding replication, and forecasting via token flows and bounty systems.
- *Epistemic Commons* represents the public-facing interface where civic participation, collective forecasting, and on-chain auditability restore trust and invite broader engagement in the scientific process.

These components reimagine the scientific method as a decentralized, incentivized, and publicly verifiable system of knowledge production.

## 3.2 Mapping Challenges to Mechanisms

Episteme addresses the core dysfunctions of modern science through integrated, protocol-level solutions. This section links each systemic problem to a specific mechanism in Episteme’s design, showing how the platform transforms structural weaknesses into verifiable, incentive-aligned workflows.

### 3.2.1 Overcoming Data Fragmentation: AI Oracles and Provenance

The fragmentation of research outputs across institutional silos, paywalled repositories, and heterogeneous formats remains a fundamental impediment to cumulative science (Leonelli, 2016; Borgman, 2015). Episteme addresses this structural problem through AI-powered oracles that continuously ingest, normalize, and classify literature, datasets, and code. Each artifact is assigned cryptographic provenance (content hashes, licenses, versioning), enabling automated verification of integrity and attribution. Rather than treating knowledge as static publications, Episteme transforms hypotheses into living epistemic states, dynamically updated as new evidence arrives. This design produces coherence without centralization: indexing remains transparent, contestable, and auditable, mitigating the epistemic opacity that has historically enabled selective reporting and data exclusion.

### 3.2.2 Democratizing Scientific Funding: Tokenized Claims and Markets

To counter this bottleneck, Episteme introduces tokenized hypothesis markets, where claims function as tradable contracts. Prices in these markets reflect collective credence, aggregating dispersed information into Bayesian-like probability estimates (Wolfers & Zitzewitz, 2004; Dreber et al., 2015). Liquidity and fees channel resources toward questions of highest uncertainty and societal impact, aligning capital flows with epistemic relevance. In doing so, Episteme redistributes funding away from entrenched hierarchies toward a decentralized, demand-driven allocation mechanism, while preserving accountability through transparent market resolution.

### 3.2.3 Incentivizing Validation and Replication: Rigor as First-Class Good

Replication, falsification, and null results remain systematically under-incentivized within the prestige economy of science, despite being cornerstones of epistemic reliability (Ioannidis, 2005; Camerer et al., 2016, 2018; Franco, Malhotra, & Simonovits, 2014).

Episteme reverses this logic by making rigor economically valuable. Market protocols earmark fee flows and bounties for replication attempts; successful replications or falsifications trigger “credence shocks” that directly affect settlement. Auditors and replicators are rewarded with attribution and compensation, while null findings are inscribed as priced signals rather than relegated to the “file drawer” (Makel, Plucker, & Hegarty, 2012; Munafò et al., 2017). Through this mechanism, validation becomes a first-class epistemic good, integrated into the incentive structure rather than marginalized by it.

#### 3.2.4 Safeguarding Data Integrity: Verifiable Provenance

Episteme addresses this structural weakness by embedding machine-verifiable provenance anchors into every stage of the hypothesis lifecycle. Each dataset, model, and replication attempt is hashed, time-stamped, and recorded on-chain, binding the claim to its evidentiary substrate. Licenses and metadata are attached immutably, enabling attribution while reducing the risk of silent data drift. By normalizing content through AI-powered oracles and linking updates to cryptographic signatures, Episteme ensures that claims cannot detach from the evidentiary scaffolding on which they rest. This architecture transforms integrity from a reputational quality to a cryptographic guarantee: datasets are not merely available, but auditable, tamper-evident, and reproducible *ex post*. In doing so, it restores confidence in the scientific record at a time when trust is increasingly strained (Zhang et al., 2021).

#### 3.2.5 Public Engagement Without Spectacle: Participatory Forecasts

Science’s opacity has widened the gulf between expert inquiry and public understanding, fueling skepticism and eroding trust (Pew Research Center, 2020; Dunwoody, 2014). Episteme reconfigures this relation by introducing participatory forecasting mechanisms that allow non-specialists to contribute credence judgments, while still grounding engagement in formal epistemic tools such as calibration scoring and structured evidence feeds (Gibbons, 1999; Habermas, 1989). Unlike shallow “gamification,” participation here is instrumental: it improves forecast accuracy, channels resources to replication, and fosters civic epistemology. By embedding publics into transparent market processes, Episteme cultivates literacy through practice, transforming spectators into epistemic co-participants.

#### 3.2.6 Trust and Transparency: On-Chain Records

The crisis of trust in science, rooted in opaque peer review, unverifiable data, and retrospective retractions, demands structural remedies (Ioannidis, 2005; Zhang et al., 2021). Episteme anchors all critical artifacts (market specifications, oracle drafts, verifier votes, dispute outcomes) on-chain, where they are immutable and independently auditable *ex post*. In this model, transparency is not a rhetorical aspiration but a data structure: every epistemic decision is recorded, attributed, and subject to public inspection. By rendering

the validation pipeline visible, Episteme restores procedural legitimacy and ensures that the authority of science rests on verifiable processes rather than institutional trust alone.

### 3.2.7 Accelerating Impact: Evidence-to-Funding Loops

Episteme compresses this timeline by creating real-time feedback loops between evidence and capital. Hypothesis markets update continuously as new data streams into AI oracles; prices shift immediately to reflect changing credence. These dynamics channel liquidity toward questions where marginal evidence has the highest impact, ensuring that signals are not archived passively but priced and acted upon. Early findings no longer await bureaucratic approval, they directly modulate the flow of attention, resources, and replication efforts.

In this model, discovery and application are not sequential but recursive. Each publication or dataset enters the market as live input, affecting funding decisions within hours or days rather than years. By collapsing the evidence–impact lag, Episteme transforms knowledge production into an economy of real-time epistemic responsiveness, aligning the tempo of science with the urgency of societal challenges.

<b>Table 1. Episteme systemic responses to core failures in the scientific ecosystem</b>	
<b>Problem</b>	<b>Episteme Response</b>
Fragmented, delayed synchronization	AI oracles continuously index and structure global research with verifiable provenance.
Validation under-incentivized	Markets pay for accuracy; replication and falsification earn explicit rewards.
Conservative, siloed funding	Open hypothesis markets route capital to under-served, high-uncertainty domains.
Public exclusion	Participatory forecasting and transparent records enable civic epistemology.
Trust deficit	Immutable audit trails: resolution logic, data sources, and outcomes are on-chain and reviewable.
Delayed impact	Dynamic prices link early evidence to funding flow and prioritization.
Reproducibility crisis	Replication incentives embedded in the market life cycle and fee flows.

**Table 1. Episteme’s systemic responses to core failures in the scientific ecosystem.** *This table maps seven structural problems in modern science to Episteme’s targeted solutions. Through AI-oracled indexing, incentive-aligned markets, and on-chain auditability,*

*Episteme restructures knowledge production to prioritize verifiability, participation, and timely impact.*

## 3.3 Innovation and Positioning

### 3.3.1 Competitor Landscape

Existing projects illuminate fragments of the problem but stop short of systemic integration. General Web3 prediction markets such as *Polymarket*<sup>1</sup> and *Hedgehog*<sup>2</sup> have proven liquidity and crowd calibration, but they operate without domain-specific provenance or incentives for replication. Science-focused pilots like *SciCast*<sup>3</sup> and *Stadium.Science*<sup>4</sup> showed that collective forecasts can rival expert judgment, yet they struggled with thin engagement and reliable settlement. DeSci DAOs such as *VitaDAO*<sup>5</sup>, *Molecule*<sup>6</sup>, and *BIO Protocol*<sup>7</sup> mobilize community funding for biomedical niches, but their scope is narrow and focused on IP rather than claim-level validation. Gamified platforms like *Foldit*<sup>8</sup> or *Eterna*<sup>9</sup> demonstrate the value of citizen participation, but they remain centralized and disconnected from incentive-aligned provenance. Data-token projects such as *Ocean Protocol*<sup>10</sup> and *Numerai*<sup>11</sup> tokenize datasets or financial signals, but their orientation is toward markets for data or trading models, not epistemic claims.

### 3.3.2 Episteme's Distinctives

Episteme integrates these disparate strands into a claim-native architecture for science. Its innovations can be summarized along five axes:

- *Hypotheses as primitives*. Claims are minted as tokens with canonical specifications, provenance metadata, and resolution criteria embedded from inception.
- *AI-assisted resolution*. Probabilistic judgments are drafted by AI oracles, continuously updated as new evidence arrives, and contestable through human fallback.
- *Validation economics*. Replication, falsification, and auditing are rewarded through bounties, settlement shocks, and fee-routing, embedding rigor into the incentive structure.
- *Polycentric governance*. Protocol rules, validator DAOs, and domain-specific panels interact to provide layered dispute resolution and adaptive oversight.

---

<sup>1</sup> <https://polymarket.com/>

<sup>2</sup> <https://www.hedgehog.markets>

<sup>3</sup> <https://scicast.org>

<sup>4</sup> <https://www.stadium.science/>

<sup>5</sup> <https://www.vitadao.com/>

<sup>6</sup> <https://molecule.xyz>

<sup>7</sup> <https://www.bio.xyz>

<sup>8</sup> <https://fold.it/>

<sup>9</sup> <https://www.eterna.com/>

<sup>10</sup> <https://oceanprotocol.com/>

<sup>11</sup> <https://www.numer.ai/>

- *Public-good primacy.* Resolution artifacts, provenance trails, and oracle drafts remain open by default; evidence cannot be privatized.

Beyond these architectural pillars, Episteme redefines the *market logic* of science. Five differentiators highlight its positioning:

- *AI-fast truth.* Hypotheses are resolved at machine speed, enabling instant and scalable updates that keep pace with research.
- *Science as an asset class.* Research becomes investable, verifiable, and liquid; hypotheses trade as economic instruments rather than static publications.
- *Built-in funding.* Every trade contributes fees that route back into research, turning epistemic activity into a continuous funding stream.
- *Belief, live-tracked.* Market prices act as transparent barometers of collective evidence and shifting credence.
- *Volatility, harnessed.* Instead of being treated as noise, disagreement and price swings are reframed as discovery signals, surfacing uncertainty and driving replication.

Episteme positions itself as a *full-stack epistemic economy*. Where existing efforts isolate liquidity, funding, or participation, Episteme fuses them into one infrastructure: a system where hypotheses are tradable assets, their validity continuously priced, their resolution auditable on-chain, and their incentives tuned toward truth rather than prestige.

**Table 2. Comparative positioning of Episteme across adjacent categories**

Category	Examples	Strengths Shown	Limitations	Episteme's Advance
General Prediction Markets	Polymarket, Hedgehog	Liquidity, real-time forecasting	No provenance; no replication incentives	Claim-native design with AI-oracle resolution
Science Forecasting	SciCast, Metaculus	Demonstrated accuracy of crowd predictions	Thin liquidity; poor settlement mechanisms	Embedded market incentives + verifier fallback
DeSci Funding DAOs	VitaDAO, Molecule, BIO Protocol	Mobilized biotech niches, tokenized IP	Narrow scope; IP/IPT focus	Broad, claim-level hypothesis markets across fields
Gamified Citizen Science	Foldit, Eterna	Mass engagement, participatory problem-solving	Centralized; no incentive alignment	Participatory markets with scoring rules + on-chain audit

Data/Token Projects	Ocean Protocol, Numerai	Data tokenization, model training economies	Oriented to datasets and finance, not truth	Hypothesis-level tokenization + public-good primacy
---------------------	-------------------------	---	---	---

**Table 2. Comparative positioning of Episteme across adjacent categories.** *The table situates Episteme within the broader landscape of prediction markets, science forecasting platforms, DeSci funding DAOs, citizen science initiatives, and data/token projects. While existing models demonstrate strengths in liquidity, engagement, and niche funding, they often lack provenance, replication incentives, or hypothesis-level resolution. Episteme advances the field by combining AI-oracled claim verification, embedded incentive mechanisms, and public-good alignment at the level of epistemic claims.*

### 3.4 Design Philosophy

Episteme is not simply a technical platform; it is an epistemic institution, designed to realign the foundations of scientific legitimacy. Its architecture encodes a set of philosophical commitments, transforming science from a prestige economy into a knowledge economy. What follows are the core doctrines that guide this design, each rooted in long-standing traditions in philosophy of science and epistemology.

#### 3.4.1 Ontology: From Papers to Claims

For centuries, the scientific paper has functioned as the canonical unit of knowledge, reinforced by journals, citation indices, and editorial hierarchies. Yet this print-era architecture is increasingly misaligned with the speed and complexity of modern science. Episteme redefines the ontology of knowledge: hypotheses become precisely specified claims with observable endpoints, admissible evidence streams, bounded timeframes, and explicit resolution criteria. Evidence is stored as content-addressed, versioned artifacts, making it transparent and auditable.

This move resonates with Judea Pearl’s formalization of causal inference (Pearl, 2009) and Donald Rubin’s causal model (Rubin, 1974), both of which insist that claims must be specified in terms of observable outcomes and admissible interventions. By formalizing hypotheses this way, Episteme makes causality explicit, transforming narrative propositions into programmable objects.

##### 3.4.1.1 Hypothesis (H)

In this model, each hypothesis is minted as a precisely specified claim. It carries formal attributes: observable endpoints, admissible data sources, a bounded timeframe, explicit resolution criteria, and a structured provenance graph. This specification transforms hypotheses from narrative propositions into programmable objects whose boundaries and adjudication are clear at inception.

### 3.4.1.2 Evidence (E)

Supporting material is organized as content-addressed artifacts. Datasets, code, experimental protocols, replication attempts, and independent audits are stored with versioning, licensing metadata, and cryptographic hashes. Unlike the static “supplementary materials” appended to papers, evidence here becomes a dynamic layer—transparent, modular, and continuously auditable.

### 3.4.1.3 Belief $P(H|E)$

Credence in a hypothesis is expressed through market prices under a fixed numéraire. This mechanism operationalizes Bayesian updating: as new evidence (E) is introduced, collective belief  $P(H|E)$  is revealed through positions taken by forecasters and participants. Price discovery thus acts as a public, quantitative measure of epistemic confidence, calibrated not by prestige hierarchies but by open participation under shared rules.

### 3.4.1.4 Resolution (R)

Hypotheses culminate in resolution, a scalar, evidence-backed verdict anchored in written rationales, citations, and accountable signatories. Resolution events are not hidden within editorial boards but conducted in a transparent, contestable process, integrating both AI-oracle drafts and human adjudication.

Together, these four primitives (H, E,  $P(H|E)$ , R) constitute Episteme’s ontology, shifting the center of legitimacy from prestige economies to provenance economies.

## 3.4.2 First Principles

Episteme is guided by a set of axiomatic commitments that distinguish its epistemic architecture from both traditional science and prior experiments in decentralized knowledge systems. These principles establish the normative bedrock upon which mechanisms are designed and legitimacy is conferred.

### 3.4.2.1 Auditability over authority

Epistemic legitimacy derives not from institutional prestige or reputational capital but from transparent rules, accessible logs, and publicly available reasons. Every resolution, update, and dispute leaves an immutable record, allowing authority to be exercised through process rather than position (Zhang et al., 2021).

### 3.4.2.2 Incentive compatibility

Participants must earn more by being accurate, or by contributing to others’ accuracy, than by cultivating popularity or rhetorical dominance. This echoes the principle of mechanism design (Hurwicz, 1973; Myerson, 1999), aligning private incentives with collective epistemic welfare.

### 3.4.2.3 Pluralism with convergence

The system permits diverse priors, methodologies, and heuristics. Yet, as evidence accumulates and markets update, credences converge into probabilistic consensus. This design balances the epistemic benefits of diversity (Page, 2007) with the discipline of aggregation.

### 3.4.2.4 Human accountability

While AI oracles draft probabilistic judgments and assist in synthesis, final responsibility for contested outcomes remains human. Contestability ensures that resolution is never delegated to an unaccountable algorithm but remains grounded in deliberative legitimacy.

### 3.4.2.5 Public-good primacy

Evidence must remain reusable. Private rewards (trading profits, replication fees, or attribution credits) cannot privatize the artifacts themselves. By keeping resolution data and provenance trails open, the system sustains knowledge as a global commons (Ostrom, 1990).

## 3.4.3 Market Epistemology

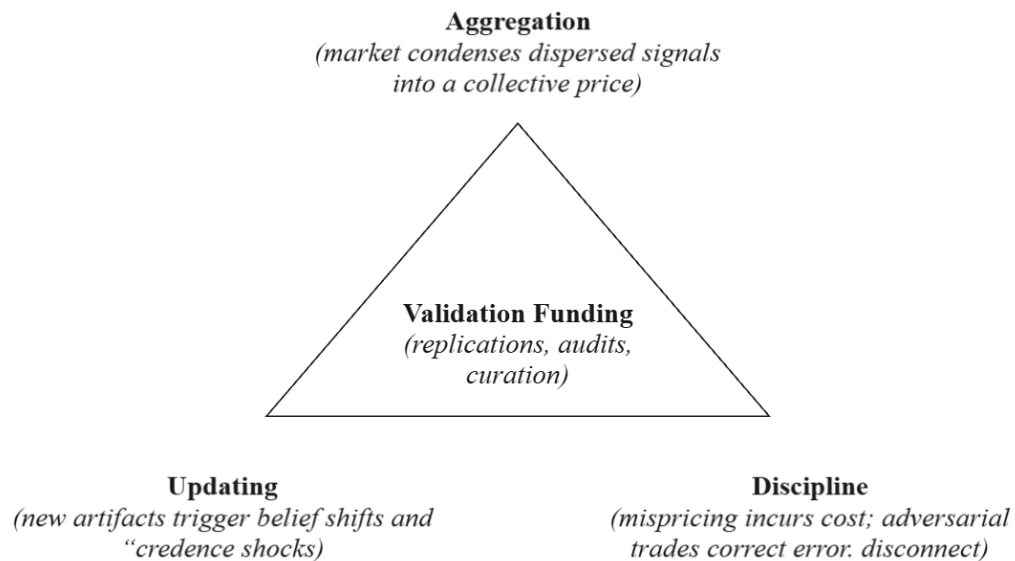
Prediction markets within Episteme function as epistemic engines. Their design provides three core functions:

- *Aggregation.* Prices summarize dispersed private signals and public evidence. Market volatility does not indicate noise to be suppressed but marks the arrival of new information and its assimilation into collective belief (Wolfers & Zitzewitz, 2004; Osband, 2025).
- *Updating.* New artifacts (replication attempts, null results, independent audits) produce measurable shifts in belief. Replications that succeed strengthen credence; falsifications or anomalies induce sharp “credence shocks.” In this way, Bayesian updating is instantiated through liquidity.
- *Discipline.* Mispricing is costly. Biases, misinformation, or misplaced confidence translate into financial losses. Adversarial trading converts epistemic error into someone else’s opportunity, transforming distortion into a corrective force (Hanson, 1995).

Markets, however, are not self-sufficient epistemic machines. Metrics are prone to Goodhart’s law: once a measure becomes a target, it risks distortion (Strathern, 1997). To counteract this, Episteme explicitly routes fee flows to pay for validation externalities—replication studies, audits, and evidence curation. These systemic “checks” ensure that credibility is not merely signaled but actively sustained.

This dynamic aligns with Karl Popper’s principle of falsification (Popper, 1959), Thomas Kuhn’s analysis of paradigm shifts (Kuhn, 1962), and Imre Lakatos’s conception of

progressive versus degenerative research programmes (Lakatos, 1970). Episteme makes these epistemic dynamics continuous and transparent: crises and corrections are registered as price signals.



**Figure 4. Market Epistemology: Aggregation – Updating – Discipline.** *Prediction markets function as epistemic engines: aggregating signals, updating beliefs, and disciplining error through incentives.*

### 3.4.4 Applied Doctrines

#### 3.4.4.1 Epistemic Robustness

Episteme is designed under the assumption that manipulation, collusion, and bias are not rare anomalies but ordinary operating conditions of knowledge systems. Robustness emerges through Sybil resistance (via Humanode’s cryptobiometrics), adversarial market incentives, open audit logs, and slashing mechanisms that penalize negligence or misconduct. Bias and error are not eliminated but converted into opportunities for correction. Robustness is measured not only by average accuracy but by the system’s capacity to remain reliable under epistemic stress.

#### 3.4.4.2 Value of Information as a Funding Primitive

Reliability in science requires active financing, not passive hope. Episteme routes a portion of market fees into replication studies, audits, and evidence curation, allocating resources according to the expected value of information (VOI). This principle directs scarce validation efforts toward areas where marginal epistemic gain is highest. By embedding VOI into its economics, Episteme ensures that replication and verification become sustained, incentivized activities rather than neglected side projects.

#### 3.4.4.3 Mechanism Falsifiability (Meta-Epistemics)

Episteme extends the principle of testability to itself. System-level metrics, resolution error rates, forecaster calibration, and dispute outcomes—are continuously monitored and can trigger governance responses when thresholds are breached. This reflexivity treats the protocol as a hypothesis subject to empirical evaluation and revision, ensuring that the institution evolves under the same epistemic discipline it demands of claims within the system.

#### 3.4.4.4 Temporal Ontology of Truth

Truth is not static but historical. Friedrich Nietzsche described it as a “mobile army of metaphors” (Nietzsche, 1873/1999), Alfred North Whitehead conceived of reality as process (Whitehead, 1929), and Kuhn demonstrated that paradigms evolve through crises (Kuhn, 1962). Episteme encodes temporality directly: hypotheses are assigned half-lives, belief trajectories mapped through volatility curves, and resolution criteria versioned. This makes the time-dependence of knowledge explicit, distinguishing settled consensus from live fronts of inquiry.

#### 3.4.4.5 Epistemic Justice and Access

Legitimacy must expand beyond elite institutions. Miranda Fricker’s concept of epistemic injustice highlights how exclusion distorts credibility (Fricker, 2007), while Habermas emphasizes the role of the public sphere in grounding legitimacy (Habermas, 1989). Episteme encodes justice and access through Humanode’s one-human-one-account primitives, preventing plutocratic capture, and through open licensing and multilingual provenance, ensuring global reusability. Scientific legitimacy is redistributed to a wider epistemic public.

#### 3.4.4.6 Mythotechnics of Legitimacy

Science is not only method but also myth and theater. Hans Blumenberg argued that myth functions as a coping mechanism when rational systems reach their limits (Blumenberg, 1985), while Bruno Latour described science as a staged performance of credibility (Latour, 1987). Episteme acknowledges this dimension by integrating AI agents as epistemic personas: participants that, like humans, wager, forecast, and contest. These agents dramatize epistemic struggle, making credibility legible not just as numbers but as narratives. Episteme’s first such agent, Efi, exemplifies this role: she not only participates in markets but also narrates and embodies the cultural layer of legitimacy. In this way, authority rests not solely on rules and resolution, but also on shared myths that render verification visible, participatory, and culturally resonant.

## 4. Architecture

*“The ideal community of scientists resembles a body politic that works according to the economic principles of independent initiatives coordinated ‘as by an invisible hand’ in the joint endeavor of discovery.”*

— Michael Polanyi, *“The Republic of Science”*, 1962.

Episteme is conceived as an epistemic computing stack: a layered architecture designed to synchronize knowledge, price claims, and settle outcomes under transparent and auditable rules. The guiding objective is to make resolution, not rhetoric, the system’s central product, while preserving contestability, provenance, and due process.

### 4.1 Design Goals

Episteme’s architecture rests on five design principles, each informed by both epistemic theory and mechanism design:

- *Claim-native orientation.* Hypotheses are treated as first-class objects, specified with canonical metadata (observable endpoints, admissible sources, resolution criteria). This responds to long-standing critiques of paper-centric infrastructures, which obscure the claim level under narrative exposition (Latour & Woolgar, 1979).
- *Auditability by default.* Every material step—data ingestion, oracle drafting, validator vote, and settlement—emits machine-verifiable artifacts (hashes, licenses, timestamps). This echoes the transparency standards advocated by the FAIR principles (Wilkinson et al., 2016) and STAR protocols (Nosek et al., 2015).
- *Contestability as norm.* AI is tasked with proposing probabilistic drafts, but humans remain empowered to challenge, arbitrate, and record reasons. This design embodies Foucault’s insight that knowledge is not only produced but continuously contested (Foucault, 1972).
- *Incentive alignment.* Fee flows are explicitly routed to sustain replication, falsification, audits, and curation, activities undervalued in prestige economies (Ioannidis, 2005; Munafò et al., 2017). By internalizing these externalities, the protocol ensures that rigor is not only virtuous but profitable.
- *Polycentricity.* Following model of polycentric systems (Ostrom, 1990), Episteme distributes legitimacy across multiple oversight bodies, reducing the risks of information dominance, capture or unilateral authority.

### 4.2 Roles

Episteme’s architecture rests on distinct but interdependent participant roles, each with incentives aligned toward epistemic reliability. Together, they ensure that individual strategies aggregate into collective truth-seeking dynamics rather than rent-seeking or collusion.

### 4.2.1 Authors

Authors initiate epistemic activity by specifying hypotheses in canonical form. They bear reputational and financial exposure: poorly specified claims are less likely to attract liquidity, while precise, decision-relevant formulations are rewarded through attribution and fee shares. Authorship thus becomes an economic signal of quality.

To safeguard attribution, Episteme introduces verification badges for researchers, institutions, and labs. Verified authorship signals that a hypothesis originates from a recognized source, preventing misrepresentation and building trust with participants. Researchers may claim verified status by linking institutional affiliations, ORCID IDs, or peer-reviewed publications. Once verified, they automatically receive authorship credit and fee shares for their hypotheses, even if minted initially by a curator.

Episteme thus distinguishes three pathways of authorship:

- *Originating Authors*. Researchers who first formulate a hypothesis. With verification, their authorship is cryptographically anchored to their identity.
- *Curator–Minters*. Participants who bring external hypotheses onto the platform. They are recognized for structuring and linking provenance metadata, but priority flows back to originating authors when verified.
- *Derivative Authors*. Participants who extend hypotheses by modifying parameters (e.g., timeframe, endpoints). Their contributions are versioned, linked, and rewarded without erasing original credit.

Automated similarity checks, verifier review, and verification badges ensure transparent provenance. In this way, Episteme protects researchers' intellectual credit while allowing curators and enthusiasts to broaden participation.

### 4.2.2 Forecasters

Forecasters supply liquidity and information by taking “For” or “Against” positions in hypothesis markets. Their profits depend on calibration and marginal information gain, not herding behavior. Well-calibrated beliefs are rewarded as prices converge toward truth, transforming probabilistic honesty into financial utility.

### 4.2.3 Verifiers

Verifiers arbitrate disputes, review admissible evidence, and confirm resolution events. Correct adjudication is rewarded, while negligence or collusion is penalized through slashing. Verifiers embody adversarial accountability: their authority is contestable and transparent.

## 4.2.4 Replicators

Replicators operationalize reproducibility. By preregistering protocols, publishing datasets, and confirming or falsifying claims, they generate market-moving “credence shocks.” Successful replications or falsifications redistribute wealth from miscalibrated to accurate forecasters, compensating replicators directly for their epistemic labor.

## 4.2.5 Curators and Indexers

Curators maintain schemas, taxonomies, and provenance graphs that structure hypotheses and evidence flows. They are rewarded when improvements in metadata measurably enhance resolution accuracy. Curators thus sustain the coherence and navigability of the system.

## 4.2.6 AI Agents

Episteme envisions markets not just as human assemblies but as hybrid ecologies where AI agents act as epistemic citizens. AI agents extend epistemic capacity. They ingest data at scale, generate probabilistic updates, and participate in markets under the same rules as humans: staking, exposure to slashing, and transparent reasoning trails. Following Norbert Wiener’s early warnings on autonomous systems (1960), Stuart Russell’s work on alignment (2019), and Luciano Floridi’s account of distributed agency (2013), Episteme ensures that agents are integrated without being unaccountable. Their outputs remain contestable, subject to verifier oversight and market correction, so that automation amplifies epistemic reliability rather than undermining it.

## 4.3 Layered Overview

Episteme’s architecture is structured as a five-layered stack, with each layer fulfilling a distinct epistemic function while reinforcing the others to create a resilient, composable system.

- Base Layer (Humanode L1)

The protocol operates on a Sybil-resistant, EVM-compatible Humanode chain. While Episteme currently leverages Humanode primarily as a blockchain settlement layer, its biometric primitives open pathways toward one-human-one-account citizen science and one-human-one-vote governance in future iterations.

- Knowledge Layer (Scientific Database & Provenance)

AI-driven oracles ingest literature, datasets, and code; normalize entities; and anchor provenance using content hashes, licenses, and versioning. This ensures that hypotheses and evidence are tethered to verifiable artifacts rather than reputation or authority alone.

- Asset Layer (Tokenized Claims & Positions)

Researchers and enthusiasts submit hypotheses directly into the system. Each hypothesis is minted as an NFT, while “For/Against” positions circulate as tradable tokens. A bonding curve mechanism governs pricing, ensuring liquidity and dynamic adjustment as market participation shifts. Fee routing, staking, and bounty mechanisms encode the economics of validation, aligning incentives toward replication, falsification, and evidence-based challenges.

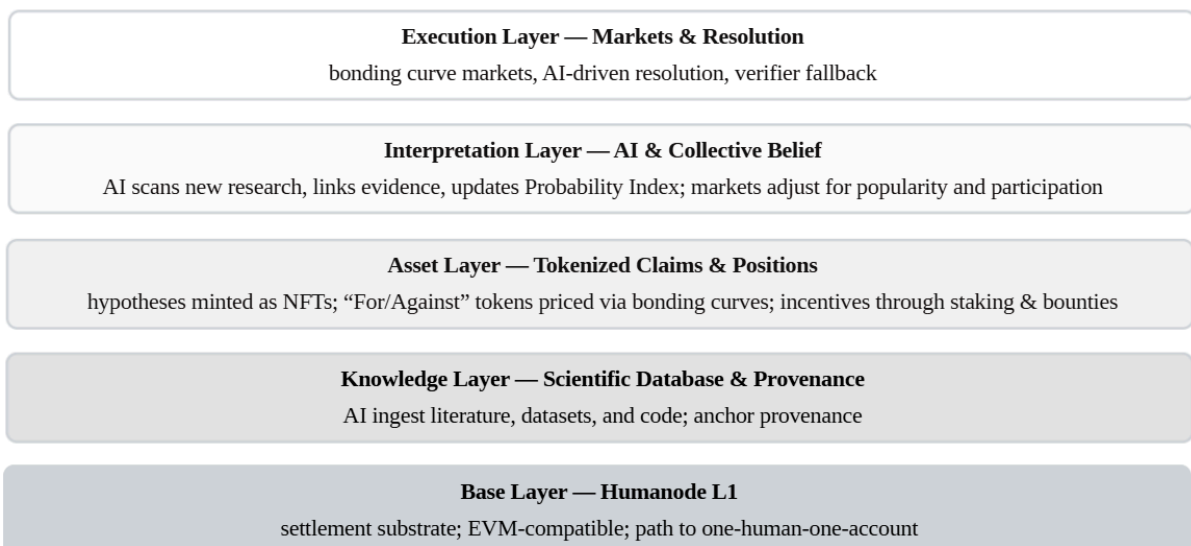
- Interpretation Layer (AI & Collective Belief)

AI and probabilistic models evaluate submitted claims by linking admissible evidence, running Bayesian updates  $P(H|E)$ , and producing explanatory outputs with machine-verifiable references. These inputs generate a continuously updated Probability Index for each hypothesis, while volatility metrics capture both the uncertainty of evidence and shifts in collective belief. Interpretation thus integrates machine-based inference with human participation, Produces the Probability Index and volatility, reflecting both epistemic updates and shifts in collective confidence.

- Execution Layer (Markets & Resolution)

Bonding curve contracts, and settlement mechanisms enable adversarial pricing of truth claims. AI-driven resolution serves as the default path, automatically adjudicating outcomes against admissible evidence, while a verifier tribunal provides human-in-the-loop fallback for disputes

This modularity ensures that identity, data, interpretation, assets, and resolution remain composable and auditable, preventing any single point of epistemic failure.



**Figure 5. Episteme’s five-layer stack.** From settlement (Humanode L1), through scientific data synchronization (Knowledge Layer) and hypothesis tokenization (Asset Layer), to AI-driven interpretation (Interpretation Layer) resolution (Execution Layer), the stack ensures composable, auditable, and economically aligned truth production.

### 4.3.1 Base Layer: Humanode L1

As the underlying blockchain infrastructure, Episteme leverages Humanode for its unparalleled commitment to decentralization, fairness, and Sybil resistance in the Web3 ecosystem. Humanode operates on a cryptobiometric consensus mechanism, where each validator is anchored to a verified, living human being rather than capital holdings or computational power. This ensures that every participant in the network is unique, mitigating the risks of centralization, collusion, or manipulation.

Humanode's decentralization is empirically demonstrated by its Nakamoto Coefficient of 341, making it the most decentralized blockchain network in existence. For comparison, Mina achieves 151, Polkadot 92, Solana 18, and Binance only 7. This exceptional degree of validator distribution provides Episteme with a settlement layer that is highly resistant to capture and robust against adversarial interference.

This foundation directly supports Episteme's mission: to create a platform where scientific inquiry and prediction markets remain open, equitable, and resilient. By removing capital-weighted control, Humanode ensures that no powerful entity can distort epistemic outcomes. In future iterations, Episteme will extend these primitives to enable one-human-one-account citizen science and one-human-one-vote governance, further aligning the infrastructure of truth with democratic principles.

To ensure seamless integration and adoption of technology within existing protocols, Episteme benefits from full Ethereum Virtual Machine (EVM) compatibility. The Humanode network includes an EVM pallet that allows the execution of Solidity smart contracts and the use of existing developer tools, making it straightforward for Episteme to operate within Ethereum's ecosystem while taking advantage of Humanode's unique features. The EVM compatibility is powered by SputnikVM, which comprises four essential modules: evm, evm-core, evm-runtime, and evm-gasometer. These components allow the Humanode network, and by extension Episteme, to execute Ethereum-compatible contracts while addressing certain limitations of Ethereum, such as transaction fee volatility.

One of the critical innovations Humanode brings to Episteme's platform is a cost-based fee system that replaces the conventional gas-based system. This ensures that transaction fees remain stable in USD terms, despite the volatility of the underlying native token, eliminating the unpredictability commonly associated with DeFi platforms. This stability is crucial for users placing bets or trading scientific assets, ensuring they are not adversely affected by sudden fluctuations in token prices.

Mainly, Humanode functions not merely as Episteme's technical substrate, but as its ethical anchor, ensuring that the very base of the system reflects the same commitment to fairness and decentralization that the platform seeks to bring to science itself.

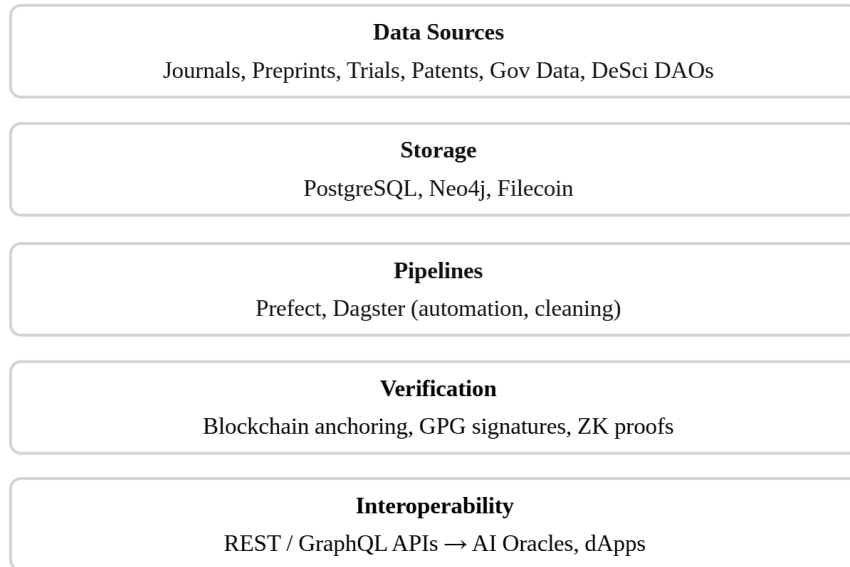
### 4.3.2 Knowledge Layer: Scientific Database and Provenance

The Knowledge Layer forms the epistemic substrate of Episteme, binding hypotheses to verifiable artifacts rather than to reputation or authority. At its core, it integrates heterogeneous sources of scientific knowledge into a unified database that is both digestible by AI and auditable by humans. It draws upon peer-reviewed journals and indexes such as PubMed, Scopus, and Crossref; preprint servers including arXiv, bioRxiv, and chemRxiv; clinical trial registries like ClinicalTrials.gov and the EU Clinical Trials Register; patent repositories such as USPTO, WIPO, and Google Patents; open scientific archives including Zenodo, Dryad, and Figshare; and large public datasets from agencies such as NASA, WHO, and the CDC. In early phases, additional contributions will come from Web3 research DAOs and institutional partners, ensuring that the database captures the full spectrum of contemporary scientific production, from preliminary preprints to regulated trial data.

These records are preserved within structured storage systems (PostgreSQL) and enriched by a knowledge graph architecture (Neo4j) that models relationships between hypotheses, researchers, institutions, and results. Such relational modeling enables claims to be situated within a wider epistemic network, revealing flows of evidence and influence that shape scientific discourse. Large raw datasets and supplementary materials are stored on decentralized systems (Filecoin), ensuring persistence, accessibility, and resistance to censorship. Automated data pipelines (Prefect, Dagster) orchestrate ingestion, cleaning, and integration across heterogeneous sources, while specialists intervene in the early phases to filter noise and prepare AI-readable corpora. Over time, this process evolves into a largely autonomous ingestion system, capable of sustaining itself without compromising accuracy.

To guarantee integrity, provenance is anchored on-chain through cryptographic mechanisms such as blockchain anchoring, GPG signatures, Merkle proofs, and zero-knowledge attestations. Each artifact, whether a preprint, dataset, or patent, is thus immutably tied to its origin, rendering tampering both detectable and economically irrational. This layer further exposes standardized APIs for interoperability, allowing external agents, oracles, and decentralized applications to query and interact with the database directly.

Through this architecture, the Knowledge Layer transforms fragmented, siloed scientific outputs into a transparent, auditable, and relationally rich substrate. It ensures that every hypothesis entering Episteme is rooted in verifiable evidence, structured for machine-assisted interpretation, and preserved within a trustless system of provenance.



**Figure 6. The Knowledge Layer stack** *integrates diverse data sources, structured storage, automated pipelines, cryptographic provenance, and interoperability into a unified epistemic substrate.*

### 4.3.3 Asset Layer: Tokenized Claims & Positions

The Asset Layer transforms hypotheses into verifiable, tradable primitives that form the core of Episteme’s epistemic markets. Anyone may submit a hypothesis, ensuring accessibility and broad participation, while researchers, institutions, and labs can attain verified status, adding reputational weight to their claims. Upon submission, each hypothesis is minted as a Hypothesis NFT, encoding canonical parameters such as wager type (binary, multinomial, or probabilistic), admissible data sources, timeframe, and resolution criteria. This formalization guarantees that hypotheses are born as standardized, machine-readable objects rather than ambiguous narratives.

To prevent duplication, AI algorithms cross-reference each new submission against the full database of active and historical hypotheses, checking for semantic and contextual similarity. Where closely related claims exist, a versioning system preserves their distinctiveness while linking them as parallel research tracks. This ensures uniqueness, transparency, and cumulative knowledge without redundancy.

Once minted, each NFT anchors its market through Position Tokens, circulating as “For” or “Against” stakes. Pricing is governed by bonding curves, which provide continuous liquidity and dynamic adjustment as market participation evolves. Through this design, the Asset Layer ensures that every hypothesis enters Episteme not as a static claim but as a living epistemic asset, unique, verifiable, parameterized, and positioned at the intersection of knowledge representation and market participation.

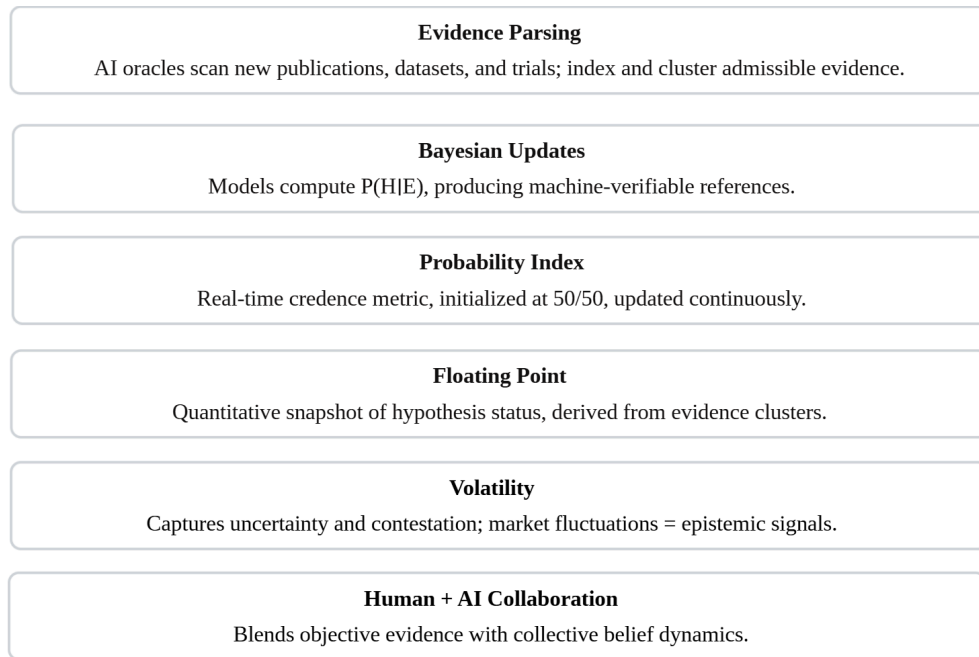
#### 4.3.4 Interpretation Layer: AI, Evidence, and Collective Belief

The Interpretation Layer transforms static tokenized hypotheses into living epistemic signals. Once a hypothesis is minted, it is continuously evaluated by AI oracles and probabilistic models that parse admissible evidence, update Bayesian probabilities, and generate machine-verifiable references. This process grounds every claim not only in human submission but in a dynamic interplay between evidence, markets, and collective judgment.

At the core of this layer lies the Probability Index, which represents the evolving credence assigned to each hypothesis. Markets initially open at 50/50 parity, establishing a neutral baseline, but the Probability Index quickly diverges as AI-driven updates and human participation accumulate. AI agents continuously scan new publications, datasets, and trial results, indexing relevant information and clustering it into coherent evidence streams. From these clusters, the system generates a floating point value—a quantitative snapshot of the hypothesis at a given moment, reflecting both the weight of the evidence and the credibility of its sources.

The Probability Index thus becomes a real-time barometer of epistemic reliability. Volatility within the index captures how contested or uncertain a claim is, signaling to participants whether a market is stable or in flux. What might traditionally be dismissed as noise is reframed here as an epistemic signal: sudden fluctuations reveal the arrival of new data, controversy in interpretation, or shifts in collective belief. In this sense, the Interpretation Layer integrates both objective evidence and subjective participation, blending machine inference with human wagers into a continuously updated portrait of scientific truth.

Through this architecture, interpretation is no longer a static peer review event but a perpetual process of updating, indexing, and contestation. Hypotheses are not merely minted and traded; they are monitored, revised, and reframed in real time, ensuring that science within Episteme is always alive to new data and open to challenge.



**Figure 7. The Interpretation Layer** *transforms hypotheses into living epistemic signals by integrating AI-driven updates, human participation, and volatility as a measure of uncertainty.*

### 4.3.5 Execution Layer

The Execution Layer operationalizes Episteme’s hypothesis markets, transforming tokenized claims into adversarially tested credibility signals. At this stage, economic incentives, AI-driven oracles, and human verifiers converge to enforce truth as both a market outcome and an epistemic process.

The core of this layer is resolution. By design, Episteme defaults to AI-driven resolution, in which oracles analyze admissible evidence, cross-check it against predefined criteria, and produce a draft outcome with citations. This ensures that most markets can be settled transparently, efficiently, and with reference to the scientific record. Where disputes arise, whether due to ambiguous evidence, contested methodology, or oracle uncertainty, a verifier tribunal of human experts provides a fallback. These verifiers adjudicate within a challenge window, issuing a Resolution Record that contains citations, reasoning, and accountable signatories. The result is a layered adjudication process that balances automation with human oversight.

To accommodate the inherent uncertainty of science, the Execution Layer also includes adaptive resolution mechanisms. If a wager reaches its time limit but the evidence remains inconclusive, the system may (1) distribute partial payouts based on the current Probability Index, (2) extend the wager until additional data becomes available, or (3) suspend the wager temporarily while preserving the ability to trade positions. In especially complex or contentious cases, community governance, including verified researchers and token

holders, may decide whether to extend, restructure, or terminate the market. These adaptive pathways ensure that execution remains fair, transparent, and resilient to the irregular rhythms of scientific discovery.

Participants are incentivized throughout. Traders earn returns for accurate forecasts; authors of hypotheses receive attribution and a share of fees; replicators are rewarded when their studies materially affect market credence; verifiers are compensated for accurate resolutions and risk slashing for negligence or collusion. These flows ensure that every role, forecasting, testing, curating, adjudicating, contributes to epistemic reliability. Exit mechanisms further support dynamic participation: users can sell or transfer their position tokens before resolution, allowing early profit-taking, risk management, or reallocation of capital as the Probability Index evolves.

Through this design, the Execution Layer closes the loop between knowledge and economics. It ensures that every hypothesis moves from submission to trading to resolution under transparent and adversarial conditions, producing outcomes that are financially enforced and epistemically grounded. In this way, Episteme replaces episodic peer review with living markets of credibility, where truth is continuously tested, priced, and contested in public view.

## 4.4 Token Ecosystem

Episteme’s token economy is designed as a modular, functionally separated system, where each token serves a distinct epistemic and economic role. This separation of functions prevents concentration of power, clarifies incentive flows, and ensures that markets operate transparently as engines of collective inquiry.

At the core lies *\$EPI (Episteme Token)*, an ERC-20 utility and governance token. *\$EPI* serves as the medium of exchange for market entry, transaction fees, and incentive distribution, while also anchoring participation in governance. Holders of *\$EPI* may vote on protocol upgrades, market standards, and dispute frameworks, making it both the economic and political substrate of the system.

Each hypothesis is instantiated as a *\$EWH (Episteme Wager Hypothesis)*, an ERC-721 non-fungible token that encodes canonical parameters—observable endpoints, admissible evidence streams, timeframe, and resolution criteria. By structuring hypotheses as NFTs, Episteme ensures uniqueness, auditability, and provenance, while enabling composability with other DeFi and DeSci infrastructures.

Belief in these hypotheses is expressed through *\$EWP (Episteme Wager Position)* tokens, also implemented as ERC-721 assets. Each *\$EWP* represents a staked position (“For” or “Against”) collateralized in *\$EPI*. Their price is governed by bonding curves, ensuring continuous liquidity and rational price discovery. In this way, *\$EWP* tokens embody epistemic judgments as tradable financial instruments, transforming credibility itself into a liquid, contestable signal.

Finally, the system operates atop \$eHMND (Humanode Token), the native token of Humanode L1, which functions as the gas token for settlement and consensus. Beyond transaction fees, Humanode’s cryptobiometric consensus guarantees that participation in Episteme is Sybil-resistant, anchoring the epistemic economy in human uniqueness rather than capital concentration.

Together, these four tokens constitute a layered architecture: \$EPI as governance and utility, \$EWH as hypothesis-shell, \$EWP as tradable epistemic stance, and \$eHMND as the Sybil-resistant settlement token. This modular separation not only secures functional clarity but also encodes a deeper principle: that the pursuit of truth requires distinct yet interlocking instruments (identity, claim, position, and value) woven into a coherent epistemic economy.

Table 3. Episteme Token Stack		
Token	Type	Role
\$EPI (Episteme Token)	ERC-20	Used to enter markets, pay fees, and vote in governance
\$EWH (Episteme Wager Hypothesis)	ERC-721	Non-fungible token representing a hypothesis (the market shell)
\$EWP (Episteme Wager Position)	ERC-721	Position tokens (“For” / “Against”) tied to each hypothesis market
\$eHMND (Humanode Token)	Native token	Gas fee token for all transactions.

**Table 3. Core functional token stack:** *\$EPI as utility/governance, \$EWH as hypothesis primitive, \$EWP as tradable epistemic stance, and \$eHMND as gas token.*

Finally, Episteme also acknowledges *\$EFI (Epistemic Future Intelligence)*, a cultural and agentic token linked to Efi, the protocol’s first AI agent and oracle persona. Unlike the core functional tokens, \$EFI anchors the social, cultural, and participatory dimension of Episteme. Efi publishes forecasts, interacts with communities, and embodies the platform’s identity as not merely an infrastructure but a living epistemic culture. As she evolves, \$EFI will secure her governance and unlock advanced features, rewarding holders and enabling participation in her growth. In this sense, \$EFI symbolizes Episteme’s recognition that legitimacy in science is not only technical but also narrative, cultural, and affective—a bold experiment in decentralized AI at the frontier of scientific engagement.

## 4.5 Hypothesis Markets

At the heart of Episteme lies the *Hypothesis Market*, a tradable contract on a scientific claim. In these markets, financial incentives, AI-driven updates, and collective intelligence

converge to transform passive observation into active epistemic participation. Where traditional science relies on episodic peer review, Episteme introduces a continuously priced system of credibility: each hypothesis is tokenized, monitored, and resolved under transparent, auditable rules.

### 4.5.1 Core Components

Four components distinguish Episteme’s hypothesis markets:

- *Probability Index*. A real-time metric of collective credence, expressed as the price of “For” versus “Against” positions. Functions as a public barometer of reliability.
- *Bonding Curve*. The mathematical rule governing token pricing as demand shifts. Variants such as LMSR or CFMM guarantee liquidity while bounding systemic risk.
- *AI Multiplier*. An adaptive mechanism whereby AI oracles adjust the Probability Index as new evidence streams arrive, operationalizing Bayesian inference at scale.
- *Resolution Oracles*. Hybrid modules (AI default + human fallback) that validate outcomes, generate rationales with citations, and trigger on-chain settlement.

Table 4. Hypothesis Market Core Components	
Component	Description
Probability Index	Reflects real-time confidence in the hypothesis outcome
Bonding Curve	Governs token pricing as demand changes
AI Multiplier	Continuously adjusts odds based on research feeds and Bayesian inference
Resolution Oracles	Validate outcomes and trigger settlement on-chain

**Table 4. Core components of Episteme’s Hypothesis Markets.** *Each market integrates four mechanisms: the Probability Index (collective confidence), the Bonding Curve (dynamic pricing), the AI Multiplier (continuous evidence-based updates), and Resolution Oracles (on-chain settlement). Together, these components transform hypotheses into liquid, continuously adjudicated signals of credibility.*

### 4.5.2 Lifecycle of a Hypothesis Market

The lifecycle of a hypothesis market mirrors the scientific method while embedding cryptoeconomic incentives.

#### 4.5.2.1 Creation

A researcher or participant specifies a claim, which is minted as a Hypothesis NFT (\$EWH). Each \$EWH encodes canonical parameters, observable endpoints, admissible data sources, timeframe, and resolution criteria, ensuring machine-readable, auditable

contracts rather than ambiguous narratives. Verified researchers and institutions may receive reputational badges, signaling credibility and anchoring trust. Duplicate hypotheses are prevented through AI cross-referencing and versioning.

#### 4.5.2.2 Activation

Once listed, the hypothesis opens to participation. Market actors stake belief through Position Tokens (\$EWP), denominated “For” or “Against” and collateralized in \$EPI. An initial Probability Index defaults to 50/50 unless otherwise adjusted by pre-analysis, providing a neutral baseline.

#### 4.5.2.3 Monitoring

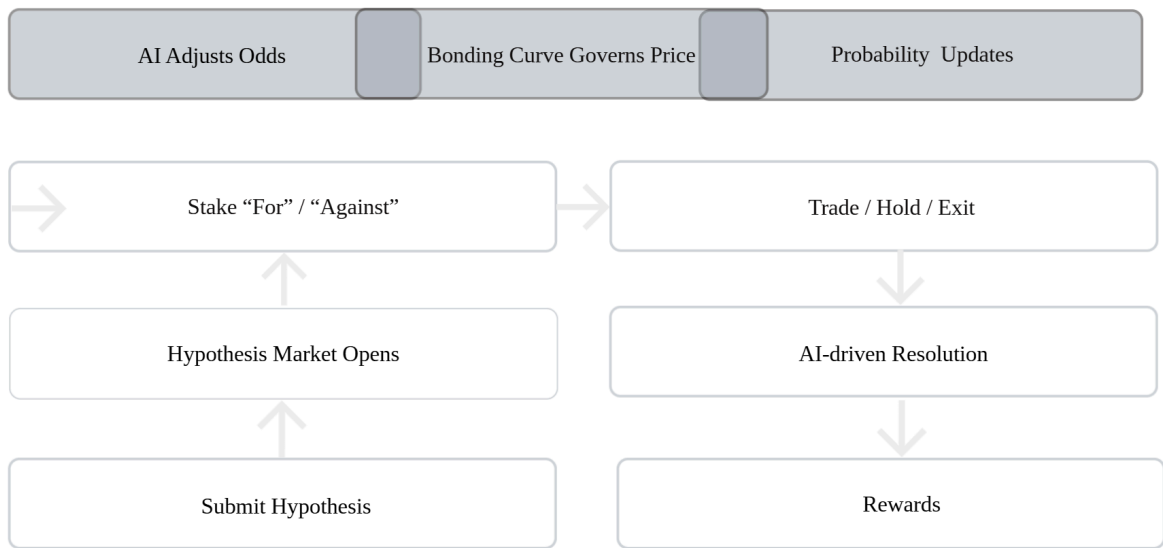
Throughout the lifecycle, AI oracles continuously ingest and cluster new literature, datasets, and experimental results. From these inputs, they generate Bayesian updates with machine-verifiable citations. The Probability Index adjusts accordingly, while volatility captures uncertainty and shifts in collective belief. This functionality aligns with Hanson’s (1995) thesis that markets can serve as institutional alternatives to static peer review.

#### 4.5.2.4 Trading

Participants may freely buy, sell, or hold \$EWP tokens over time. Bonding curves ensure liquidity and rational price discovery. Market volatility, far from being dismissed as noise, becomes an epistemic signal: the arrival of new data registers as price fluctuation, echoing Osband’s (2025) account of markets as rational learning systems. Participants retain flexible exit options, able to liquidate or transfer positions before final resolution.

#### 4.5.2.5 Resolution

At the close of a predefined timeframe, or upon meeting explicit criteria, the market proceeds to resolution. AI-driven oracles produce a draft outcome with citations, subject to challenge during a dispute window. In case of ambiguity, a verifier tribunal issues a Resolution Record with accountable signatories. Outcomes trigger settlement: winning positions are rewarded, authors and replicators receive attribution and bonuses, and verifiers earn fees or face slashing for negligence. Replication and falsification, often undervalued in traditional science (Fanelli, 2018), are incentivized as income-generating activities (Camerer et al., 2016).



**Figure 8. Lifecycle of a Hypothesis Market.** *A researcher submits a hypothesis, which is minted into a market. Participants stake “For” or “Against” positions, governed by bonding curves and dynamically adjusted by AI updates to odds and probabilities. Traders can buy, sell, or exit positions, while AI-driven resolution (with human fallback) determines outcomes and allocates rewards. This process transforms hypotheses into continuously priced, auditable signals of scientific credibility.*

### 4.5.3 Incentive Flows

Episteme’s incentive architecture is designed to align financial gain with epistemic reliability, ensuring that profitability coincides with improvements in accuracy, provenance, and replicability. Attempts to extract value without contributing to information or verification are structurally unprofitable, creating a market environment where truth-seeking is the only sustainable strategy.

#### 4.5.3.1 Authors

Hypothesis authors initiate epistemic activity by specifying claims in canonical form, with explicit endpoints, admissible evidence streams, and resolution criteria. They receive attribution and a continuing share of downstream market activity tied to their hypotheses. Precision and decision-relevance attract liquidity and engagement, while vague or ill-formed claims remain underfunded, making authorship itself an economic signal of quality.

#### 4.5.3.2 Forecasters

Forecasters stake “For” or “Against” positions, with rewards linked not to volume or popularity but to calibration and marginal information gain. Signals that shift credence toward truth, particularly those causing sharp updates (“credence shocks”), are privileged

over consensus-seeking behavior. This ensures that the market rewards epistemic novelty and correction rather than herding.

#### 4.5.3.3 Verifiers

Verifiers arbitrate disputes, review admissibility of evidence, and finalize resolution outcomes. Compensation is tied to accuracy, timeliness, and the quality of written reasoning, while negligence or collusion is penalized through slashing. Open logs and challenge windows keep their authority transparent and contestable, preventing verifier capture.

#### 4.5.3.4 Replicators

Replicators preregister protocols, conduct confirmatory or falsifying studies, and publish underlying data and code. When their results materially shift market credence, they are compensated according to the expected value of information (VOI). Replication thus becomes an income-generating activity rather than an underfunded academic duty, institutionalizing verification as an economic primitive.

#### 4.5.3.5 Curators & Indexers

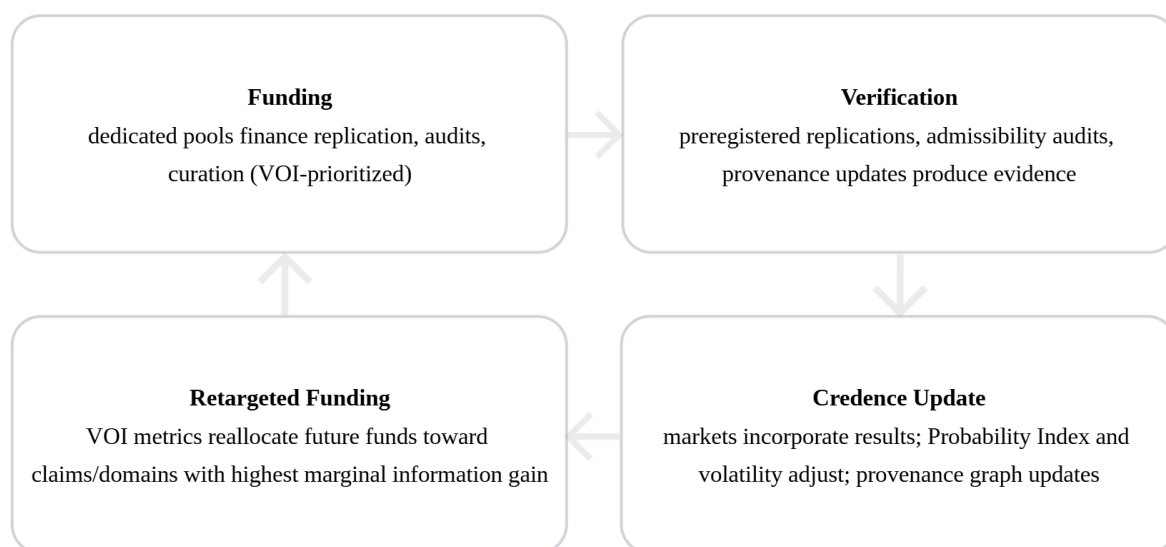
Curators maintain schemas, ontologies, and provenance graphs that structure the epistemic substrate. They are rewarded when improvements measurably reduce resolution error or accelerate credible settlement, ensuring that infrastructural refinement is directly tied to epistemic performance.

#### 4.5.3.6 The Incentive–Verification Loop

Trading activity funds replication, audits, and curation; completed verifications feed back into market prices and provenance records; updated credence then reallocates future funding by VOI. In this loop, volatility is not dismissed as noise but harnessed as a discovery signal, turning uncertainty into resource allocation toward domains where truth remains contested.

#### 4.5.3.7 Operational Metrics & Governance Triggers

To preserve robustness, the protocol continuously monitors resolution latency, dispute frequency, validator reversal rates, market depth, calibration distributions, VOI spend-to-impact ratios, and the integration of null or negative results. Pre-set thresholds on these indicators act as governance triggers: adaptive weighting of marginal information gain, temporary validator rotation, increased replication bounties in high-risk domains, or upgrades to dispute resolution rules. By embedding reflexivity into its governance, Episteme ensures that reliability improves as the system learns.



**Fig. 9. Incentive–Verification Loop: Funding → Verification → Credence Update → Retargeted Funding.** *The protocol converts market activity into financed verification and feeds verified results back into prices and provenance. Updated credence then redirects future funds by expected value of information, creating a self-correcting, reliability-maximizing cycle.*

#### 4.5.4 Anticipating Failure Modes

Episteme treats failure not as an aberration but as a normal operating condition of epistemic systems. Markets inevitably attract bias, collusion, and noise; robustness therefore requires counterweights built directly into the mechanism design. Premature convergence, or herding, can suppress minority but accurate signals, yet this risk is countered by rewarding marginal information gain and calibration over agreement, so that dissent which improves forecasts is recognized as economically valuable. The longstanding problem of cherry-picking and suppression of null results is addressed through provenance audits, preregistration, and explicit bounties for negative or refuting evidence, transforming falsification into a profitable activity.

Episteme also guards against topic glamour bias, in which fashionable domains attract disproportionate funding while socially vital but neglected fields remain under-resourced. By earmarking replication pools and employing quadratic matching guided by the expected value of information, resources are redistributed toward underserved areas where epistemic gains are highest. Manipulation and collusion, whether through coordinated trading or validator capture, are constrained through Humanode’s Sybil-resistant identity system, combined with staking, slashing, transparent dispute windows, and auditable reasoning trails that make collusion both costly and contestable.

Speculative volatility presents another danger, where noise masquerades as knowledge and short-term speculation displaces truth. Here, volatility is treated as an epistemic signal only when anchored to admissible evidence, and trading activity is explicitly routed into

replication and audits so that market motion remains disciplined by verification. Risks of metric gaming, consistent with Goodhart’s law, are addressed through multi-objective scoring that combines calibration, information gain, and resolution impact, supported by randomized audits and periodic metric rotation to resist over-optimization.

As automated agents enter markets, the threat of agent gaming and data poisoning grows. Episteme requires agents to abide by the same staking, slashing, and transparency rules as humans, with admissibility checks and cross-source consensus mitigating adversarial data. Long-horizon illiquidity is managed through staged milestones and interim evidence gates, which allow partial settlement and continuous verification funding, preventing slow-resolving hypotheses from falling dormant. To further reinforce integrity, verifiers must file conflict-of-interest attestations and undergo random spot audits, while reputation scores decay in the absence of accurate calls. A dedicated bounty pool guarantees compensation for null and disconfirming results, permanently pinned in provenance to preserve reusability.

The cumulative effect is an adversarial equilibrium: every distortion creates an opportunity for informed correction. By financing validation, rewarding novelty, and penalizing manipulation, Episteme ensures that reliability does not emerge in spite of contestation, but through it.

<b>Table 5. Failure Modes and Countermeasures in Episteme</b>		
<b>Failure Mode</b>	<b>Mechanism of Distortion</b>	<b>Countermeasure in Episteme</b>
Herding & Information Cascades	Premature convergence suppresses minority but accurate signals	Rewards weighted by marginal information gain; calibration prioritized over consensus
Cherry-Picking & Null Suppression	Selective reporting and underproduction of disconfirming evidence	Provenance audits; preregistration; explicit bounties for negative or falsifying results
Topic Glamour Bias	Overfunding fashionable domains; neglect of critical but low-attention fields	Domain-earmarked replication pools; quadratic matching guided by VOI
Manipulation & Collusion	Coordinated trading or verifier capture distorts settlement	Sybil resistance; staking and slashing; transparent dispute processes; open logs
Speculative Volatility	Noise masquerading as knowledge; short-term speculation dominates	Volatility treated as epistemic signal only when evidence-anchored; trading fees fund replication & audits

Metric Gaming (Goodhart's Law)	Optimizing visible metrics (volume, attention) over epistemic reliability	Multi-objective scoring (calibration + info gain + resolution impact); randomized audits; metric rotation
Agent Gaming & Data Poisoning	Automated agents amplify errors or seed adversarial data	Agents subject to staking/slashing; admissibility checks; cross-source consensus validation
Long-Horizon Illiquidity	Thin participation in slow-to-resolve claims	Staged milestones; interim evidence gates; partial settlements to sustain liquidity
Integrity Risks	Conflicts of interest, verifier stagnation, missing nulls	COI attestations; random spot audits; decaying reputations; dedicated bounty pool for null/disconfirming results

**Table 5. Anticipated failure modes in Episteme and their embedded countermeasures.** *Rather than treating bias, collusion, or volatility as anomalies, the protocol encodes them as normal operating conditions, ensuring that adversarial pressure fuels correction and reliability.*

#### 4.5.5 Epistemic Significance

The hypothesis market represents not just a funding tool but a reconfiguration of epistemic legitimacy. By making hypotheses tradable, adversarial, and continuously adjudicated, Episteme shifts science away from static publication toward living, contestable credibility signals. Markets provide three epistemic functions: aggregation of dispersed signals, updating in response to new evidence, and discipline through adversarial correction.

Empirical precedents support this vision. Dreber et al. (2015) showed that prediction markets accurately forecast replication outcomes in psychology; Camerer et al. (2016, 2018) replicated these results in economics; and Holzmeister et al. (2024) extended the approach to decision markets, where trader-driven prioritization improved replication success rates. Collectively, these studies demonstrate that markets can serve as epistemic engines, transforming volatility into discovery signals and aligning financial incentives with scientific rigor.

In this sense, Episteme realizes what Murphy (2023) calls a meta-episteme: a structural transformation of knowledge production in which truth is not decreed by journals but continuously priced, tested, and revised in public view.

## 5. Human–AI Epistemics: Adjudication & Contestability

Episteme treats AI not as an oracle of record but as an epistemic amplifier inside a transparent, adversarial, and ultimately human-accountable process. Large models synthesize corpora, propose probabilistic updates, and attach machine-verifiable citations; humans decide contested outcomes under explicit rules. The design objective is *explainable convergence*: the outcome that is both best supported and best explained prevails, rather than the one emitted with the highest confidence.

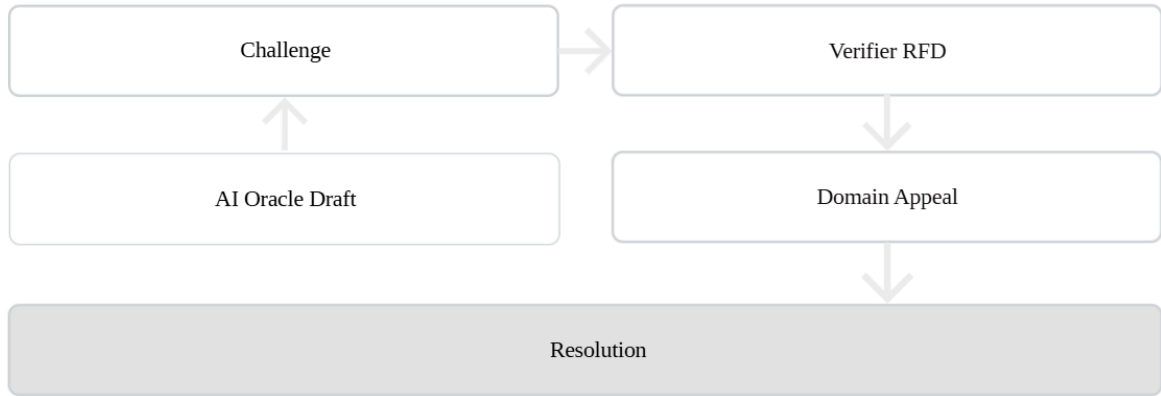
Unlike purely human peer review, which suffers from latency, bias, and opacity, Episteme’s hybrid loop provides rapid probabilistic drafts continuously exposed to challenge. Unlike AI-only adjudication, which risks hallucination and hidden confounders, Episteme embeds accountability, explanation, and human override. The result is a system that is faster than traditional science yet more trustworthy than unmediated automation.

### 5.1 Process Architecture

The adjudication process follows a structured sequence:

- *Oracle Draft*: AI oracles ingest literature, datasets, and code to propose a probabilistic draft with citations and uncertainty annotations.
- *Challenge Window*: Participants may dispute the draft by submitting counter-evidence, provenance objections, or alternative likelihood estimates.
- *Verifier Reasons for Decision (RFD)*: Verifiers review the draft and objections, issuing a Reason for Decision that explicitly addresses dissenting arguments.
- *Domain Appeal (optional)*: In specialized or high-stakes cases, appeals escalate to domain panels.
- *Resolution*: Final outcomes are recorded on-chain with transparent rationales and preserved minority opinions.

This staged architecture ensures that authority is exercised not through reputation or opaque processes but through reasons, evidence, and accountability.



**Figure 9. Resolution pipeline.** Oracle drafts provide probabilistic updates, challenge windows enable contestation, verifiers issue Reasons for Decision (RFD), and domain appeals refine context-specific cases before final resolution.

## 5.2 AI oracles as first pass, not final say

AI agents continuously ingest literature, code, datasets, and preregistrations to produce three artifacts: (i) structured claim digests and dependency graphs, (ii) proposed likelihood updates to  $P(H|E)P(H|\text{mid } E)$  with explicit premises, and (iii) draft resolution memos linking admissible evidence. Each artifact is schemed (assumptions, admissibility checks, exclusion reasons, and uncertainty drivers are enumerated) so that disagreement targets *specific* premises rather than opaque outputs, reflecting Alvin Goldman’s (1999) call for *social epistemics* grounded in structured testimony: reasons and evidence must be made explicit to enable critique.

## 5.3 Traceability by default

Every oracle action is logged: model family/version, weight provenance, prompt template, retrieval sources (anchored by content hashes), and emitted rationale. These records are content-addressed and anchored on-chain, enabling reproduction, comparison across model versions, and post-hoc error analysis. By exposing provenance, Episteme resists automation bias, ensuring that outputs are not taken at face value but evaluated through transparent traceability.

## 5.4 Designed contestability

Oracle drafts open a time-bounded challenge window. Any participant may (a) submit counter-evidence, (b) flag provenance defects (e.g., leakage, inadmissible data), or (c) propose an alternative likelihood with reasons. Verifiers are obligated to address salient objections in a formal RFD before signing. Failure to engage material objections is slashable. Challenge tooling highlights premise conflicts, missing artifacts, and citation mismatches, making contestation precise and productive.

## 5.5 Human accountability and layered review

Resolution proceeds through layers: oracle draft → public challenge → verifier panel → optional domain appeal. Signatories are accountable for their votes, and minority opinions are preserved with citations. Contestability ensures that authority is not monopolized by opaque models or elite verifiers but distributed across plural roles, reflecting Kitcher's (1993) model of *well-ordered science*.

## 5.6 Calibration and performance guarantees

Both oracles and verifiers are evaluated ex post. Proper scoring rules (e.g., Brier/log score) are tracked by domain and horizon; reversal rates on appeal, dispute density, and time-to-finality are monitored; contribution to resolution accuracy can be estimated via Shapley-style or counterfactual impact measures. Poorly calibrated models lose ensemble weight and fee share; consistently performant contributors gain routing weight. Forecasters and challengers who surface decisive counter-evidence are paid for marginal information gain, aligning incentives to scrutinize rather than rubber-stamp.

## 5.7 AI-specific failure modes and mitigations

To reduce brittleness and hallucination, Episteme constrains generation to admissible, licensed sources via retrieval-augmented generation; employs cross-model ensembling with disagreement detection and abstention; runs red-team prompts and holdout corpora to detect leakage and overfit; hardens against prompt injection and data poisoning in quarantined sandboxes; and applies domain checklists for recurrent statistical and causal errors. Where uncertainty is irreducible, the protocol prefers graded resolutions with explicit residual risk and follow-up evidence gates.

## 5.8 Normative boundary

AI may propose and summarize; humans decide contested outcomes. Prompts, versions, logs, and rationales are treated as public goods to enable critique and replication. Operationally: explanations outrank assertions. This boundary keeps the *Interpretation Layer* maximally efficient while preserving legitimacy at the *Execution Layer*, where decisions bear economic and epistemic consequence.

## 5.9 Agent ecology

AI agents within Episteme, including autonomous forecasters or curators, are treated as participants subject to the same incentives and penalties as humans. They must stake, face slashing for misconduct, and disclose provenance of their reasoning. This parity preserves fairness and ensures that epistemic legitimacy derives not from anthropocentric privilege but from transparent, auditable contribution.

Episteme measures the reliability of its Human–AI epistemics through ex post calibration, dispute resolution times, appeal frequency, and verifier reversal rates. Crossing pre-set thresholds (e.g., excessive error rates or unresolved disputes) triggers governance reviews and protocol adjustments. In this way, the system treats itself as a living hypothesis, an institution continuously tested and revised in line with its own epistemic commitments.

## 6. Governance: Polycentric Oversight

Episteme’s governance philosophy is grounded in the recognition that knowledge systems require not only epistemic but also political legitimacy. If truth is to be continuously adjudicated, the process of adjudication itself must be transparent, contestable, and resistant to capture. Following *polycentric theory* (Ostrom, 1990), Episteme distributes authority across multiple centers of decision-making rather than concentrating it in a single governing body. This plurality ensures resilience, prevents monopolization, and mirrors the distributed nature of the scientific enterprise itself.

### 6.1 Decentralization: From Core Team to Polycentric Order

Decentralization for Episteme is not an ornamental aspiration but the core ethos of the system. In science, as in finance, centralization tends to concentrate epistemic authority in a few hands—editorial boards, grant agencies, or corporate R&D divisions, limiting both transparency and participation. Episteme proposes instead that the governance of discovery itself be community-driven, transparent, and irreducibly plural.

The practical basis for this vision lies in its integration with Humanode’s Sybil-resistant identity layer, which ensures that one node corresponds to one unique human being. This allows Episteme to instantiate the most elementary democratic principle, one human, one voice (1h1v), without the distortions of plutocracy or the vulnerabilities of traditional pseudonymous systems. Within this framework, identity becomes a means of anchoring epistemic citizenship rather than a resource for exploitation.

Alongside identity-anchored legitimacy, Episteme introduces token-based governance. Governance tokens are distributed not arbitrarily but through participation: researchers, bettors, verifiers, and data providers accrue tokens by contributing to the epistemic economy. These tokens confer voting rights over operational parameters, system upgrades, and budgetary allocations. In this way, financial exposure is tied to epistemic responsibility, ensuring that those who benefit from the system also help steer it.

This dual structure—identity on one side, stake on the other—creates the conditions for what political theory would call polyarchy (Dahl, 1971): a pluralist regime in which no single modality of power dominates. Identity protects against plutocracy, stake prevents apathy by aligning incentives, and futarchy introduces a further layer of foresight and accountability.

The trajectory of decentralization is explicitly phased. In its earliest stage, core development teams necessarily retain a degree of custodianship, guiding system design and security. In the intermediate stage, authority is shared in a hybrid regime, with token holders and 1h1v participants taking increasing responsibility for strategic decisions. In the mature stage, Episteme converges toward full decentralization, with the DAO structure assuming primacy and the community collectively authoring the platform's trajectory.

## 6.2 Protocol Layer: A Constitutional Baseline

At the foundation, Episteme encodes a narrow set of *non-contestable minima*. These specify admissible endpoints (binary outcomes, graded probabilities, interval forecasts), provenance standards (cryptographic hashing, licensing, versioning), and due-process requirements (notice → challenge → reasoned decision → appeal). These are not subject to daily flux. Like a constitutional charter, they can only be amended through bicameral concurrence (requiring approval by both the OHOV chamber and the token house), public notice, and mandatory time-locks.

The analogy to constitutionalism is deliberate. Just as political constitutions anchor the stability of expectations in pluralist societies, the protocol baseline guarantees stability across epistemic domains. Without this, the project risks devolving into ad hoc rule-making vulnerable to capture and erosion of legitimacy.

## 6.3 Verifier DAO: Procedure as Living Case Law

The Verifier DAO functions as Episteme's judiciary. Verifiers, randomly assigned to prevent collusion, must disclose conflicts of interest and rotate regularly. Each decision must be justified by a written RFD: evidence considered, objections addressed, standards applied, and reasoning explained. Dissenting opinions are preserved, ensuring that minority views are not erased but remain part of the epistemic archive.

Over time, these decisions form a living body of case law. The jurisprudence is not static but evolves incrementally, much like common law, as precedents are cited, contested, and occasionally overturned. This mode of governance embodies Habermas's insight that legitimacy is generated not by outcomes alone but by the transparency of reasons and the fairness of procedure.

## 6.4 Domain Panels: Specialized Doctrine

Where hypotheses demand field-specific nuance (clinical endpoints, climate attribution, quantum thresholds) Episteme convenes rotating domain panels. These panels publish interpretive doctrines: what sources are admissible, what thresholds of sufficiency apply, and how to interpret edge cases. Their role is contextualization, not gatekeeping. Bound by constitutional minima and overseen by the Verifier DAO, they cannot override due process.

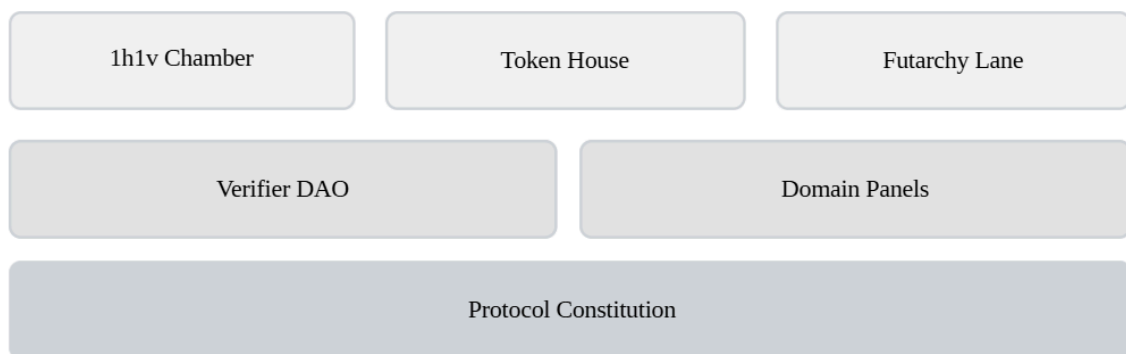
Panels are periodically audited against appeal outcomes, and persistent reversals trigger rotation. Expertise is thus included but never enthroned. As Michel Foucault observed in his analysis of regimes of truth, expertise is always ambivalent: a necessary resource, but one that must be continuously subjected to checks lest it consolidate into domination (Foucault, 1980).

## 6.5 Hybrid Modalities: Identity, Stake, and Foresight

Episteme’s governance stack interlocks three modalities:

- *One-Human-One-Vote (1h1v)*, instantiated via Humanode’s cryptobiometric system, anchors legitimacy in the equality of persons. It ensures that capture by wealth or capital is structurally impossible.
- *Token Governance* ties power to responsibility through stake, embodying the principle that those exposed to risk have an incentive to govern wisely.
- *Futarchy*, carefully bounded, channels distributed foresight by allowing markets to adjudicate which policies are most likely to advance epistemic KPIs.

The balance is delicate. 1h1v alone risks inefficiency and populism; token governance risks plutocracy; futarchy risks metric capture. But by constraining each within a polycentric order, Episteme designs a system in which their respective weaknesses are offset by one another’s strengths. Constitutional changes require 1h1v and token concurrence. Allocative questions are routed through futarchy, but only within 1h1v-defined KPI frameworks and token-capped budgets. Emergency powers are temporary, exercised by sortition-based safety councils, and sunset unless ratified bicamerally.



**Figure 10. Polycentric governance in Episteme.** *Constitutional rules anchor the system; Verifier DAO and Domain Panels ensure procedural and domain-specific oversight; and a hybrid stack of 1h1v (Humanode Vortex Chamber), Token House, and bounded futarchy provides adaptive, contestable, and forward-looking governance.*

### 6.5.1 1h1v House via Humanode Vortex (Identity)

Episteme implements 1h1v using Humanode’s Vortex governance, which guarantees one person = one vote through cryptobiometric uniqueness. We instantiate an Episteme Chamber within Vortex to deliberate:

- Constitutional items, ethics/fairness guardrails, and KPI selection for futarchy.
- Optional sub-chambers (Methods / Ethics / Infra) for scoped deliberation while preserving 1p1v.
- Quadratic voting may be enabled to express intensity without enabling majority domination.
- Public reasons for vote are recorded; cooling-off periods and quorum/supermajority rules apply to constitutional changes.

### 6.5.2 Token House (Stake)

The \$EPI chamber governs operational parameters and budgets within constitutional caps:

- Delegated voting (liquid democracy) with light re-delegation friction to deter flash swings.
- Commit-reveal on sensitive ballots; COI disclosures; slashing for bribery/collusion.
- Conviction voting is optional for continuous funding flows (e.g., maintenance grants).

### 6.5.3 Futarchy for Bounded Allocations (Foresight)

Finally, the system integrates Robin Hanson’s proposal of futarchy, “vote on values, bet on beliefs” (Hanson, 2000), but within carefully bounded domains. Where outcomes are measurable and values sufficiently shared, prediction markets can guide allocative decisions more effectively than committees. Yet futarchy’s known vulnerabilities—metric manipulation, liquidity distortions, and the temptation to outsource values to markets—are constrained by surrounding constitutional and procedural checks. The result is not wholesale adoption of futarchy but its domestication into a polycentric order where it complements rather than supplants other forms of legitimacy.

## 6.6 Reflexivity and Adaptivity

Unlike static constitutions, Episteme is reflexive. Governance performance is measured continuously: legitimacy through turnout and diversity indices, integrity through reversal rates and dispute outcomes, throughput through decision latency, and epistemic impact through changes in resolution reliability. Failures trigger adaptive responses: narrowing the scope of futarchy, convening deliberative mini-publics to break 1h1v deadlocks, adjusting quorum and delegation incentives when token governance shows concentration.

Governance is thus subjected to the same cycle of hypothesis, test, revision that defines science itself. It is epistemic not only in what it governs but in how it governs itself.

## 6.7 Risks and Mitigations

No governance system is immune to failure. The history of both political and scientific institutions shows that concentration of power, collusion among elites, overzealous correction, and ethical drift are recurrent hazards. Episteme's design acknowledges these risks and embeds countervailing mechanisms at multiple levels of its polycentric order.

### 6.7.1 Plutocracy and Capture

In token-based systems, there is always the danger that wealth translates directly into control. Token whales can concentrate voting power, undermine broad legitimacy, and bend epistemic priorities toward narrow financial interests. Episteme mitigates this by insisting that no constitutional change can be made without bicameral concurrence: both the 1h1v chamber (anchored in biometric personhood) and the token house must approve. This ensures that capital cannot override identity-based legitimacy. Further, stake caps and quadratic voting costs are applied to temper outsized influence, redistributing marginal power back toward the many.

### 6.7.2 Verifiers Cartels

A system reliant on procedural arbiters risks capture if verifiers collude or are consistently drawn from a narrow cohort. This can lead to predictable biases, rubber-stamp rulings, or corrupt bargains. Episteme defends against cartelization by randomized verifier rotation, mandatory conflict-of-interest disclosures, and the requirement that every ruling be accompanied by a publicly auditable RFD. Collusion or negligence is slashable, and repeated reversals in appeal processes trigger automatic removal. These measures create a jurisprudential culture where dissent is preserved and majority conformity cannot conceal procedural failure.

### 6.7.3 Over-Correction and Procedural Rigidity

Systems that are excessively responsive to every objection risk paralysis, while those that over-correct may swing wildly between rules, undermining stability. Episteme addresses this by bounding appellate mandates: appeals must give explicit reasons for reversal, and emergency powers carry automatic sunset clauses unless re-affirmed through bicameral approval. This ensures that interventions are proportionate and temporary, and that governance learns incrementally rather than oscillating between extremes.

### 6.7.4 Ethical Drift

Prediction markets are notoriously vulnerable to perverse incentives, such as the temptation to create markets that profit from harm. Without careful guardrails, epistemic wagering could incentivize exploitative or privacy-violating behavior. Episteme embeds ethical boundaries at the constitutional level: markets that target specific individuals, encourage harmful actions, or require data obtained through privacy violations are categorically

inadmissible. Sanctions are designed to be proportional and, where feasible, reversible, to avoid permanent exclusion for procedural missteps while maintaining deterrence for genuine abuses.

## 6.8 Governance as Truth-Making

The final point is philosophical. Episteme’s governance does not merely oversee truth-production; it is itself part of the truth-making apparatus. In Foucault’s terms, it constitutes a *regime of veridiction*: a set of rules through which statements acquire or lose the status of truth (Foucault, 1980). In Habermas’s sense, it embodies procedural legitimacy by ensuring that authority is exercised through transparent reasons. And in Hanson’s sense, it channels distributed foresight through prediction markets, while constraining their risks by embedding them in a broader polycentric architecture.

The result is a reflexive, adaptive, polycentric order in which governance and epistemology converge. To participate in Episteme is to participate not only in the contestation of hypotheses but in the continuous remaking of the rules by which truth itself is produced.

## 7. Revenue Flow & Treasury Policy

Episteme’s revenue model is designed not to extract rents but to sustain a self-funding epistemic commons. Instead of monetizing enclosure through paywalls or proprietary lock-in, value is generated whenever epistemic work occurs—when hypotheses are created, traded, verified, resolved, or analyzed—and is routed back into the very processes that make results reliable. In this sense, the protocol’s economy mirrors its epistemology: it is non-extractive, transparent, and continuously contestable.

### 7.1 Sources of Revenue


Revenue arises from diverse but complementary streams, each anchored in knowledge-production itself:

- *Market Activity*. Fees on minting and redemption of positions, secondary trading, and profitable exits. These are activity meters, accruing only when participants contribute information and calibrated risk.
- *Asset Lifecycle*. Minting, upgrading, or re-opening tokenized hypotheses incurs small lifecycle fees, discouraging spam and financing provenance stewardship.
- *Verification & Identity*. Researcher and institutional verification provide modest revenue while securing authorship and preventing Sybil capture.
- *Resolution & Disputes*. Oracle calls, verifier rulings, challenge bonds, and dispute arbitration generate fees proportionate to adjudication work. Penalties for frivolous challenges help underwrite a shared risk reserve.

- *Integrations Marketplace*. Shared revenue with third-party tools (replication labs, agent models, dashboards) plugged into Episteme’s schema and provenance rails.

## 7.2 Movement & Accounting

All fees are routed by smart contracts into named on-chain vaults: core treasury, replication/audit pool, verifier rewards, risk reserve, ecosystem grants. This separation prevents cross-subsidization, and every unit of value is traceable: from activity → vault → disbursement. Public dashboards expose balances and flows in real time, making accounting itself a form of veridiction.



**Figure 11. Revenue Flow in Episteme.** *Epistemic activity (hypothesis creation, trading, resolution, verification) generates fees that are programmatically routed into on-chain vaults. These vaults finance replication, audits, adjudication, and ecosystem grants, with disbursements governed by Value-of-Information metrics. Revenue thus circulates as a closed epistemic loop: inquiry produces value, value sustains verification, and verification renews trust in inquiry.*

## 7.3 Allocation & Accumulation

Episteme’s treasury is not a passive storehouse but a diversified epistemic engine. Revenue accumulates in multiple vaults:

- *Core Treasury* sustains operations, R&D, and protocol hardening.
- *Replication & Audit Pool* funds independent verification based on VOI.
- *Verifier & Curator Rewards* compensate for procedural work and metadata curation.
- *Risk Reserve* insures against rare failures (oracle error, verifier misconduct).
- *Ecosystem Grants* bootstrap standards, tools, and multilingual provenance infrastructure.

Allocations are adaptive: treasury governance continuously recalibrates spending toward domains and hypotheses where marginal epistemic gain is highest.

## 7.4 Distribution: Financing Reliability

Revenue is distributed not merely to cover throughput but to finance reliability.

- Protocol sustainability through maintenance, audits, and mechanism design.
- Verification-first outputs, where replications and audits receive priority based on VOI scoring.
- Participant rewards only when actions measurably improve calibration, provenance, or resolution quality.
- Integrity costs, replenishing reserves after disputes and funding anti-manipulation safeguards.
- Community & governance, including education, standards-setting, and participation tooling.

## 7.5 Policy Guards

Treasury policy is transparent, contestable, and metrics-driven:

- *Triggers*. Resolution latency, dispute reversals, calibration drift, and VOI impact ratios automatically schedule policy reviews.
- *Conflict-of-interest controls*. Verifier attestations, random audits, and reputation decay mitigate capture.
- *Equity*. Matching mechanisms and VOI earmarks direct resources to under-served fields, correcting glamour bias.
- *Agent parity*. AI agents operate under the same staking, slashing, and admissibility checks as human actors, preventing automated dominance.

## 7.6 Design Ethos

The design ethos is commons-preserving. Evidence remains open; what is monetized are the processes of coordination, validation, and interpretation. Analytics, services, and integrations generate surplus, but the epistemic artifacts themselves remain public goods. Thus, Episteme sustains a loop in which the act of inquiry funds its own verification.

# 8. Roadmap and Long-Term Vision

Episteme's trajectory is conceived as more than a sequence of software releases. Each phase is named after a classical epistemological concept, marking a philosophical journey from hypothesis to verifiable truth. This symbolic framing reflects the project's conviction that the evolution of technical infrastructure is inseparable from the evolution of epistemic culture.

## 8.1 Phased Rollout

### 8.1.1 Chōra (Internal Genesis, Q1 2025)

Borrowed from Plato’s *Timaeus*, the *chōra* is the receptacle of becoming. Episteme’s genesis phase lays the foundation: core contracts, AI oracle prototypes, and verifier logic are tested in controlled environments. This phase emphasizes robustness and interoperability rather than scale.

### 8.1.2 Eidōlon (Public Testnet, Q3 2025)

The *eidōlon* is an image, a first appearance. Here Episteme opens its interface to public trial. Users stake on mock hypotheses, experiment with trading mechanisms, and observe how AI-driven updates shift market credences. Though non-binding, this testnet is pedagogical: it familiarizes participants with the mechanics of epistemic pricing and allows early stress-testing under heterogeneity of engagement.

### 8.1.3 Krino (Adversarial Testnet, Q1 2026)

From the Greek *krinein*, to judge, Krino is the crucible for adversarial trials. Verifiers are randomly stressed, oracle pipelines face hostile data. The purpose is to surface weaknesses: validator cartels, manipulation attacks, liquidity shocks. Krino is where governance and procedural safeguards are tempered.

### 8.1.4 Aletheion (Mainnet Launch, 2026)

*Aletheia* signifies truth, disclosure. At this stage, Episteme becomes fully operational: real hypothesis markets go live, AI oracles resolve claims in production, verifiers generate jurisprudence, and incentives flow to forecasters, researchers, replicators, and curators. Episteme transitions from a speculative design to an epistemic economy.

Table 6. Episteme Roadmap (2025-2026)			
Phase	Type	Date	Milestones
Chōra	Internal Genesis	Q1 2025	Core components are built and tested.
Eidōlon	Public testnet	Q3 2025	Users explore UI, stake on mock hypotheses, test core features in a simulated environment.
Krino	Adversarial testnet	Q1 2026	Stress-test resolution logic, AI integrity, verifier DAO.
Aletheion	Mainnet launch	2026	Live hypothesis markets, AI validation layers activate, and incentives begin flowing to forecasters and researchers.

**Table 6. Episteme Roadmap 2025-2026.** *Four staged phases mark the protocol’s progression from foundational build (Chōra) through public experimentation (Eidōlon) and adversarial stress-testing (Krino) to full mainnet activation (Aletheion). Each phase symbolizes a step in the journey from hypothesis to verifiable truth.*

## 8.2 Long-Term Vision: Episteme 2030+

The horizon of Episteme extends beyond its mainnet. Its north star is the creation of an *Open Epistemic Layer* for humanity, a permissionless substrate where the production, funding, and verification of knowledge circulate with the same liquidity as digital assets today.

- *Global Knowledge Ledger*. Scientific claims, once opaque or siloed in paywalled journals, become transparently minted and tracked on-chain. AI agents synchronize global research outputs into active hypothesis markets, producing a continuously updated record of what is believed, at what probability, and why.
- *Economic Engine for Truth*. Accuracy becomes remunerative. Scientists, forecasters, and autonomous agents are rewarded not for prestige or novelty but for calibration and rigor. Replications and null results, long neglected, are financially supported through the very flows of market activity.
- *Autonomous Validation Loops*. As laboratories automate and instrumentation becomes networked, machine-driven research can feed directly into markets. IoT sensors, robotic labs, and simulation pipelines produce evidence streams that AI oracles ingest, evaluate, and route into resolution circuits. Validation becomes continuous rather than episodic.
- *Planet-Scale Scientific Commons*. Episteme lowers barriers to entry with micro-stakes and multilingual provenance. Participation is no longer confined to elite institutions but extended to citizen scientists, local research clusters, and distributed AI agents. The commons expands horizontally, integrating diverse geographies, languages, and epistemic traditions.

By 2030, the vision is of a *self-correcting epistemic infrastructure* where cycles of hypothesis, contestation, validation, and trust are accelerated, transparent, and collectively owned. Institutions of science need not be dismantled; Episteme overlays them with an open verification layer that redistributes incentives and restores reliability. The long-term wager is that truth, if properly incentivized, can become not only a philosophical aspiration but an economic reality.

<b>Vision</b>	<b>What It Enables</b>	<b>How Episteme Contributes</b>
Global Knowledge Ledger	Scientific claims minted & tracked on-chain	AI syncs global research ↔ active markets
Economic Engine for Truth	Scientists, forecasters, agents earn for accuracy	Rewards tied to validated outcomes & replication
Fair participation	Sybil-resistant citizen science	Humanode cryptobiometrics enforces one-human-one-account

Autonomous Validation Loops	Continuous machine-driven verification	IoT/lab agents stream results into markets
Planet-Scale Scientific Commons	Global participation as a shared public good	Micro-stakes, mobile UX, localization

**Table 7. Long-Term Vision of Episteme.** *Illustration of Episteme’s 2030+ horizon: a global knowledge ledger for transparent claims, an economic engine rewarding accuracy and replication, Sybil-resistant citizen science through biometric uniqueness, autonomous validation loops integrating machine agents, and a planet-scale scientific commons open to global participation.*

## 9. Discussion

### 9.1 Theory of Change

Episteme’s theory of change begins from the recognition that science suffers not only from funding bottlenecks and replication crises but also from a deeper misalignment of incentives. Activities that are epistemically essential—replication, careful audit, falsification, and transparent curation—are undervalued in the prestige economy of publications and citations. By contrast, novelty, volume, and visibility dominate, even when they do little to enhance reliability.

Episteme intervenes by transforming scientific hypotheses into tradable primitives. Once a claim is listed on-chain, it becomes the focal point of a continuous prediction market in which forecasters, researchers, and verifiers place calibrated wagers. In this design, pricing becomes more than a signal: it is a standing incentive to generate, contest, and verify evidence. Every trade is an epistemic bet that shifts market credence and exposes provenance to public scrutiny.

The feedback loop can be summarized as:

*Pricing → Verification Incentives → Provenance Exposure → Fee Recycling → Shortened Cycles → Trust → More Pricing.*

- *Pricing* of hypotheses creates liquidity around uncertainty, attracting participants to test claims.
- *Verification incentives* ensure that replication, falsification, and careful curation become income-generating rather than marginal activities.
- *Provenance exposure*—through content-addressed datasets, versioned code, and verifiable oracle drafts—shifts authority away from reputational hierarchies toward inspectable reasoning.
- *Fee recycling* routes a governed share of transaction activity into replication pools, verifier compensation, and ecosystem grants.

- *Shortened cycles* follow: claims are tested, revised, and, when necessary, retracted on faster timescales than traditional peer review allows.
- *Trust* builds cumulatively as reliability is demonstrated, not merely asserted.
- With higher trust comes more *pricing activity*, reinforcing the loop.

The intended outcome is an *antifragile epistemic commons* where reliability improves precisely through contestation and where economic flows reinforce rather than erode scientific rigor. By embedding accountability into financial and procedural structures, Episteme seeks to realign the economics of science with its epistemic purpose: the production of knowledge that is not only novel, but true.

## 9.2 Predictions and Evaluation

A credible infrastructure for truth must expose itself to falsification. Episteme can be tested against measurable predictions:

- *Latency*: replication timelines shorten relative to baselines.
- *Calibration*: market credences outperform expert panels on probabilistic accuracy.
- *Reproducibility lift*: funded replication rates rise in treated domains.
- *Integrity*: listed hypotheses show lower retraction or irreproducibility rates.
- *Diversity*: participation broadens across institutions, geographies, and career stages.

Evaluation uses a plural toolkit: scoring rules (Brier, log loss), quasi-experimental causal methods (difference-in-differences, synthetic controls), governance diagnostics (dispute/reversal rates, verifier dispersion), provenance completeness audits, and external validity checks (lag between claim listing and downstream adoption).

<b>Prediction</b>	<b>Metric</b>	<b>Methodology</b>	<b>Expected Effect</b>
<b>Latency</b> (time to replication)	Median days from claim listing → first high-quality replication	Benchmark against matched baselines; difference-in-differences across treated vs. untreated domains	Replication timelines shorten relative to controls
<b>Calibration</b> (forecast accuracy)	Brier score, log loss, calibration belts	Compare aggregated market credences to	Market probabilities outperform or match expert judgment

		expert panels on matched hypotheses	
<b>Reproducibility lift</b>	Share of hypotheses that receive funded replication	Synthetic control design; domain-level comparison of treated vs. untreated	Higher replication rates in domains with active markets
<b>Integrity</b> (claim reliability)	% of listed claims later retracted or found irreproducible	Longitudinal tracking of hypotheses listed $\geq 12$ months	Listed claims show lower retraction/irreproducibility rates
<b>Diversity</b> (participation base)	Distribution of institutions, geographies, and career stages (Gini/Simpson indices)	Contributor metadata; survey data; diversity audits	Broader and more equitable participation than journal-only pipelines
<b>Governance health</b>	Dispute frequency, appeal reversal rate, validator dispersion, time-to-finality	On-chain governance logs; statistical analysis of case law	Transparent due process; robust validator performance
<b>Provenance completeness</b>	% of hypotheses with machine-verifiable artifacts (datasets, code, licenses)	Automated content-addressed record audits; FAIR compliance checks	Higher share of reusable, reproducible research artifacts
<b>External validity</b>	Lag from claim listing $\rightarrow$ downstream uptake (clinical trials, policy, patents)	Event studies; correlation analysis with adoption timelines	Faster translation of validated claims into application

**Table 8. Evaluation Framework for Episteme** maps each prediction to its metric, methodology, and expected effect. This makes Episteme accountable to the same scientific norms it aims to transform.

## 9.3 Fairness and Epistemic Justice

Efficiency alone is insufficient; infrastructure must also serve epistemic justice. Episteme addresses equity through multiple channels:

- *Access*: low-cost, localized interfaces and translation layers broaden participation.
- *Attribution*: rewards flow not only to claimants but to replicators, curators, and data providers.
- *Distribution*: liquidity pools earmark funds for neglected domains and under-resourced regions.
- *Measurement*: routine bias audits track demographic and topical skew, with corrective levers pre-defined.

These mechanisms aim to counteract structural inequities in scientific production, aligning incentives with a normative commitment to science as a global public good.

## 9.4 Scope and Anti-Goals

Episteme is not designed as a wholesale replacement for journals, conferences, or theoretical exposition. It complements them by focusing narrowly on *claim-level credibility pricing*. The protocol excludes entertainment betting, celebrity speculation, and any markets incentivizing harm. It resists privatization of commons by keeping provenance trails open by default. Boundaries are as important as ambitions: defining anti-goals protects against mission creep and ethical drift.

## 9.5 Limitations and Open Questions

No epistemic infrastructure can escape its own limits. Episteme is both an intervention and an experiment, and its durability will depend on recognizing where fragilities lie. The following limitations are not afterthoughts but integral to the reflexive stance that a truth-making protocol must adopt: to submit itself, as much as the claims it hosts, to scrutiny, contestation, and possible revision.

### 9.5.1 Core design fragilities

#### 9.5.1.1 Speculation Versus Rigor

Financializing hypotheses introduces the danger of hype cycles. Highly visible or glamorous claims may attract liquidity far beyond their epistemic value, while less spectacular but socially critical work risks neglect. This tension mirrors the broader

economy of attention in science, where novelty often eclipses reliability. Episteme’s guardrails—VOI allocations, domain earmarks, replication bounties—may soften but cannot eliminate the gravitational pull of speculation.

#### 9.5.1.2 Futarchy Under Adversarial Pressure

Bounded futarchy offers a principled way to allocate resources, but its success rests on the stability of chosen KPIs. These metrics are vulnerable to Goodhart’s law: once targeted, they can be gamed. Adversarial actors may try to distort outcomes, while even well-intentioned participants may contest whether “resolution latency” or “replication rate” adequately measure scientific progress. The open question is whether futarchy can be bounded tightly enough to harness foresight without devolving into metric capture.

### 9.5.2 Ethical and epistemic dilemmas

#### 9.5.2.1 Ethical Conflicts in Research Participation

Decentralization complicates authorship and consent. What if one member of a research team lists a hypothesis on Episteme without notifying co-authors? What if collaborators interpret the scope of a shared project differently, staking on overlapping but misaligned endpoints? Questions of contributorship, attribution, and consent do not dissolve in on-chain provenance; they require governance rules and professional norms to ensure fairness. This is especially urgent when market incentives intersect with vulnerable populations or high-stakes biomedical claims.

#### 9.5.2.2 Complexity of Scientific Prediction

Not all science is easily compressible into binary or graded endpoints. Some phenomena resist tractable measurement, involve long causal chains, or depend on tacit expertise not reducible to datasets. Forcing precision onto inherently ambiguous problems risks producing false clarity, undermining trust in the very markets designed to strengthen it. Episteme must therefore remain humble about scope, recognizing that prediction markets illuminate certain classes of hypotheses better than others.

#### 9.5.2.3 Information Hazards and Dual Use

Some hypotheses, particularly in biotechnology or cybersecurity, may carry information hazards. A market predicting the success of a gain-of-function experiment or a vulnerability exploit could incentivize dangerous behaviors. Ex-ante domain exclusions, hazard review boards, and delayed disclosure protocols are necessary, but the ethical question remains: how to balance openness with responsibility when knowledge itself can be weaponized?

#### 9.5.2.4 Insider Information and Market Abuse

Researchers with access to non-public results could exploit information asymmetry for financial gain. While insider knowledge drives calibration in financial markets, in science it raises acute fairness and trust issues. Cooling-off periods, disclosure requirements, and slashing for undeclared conflicts may mitigate this risk, but cultural and legal expectations of research integrity complicate enforcement.

### 9.5.3 Cultural and institutional adoption barriers

#### 9.5.3.1 Cultural Adoption Barriers

Academic culture is cautious and slow to change. For many researchers, prediction markets may appear alien, reductive, or even ethically suspect. The fear of “financializing truth” could deter participation regardless of demonstrated calibration advantages. Legitimacy will thus depend not only on performance but also on careful bridge-building with journals, societies, and funders. Episteme may need to demonstrate complementarity—wrapping preprints, linking with replication initiatives—rather than framing itself as wholesale replacement.

#### 9.5.3.2 Governance Fatigue and Participation Inequality

Polycentric governance is complex. Participation costs (gas fees, time, cognitive load) may depress turnout, leading to apathy or capture by a narrow set of active delegates. Mini-publics, delegate scorecards, and incentive mechanisms may help, but the risk remains that governance becomes too intricate for broad engagement, eroding legitimacy over time.

#### 9.5.3.3 Legal and Institutional Constraints

Prediction markets occupy an uncertain regulatory space. Jurisdictions differ sharply in how they treat event contracts, and scientific claims may intersect with sensitive domains such as clinical trials or intellectual property. Similarly, hypotheses involving human subjects require ethical oversight akin to institutional review boards (IRBs). Episteme must navigate these constraints or risk exclusion from the very domains where reliable validation is most needed.

### 9.5.4 Technical and infrastructural fragilities

#### 9.5.4.1 Premature Automation in Validation Loops

Automation through AI oracles, robotic labs, and IoT sensors promises speed and scalability, but premature delegation could replicate the very opacity Episteme seeks to overcome. If models hallucinate or lab bots malfunction without transparent interpretability, automated validation could spread error rather than correct it. Incremental adoption, with human oversight preserved until reliability is demonstrated, is essential.

#### 9.5.4.2 Technical Fragilities

Episteme inherits risks common to blockchain systems: oracle poisoning, chain reorgs, validator key compromise, and non-stationarity of data. Zero-knowledge privacy, while powerful, can also obscure evidence needed for reproducibility. Compute and energy demands of large AI models and replication simulations may become prohibitive. Each of these risks requires design trade-offs: redundancy, checkpointing, selective disclosure, and sustainability metrics.

#### 9.5.4.3 Biometric Equity and Proof of Personhood

Humanode's biometric uniqueness provides Sybil resistance, but biometric systems themselves can reproduce exclusion. Variability across populations, accessibility barriers, and concerns about surveillance could marginalize contributors. Alternative proof-of-personhood mechanisms, independent audits, and strong privacy guarantees will be required to ensure that 1h1v is equitable rather than exclusionary.

#### 9.5.5 Open Research Questions

Episteme itself must be treated as a living hypothesis, subject to testing. Among the most pressing questions are:

- What mechanisms best guarantee epistemic justice across unequal geographies, languages, and institutional capacities?
- Can futarchy be bounded tightly enough to avoid KPI capture while still harnessing foresight?
- How will AI agents and human researchers coexist in markets without one eclipsing the other?
- What forms of meta-governance are required to adjust the weights of identity, stake, and foresight without paralysis?
- To what extent can market signals meaningfully influence public R&D funding policy without creating distortions of their own?
- How should authorship consent be structured when hypotheses are listed by one researcher but belong to collaborative projects?
- Can market incentives correct for the hype bias of science, or will they amplify it under liquidity pressure?

These limitations do not undermine Episteme's promise; they sharpen it. A protocol for truth cannot be naïve about its own fragilities. By treating each limitation as a standing research agenda—ethical, technical, cultural, and institutional—Episteme aspires to model the very reflexivity it demands of science: not perfection, but continuous, contestable improvement.

## 10. Conclusion

Science is not broken in its spirit but in its systems. The ideals of rigor, transparency, and collective inquiry remain intact, yet the infrastructures that sustain them—publication economies, incentive structures, funding channels—have become brittle, opaque, and misaligned. Episteme is proposed as one response: a protocol that treats epistemic reliability as a first-class good, embedding incentives for validation, contestation, and replication into the very mechanics of knowledge creation.

By transforming hypotheses into tradable primitives, Episteme reframes belief as a continuous wager. Markets do not replace peer review or scholarly interpretation; they complement them with a parallel infrastructure of calibration and accountability. AI oracles, verifier jurisprudence, and identity-anchored governance interlock to ensure that evidence remains contestable, provenance visible, and authority earned through reasons rather than incumbency.

The roadmap outlines a journey from genesis to live epistemic economy, while the long-term vision gestures toward an Open Epistemic Layer for humanity: a global ledger where scientific claims are minted transparently, tested continuously, and rewarded when verified. Such a commons does not abolish existing institutions but overlays them with a fabric of auditability, incentivizing rigor where prestige has too often privileged novelty.

Yet Episteme remains a hypothesis about science itself. Its design must be falsifiable: if markets do not shorten replication cycles, if calibration does not improve, if diversity and fairness do not increase, then the claim that this infrastructure enhances reliability must be revised or abandoned. By submitting itself to the same standards of contestability it demands of others, Episteme models a reflexive epistemology, one where governance, economics, and truth are co-produced.

The wager is bold: that by aligning incentives with epistemic quality, we can compress cycles of discovery, restore trust in science, and open participation to a planetary scale. Whether this wager succeeds is not a matter of rhetoric but of evidence. The challenge, and the invitation, is to test it together.

## 11. References

- Almenberg, J., Kittlitz, K., & Pfeiffer, T. (2009). An experiment on prediction markets in science. *PLoS ONE*, 4(12), e8500. <https://doi.org/10.1371/journal.pone.0008500>
- Almenberg, J., Kittlitz, K., & Pfeiffer, T. (2009). Experimental evidence of the wisdom of crowds in prediction markets. *Proceedings of the National Academy of Sciences*, 106(11), 456-459.
- Antweiler, W. (2012). Long-term prediction markets. *Journal of Prediction Markets*, 6(3).
- Ariely, D. (2008). *Predictably irrational: The hidden forces that shape our decisions*. HarperCollins.

- Arrow, K. J. (1962). Economic welfare and the allocation of resources for invention. In *The rate and direction of inventive activity: Economic and social factors* (pp. 609–626). Princeton University Press.
- Arrow, K., et al. (2008). ECONOMICS: The promise of prediction markets. *Science*, 320(5878), 877-878.
- Arthur, W. B. (1994). *Increasing returns and path dependence in the economy*. University of Michigan Press.
- Augur. (2019). The Augur whitepaper. Augur. Retrieved from <https://www.augur.net/>
- Baker, M. (2016). 1,500 scientists lift the lid on reproducibility. *Nature*, 533(7604), 452–454. <https://doi.org/10.1038/533452a>
- Balmaseda, M. A., Trenberth, K. E., & Källén, E. (2013). Distinctive climate signals in reanalysis of global ocean heat content. *Geophysical Research Letters*, 40(9), 1754-1759.
- Bamakan, S. M., Nezhadsistani, N., Bodaghi, O., & Qu, Q. (2022). Patents and intellectual property assets as non-fungible tokens; key technologies and challenges. *Scientific Reports*, 12(1), 1–13. <https://doi.org/10.1038/s41598-022-05920-6>
- Bartling, S., & Fecher, B. (2017). Blockchain for science and knowledge creation: Promises and potentials of distributed ledger technology in knowledge exchange. *Medium*. <https://doi.org/10.31235/osf.io/5ujd8>
- Begley, C. G., & Ellis, L. M. (2012). Drug development: Raise standards for preclinical cancer research. *Nature*, 483(7391), 531–533. <https://doi.org/10.1038/483531a>
- Björk, B.-C., Welling, P., Laakso, M., Majlender, P., Hedlund, T., & Guðnason, G. (2010). Open Access to the scientific journal literature: Situation 2009. *PLoS ONE*, 5(6), e11273.
- Benkler, Y. (2002). Coase's Penguin, or Linux and the nature of the firm. *The Yale Law Journal*, 112(3), 369–446.
- Bechhofer, S., Buchan, I., De Roure, D., Missier, P., Ainsworth, J., Bhagat, J., & Goble, C. (2013). Why linked data is not enough for scientists. *Future Generation Computer Systems*, 29(2), 599-611. <https://doi.org/10.1016/j.future.2011.08.004>
- Bell, M. J. (2006). Prediction markets for innovation: A research agenda. *Journal of Prediction Markets*, 1(1), 3-14.
- Bell, T. W. (2002). Gambling for the good, trading for the future: The legality of markets in science claims. *Chapman Law Review*.
- Bell, T. W. (2006). Prediction markets for promoting the progress of science and the useful arts. *George Mason Law Review*, 14, 37.
- Biagioli, M. (2002). From book censorship to academic peer review. *Emergences: Journal for the Study of Media & Composite Cultures*, 12(1), 11-45.
- Biddle, J. (2007). Lessons from the Vioxx debacle: What the privatization of science can teach us about social epistemology. *Social Epistemology*, 21(1), 21-39.
- Bischof, E., Botezatu, A., Jakimov, S., Suharenko, I., Ostrovski, A., Verbitsky, A., et al. (2022). Longevity foundation: Perspective on decentralized autonomous organizations for special-purpose financing. *IEEE Access*, 10, 33048–33058. <https://doi.org/10.1109/ACCESS.2022.3161392>
- Bohannon, J. (2015). Who's downloading pirated papers? Everyone. *Science*. <https://doi.org/10.1126/science.aab2181>
- Boumans, M., & Davis, J. B. (2010). *Economic methodology: Understanding economics as a science*. Palgrave Macmillan.
- Bordieu, P. (2004). *Science of science and reflexivity*. University of Chicago Press.
- Boudreau, K. J., & Lakhani, K. R. (2015). Open disclosure of innovations, incentives and follow-on reuse. Harvard Business School Working Paper.

- Borgman, C. L. (2015). *Big data, little data, no data: Scholarship in the networked world*. MIT Press.
- Brossard, D., & Lewenstein, B. (2010). A critical appraisal of models of public understanding of science: Using practice to inform theory. *Public Understanding of Science*, 19(3), 289–302.
- Brembs, B. (2018). Prestigious science journals struggle to reach even average reliability. *Frontiers in Human Neuroscience*, 12, 37. <https://doi.org/10.3389/fnhum.2018.00037>
- Buterin, V. (2014). *A next-generation smart contract and decentralized application platform*. Ethereum Whitepaper.
- Buterin, V. (2021). Moving beyond coin voting governance. Vitalik Buterin's Blog. <https://vitalik.ca/general/2021/08/16/voting3.html>
- Buterin, V., Hitzig, Z., & Weyl, E. G. (2019). Liberal radicalism: A flexible design for philanthropic matching funds. *Review of Behavioral Economics*, 6(1), 1–16.
- Buterin, V., Hitzig, Z., & Weyl, E. G. (2019). A flexible design for funding public goods. *Management Science*, 65(3).
- Camerer, C. F., Dreber, A., Holzmeister, F., Ho, T.-H., Huber, J., Johannesson, M., Kirchler, M., Nave, G., Nosek, B. A., Pfeiffer, T., Altmejd, A., Buttrick, N., Chan, T., Chen, Y., Forsell, E., Gampa, A., Heikensten, E., Hummer, L., Imai, T., ... Wu, H. (2018). Evaluating the replicability of social science experiments in *Nature* and *Science* between 2010 and 2015. *Nature Human Behaviour*, 2, 637–644. <https://doi.org/10.1038/s41562-018-0399-z>
- Carey, B. (2015). Psychologists welcome analysis casting doubt on their work. *The New York Times*. Retrieved from <http://www.nytimes.com/2015/08/29/science/psychologists-welcome-analysis-casting-doubt-on-their-work.html>
- Coase, R. H. (1937). The nature of the firm. *Economica*, 4(16), 386–405.
- Collins, H., & Pinch, T. (1993). *The golem: What everyone should know about science*. Cambridge University Press.
- Cook, J., Nuccitelli, D., et al. (2013). Quantifying the consensus on anthropogenic global warming in the scientific literature. *Environmental Research Letters*.
- Cook, J., Oreskes, N., et al. (2016). Consensus on consensus: A synthesis of consensus estimates on human-caused global warming. *Environmental Research Letters*, 11(4), 1–7.
- Cowtan, K., & Way, R. G. (2013). Coverage bias in the HadCRUT4 temperature series and its impact on recent temperature trends. *Quarterly Journal of the Royal Meteorological Society*.
- Dasgupta, P., & David, P. A. (1994). Toward a new economics of science. *Research Policy*, 23(5), 487–521.
- David, P. A. (2004). Understanding the emergence of 'open science' institutions: Functionalist economics in historical context. *Industrial and Corporate Change*, 13(4), 571–589.
- Davis P.M, Walters W.H. The impact of free access to the scientific literature: a review of recent research. *J Med Libr Assoc.* 2011 Jul;99(3):208–17. doi: 10.3163/1536-5050.99.3.008. PMID: 21753913; PMCID: PMC3133904.
- De Filippi, P., & Hassan, S. (2018). Blockchain technology as a regulatory technology: From code is law to law is code. *First Monday*, 23(12).
- DeFilippi, P., & Wright, A. (2018). *Blockchain and the law: The rule of code*. Harvard University Press.
- DeFrancesco, L., & Klevecz, A. (2022). Decentralized investor communities gain traction in biotech. *Nature Biotechnology*, 40(9), 1310–1315. <https://doi.org/10.1038/s41587-022-01459-z>

- Ding, W., Hou, J., Li, J., Guo, C., Qin, J., Kozma, R., & Wang, F. (2022). DeSci based on Web3 and DAO: A comprehensive overview and reference model. *IEEE Transactions on Computational Social Systems*, 9, 1563–1573. <https://doi.org/10.1109/TCSS.2022.3204745>
- Dreber, A., Pfeiffer, T., Almenberg, J., Isaksson, S., Wilson, B., Chen, Y., Nosek, B. A., & Johannesson, M. (2015). Using prediction markets to estimate the reproducibility of scientific research. *Proceedings of the National Academy of Sciences of the United States of America*, 112(50), 15343–15347.
- Douglas, H. (2009). *Science, policy, and the value-free ideal*. University of Pittsburgh Press.
- Dunwoody, S. (2014). Science journalism: Prospects in the digital age. *Public Understanding of Science*, 23(4), 518–533. <https://doi.org/10.1177/0963662514538339>
- Else, H. (2024). ‘Getting paid to review is justice’: Journal pays peer reviewers in cryptocurrency. *Nature*. <https://doi.org/10.1038/d41586-024-04027-4>
- Evans, D. S., & Schmalensee, R. (2016). *Matchmakers: The New Economics of Multisided Platforms*. Harvard Business Review Press.
- Fang, F. C., & Casadevall, A. (2016). Research funding: The case for a modified lottery. *mBio*, 7(2), e00422-16. <https://doi.org/10.1128/mBio.00422-16>
- Fang, F. C., & Casadevall, A. (2011). Retracted science and the retraction index. *Infection and Immunity*, 79(10), 3855–3859.
- Fang, F. C., & Casadevall, A. (2012). Reforming science: Structural reforms. *Infection and Immunity*, 80(3), 897–901.
- Fantaccini, S., Grassi, L., & Rampoldi, A. (2024). The potential of DAOs for funding and collaborative development in the life sciences. *Nature Biotechnology*, 42(4), 555–562. <https://doi.org/10.1038/s41587-024-02189-0>
- Fanelli, D. (2018). Opinion: Is science really facing a reproducibility crisis, and do we need it to? *Proceedings of the National Academy of Sciences*, 115(11), 2628–2631. <https://doi.org/10.1073/pnas.1708272114>
- Fitzpatrick, K. (2011). *Planned Obsolescence: Publishing, Technology, and the Future of the Academy*. NYU Press.
- Franco, A., Malhotra, N., & Simonovits, G. (2014). Publication bias in the social sciences: Unlocking the file drawer. *Science*, 345(6203), 1502–1505.
- Franzoni, C., & Sauermann, H. (2014). Crowd science: The organization of scientific research in open collaborative projects. *Research Policy*, 43(1), 1–20.
- Forsythe, R., et al. (1992). Anatomy of an experimental political stock market. *The American Economic Review*, 1142–1161.
- Foucault, M. (1975). *Discipline and punish: The birth of the prison*. Pantheon Books.
- Gibbons, M. (1999). Science's new social contract with society. *Nature*, 402(S6761), C81–C84. <https://doi.org/10.1038/35011576>
- Godin, B. (2006). The Linear Model of Innovation: The Historical Construction of an Analytical Framework. *Science, Technology, & Human Values*, 31(6), 639–667.
- Glänzel, W., & Schubert, A. (2004). Analyzing scientific research output. *Handbook of Quantitative Science and Technology Research*.
- Guston, D. H. (2000). Retiring the social contract for science. In *Science and Technology in Society: From the Policy Perspective to the Cultural Meaning of Science and Technology* (pp. 75–94).
- Hanson, R. (1995). Could gambling save science? Encouraging an honest consensus. *Social Epistemology*, 9(1), 3–33. <https://doi.org/10.1080/02691729508578768>
- Hanson, R. (2003). Combinatorial information market design. *Information Systems Frontiers*, 5(1), 107–119. <https://doi.org/10.1023/A:1022046409482>
- Hardin, G. (1968). The tragedy of the commons. *Science*, 162(3859), 1243–1248.

- Hayek, F. A. (1945). The use of knowledge in society. *The American Economic Review*, 35(4), 519–530.
- Heidegger, M. (1962). *Being and Time*. Trans. J. Macquarrie & E. Robinson. Harper & Row. (Original work published 1927).
- Hsu, E. (2011). Prediction markets for science. *Journal of Economic Behavior and Organization*, 80(1), 105-117. <https://ir.law.fsu.edu/articles/497/>
- Hurwicz, L. (1972). On Informationally Decentralized Systems. In *Decision and Organization*. North-Holland.
- Ioannidis, J. P. A. (2005). Why most published research findings are false. *PLoS Medicine*, 2(8), e124.
- Jacobs, P. (2014). It ain't just the heat, it's the humanity: Evidence and implications of a knowledge-based consensus on climate change. AGU Science Policy Conference.
- James, S. (2021). Proof of humanity: An explainer on how to register for Proof of Humanity and use the platform. Kleros Blog.
- Jonas, H. (1979). *The imperative of responsibility: In search of an ethics for the technological age*. University of Chicago Press.
- Kahneman, D., & Tversky, A. (1979). Prospect Theory: An Analysis of Decision Under Risk. *Econometrica*, 47(2), 263–292.
- Kant, I. (1784). *An Answer to the Question: What Is Enlightenment?*
- Kavazi, D., Smirnov, V., Shilina, S., Li, M., Contreras, R., Gajera, H., & Lavrenov, D. (2021). You are [not] a bot. arXiv preprint arXiv:2111.13189. <https://arxiv.org/abs/2111.13189>
- Kitcher, P. (1993). *The advancement of science: Science without legend, objectivity without illusions*. Oxford University Press.
- Kogut, B., & Metiu, A. (2001). Open-Source Software Development and Distributed Innovation. In *The Oxford Handbook of Group Creativity and Innovation*. Oxford University Press.
- Kronick, D. A. (1990). Peer review in 18th-century scientific journalism. *JAMA*, 263(10), 1321–1322.
- Kuhn, T. S. (1962). *The structure of scientific revolutions*. University of Chicago Press.
- Larivière, V., Haustein, S., & Mongeon, P. (2015). The oligopoly of academic publishers in the digital era. *PLoS ONE*, 10(6), e0127502. <https://doi.org/10.1371/journal.pone.0127502>
- Leonelli, S. (2016). *Data-centric biology: A philosophical study*. University of Chicago Press. <https://doi.org/10.7208/chicago/9780226416502.001.0001>
- Longino, H. E. (1994). Knowledge in social theories of science. In *Socializing epistemology* (pp. 135-157).
- Maskin, E. (1999). Nash equilibrium and welfare optimality. *Review of Economic Studies*, 66(1), 23–38. <https://doi.org/10.1111/1467-937X.00075>
- Marx, K. (1976). *Capital: A Critique of Political Economy* (Vol. 1). Penguin Books. (Original work published 1867).
- McKiernan, E. C., Bourne, P. E., Brown, C. T., Buck, S., Kenall, A., Lin, J., et al. (2016). How open science helps researchers succeed. *Elife* 5, e16800. doi:10.7554/eLife.16800
- Merton, R. K. (1973). *The sociology of science: Theoretical and empirical investigations*. University of Chicago Press.
- Miller, B. (2012). When is consensus knowledge based? Distinguishing shared knowledge from mere agreement. *Synthese*, 190(7), 1293-1316.
- Munafò, M. R., Pfeiffer, T., Altmejd, A., Heikensten, E., Almenberg, J., Bird, A., Chen, Y., Wilson, B., Johannesson, M., & Dreber, A. (2015). Using prediction markets to forecast research evaluations. *Royal Society Open Science*, 2(10), 150287. <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4632515/>

- Nakamoto, S. (2008). Bitcoin: A peer-to-peer electronic cash system. Bitcoin White Paper.
- Nelson, R. R. (1959). The simple economics of basic scientific research. *Journal of Political Economy*, 67(2), 297–306.
- Nosek, B. A., Spies, J. R., & Motyl, M. (2015). Scientific utopia II: restructuring incentives and practices to promote truth over publishability. *Perspectives on Psychological Science*, 10(6), 643–648.
- North, D. C. (1990). *Institutions, Institutional Change and Economic Performance*. Cambridge University Press.
- Open Science Collaboration. (2015). Estimating the reproducibility of psychological science. *Science*, 349(6251), aac4716. <https://doi.org/10.1126/science.aac4716>
- Ostrom, E. (1990). *Governing the commons: The evolution of institutions for collective action*. Cambridge University Press.
- Ohlhaber, P., Weyl, E. G., & Buterin, V. (2022). Decentralized society: Finding Web3's soul. SSRN. <https://doi.org/10.2139/ssrn.4105763>
- Piwowar, H., Priem, J., Orr, R., et al. (2018). The state of OA: A large-scale analysis of the prevalence and impact of Open Access articles. *PeerJ*, 6, e4375.
- Pew Research Center. (2020). Public trust in scientists. <https://www.pewresearch.org/science/2020/09/29/public-trust-in-scientists/>
- Powell, K. (2016). Young, talented and fed-up: Scientists tell their stories. *Nature*, 538(7624), 446–449.
- Polanyi, M. (1962). The republic of science: Its political and economic theory. *Minerva*, 1(1), 54–73.
- Popper, K. (1959). *The Logic of Scientific Discovery*. Hutchinson.
- Potthoff, M. (2007). The potential of prediction markets in science. *Futures*, 39(1), 45–53.
- Powell, W. W., & Grodal, S. (2005). Networks of Innovators. In J. Fagerberg, D. C. Mowery, & R. R. Nelson (Eds.), *The Oxford Handbook of Innovation* (pp. 56–85). Oxford University Press.
- Pfeiffer, T., & Almenberg, J. (2015). Prediction markets for science: Is the hype justified? *Nature*, 526(7575), 179–182.
- Prinz, F., Schlange, T., & Asadullah, K. (2011). Believe it or not: How much can we rely on published data on potential drug targets? *Nature Reviews Drug Discovery*, 10(9), 712.
- Romero, F. (2019). Philosophy of science and reproducibility crisis. *Synthese*, 196(8), 3067–3086.
- Sahlins, M. (1972). *Stone Age Economics*. Aldine de Gruyter.
- Samuelson, P. A. (1954). The Pure Theory of Public Expenditure. *The Review of Economics and Statistics*, 36(4), 387–389.
- Sen, A. (1999). *Development as Freedom*. Oxford University Press.
- Smith, A. (1976). *An Inquiry into the Nature and Causes of the Wealth of Nations*. University of Chicago Press. (Original work published 1776)
- Smith, R. (2006). Peer review: a flawed process at the heart of science and journals. *Journal of the Royal Society of Medicine*, 99(4), 178–182.
- Simon, H. A. (1957). *Models of Man: Social and Rational*. Wiley.
- Shilina, S. (2023). Decentralized science (DeSci): Web3-mediated future of science. *Paradigm Research*. Retrieved from <https://medium.com/paradigm-research/decentralized-science-desci-web3-mediated-future-of-science-2547f9a88c40>
- Shilina, S. (2023). Intellectual property (in science): The potential & advantages of NFTs [Preprint]. <https://doi.org/10.13140/RG.2.2.10291.50729/1>
- Stigler, G. J. (1961). The economics of information. *Journal of Political Economy*, 69(3), 213–225.

- Stephan, P. E. (2012). *How Economics Shapes Science*. Harvard University Press.
- Stokes, D. E. (1997). *Pasteur's Quadrant: Basic Science and Technological Innovation*. Brookings Institution Press.
- Swan, M. (2015). *Blockchain: Blueprint for a New Economy*. O'Reilly Media.
- Tirole, J. (1996). A Theory of Collective Reputations (with Applications to the Persistence of Corruption and to Firm Quality). *The Review of Economic Studies*, 63(1), 1–22.
- Thicke, M. (2016). The limits of prediction markets for scientific consensus. *Studies in History and Philosophy of Science*, 58(1), 50-58. <https://philpapers.org/rec/THIPMF-2>
- Trovò, B., & Massari, N. (2021). Ants-Review: A protocol for incentivized open peer-reviews on Ethereum. ArXiv. [https://doi.org/10.1007/978-3-030-71593-9\\_2](https://doi.org/10.1007/978-3-030-71593-9_2)
- UNESCO. (2021). *UNESCO Science Report: The Race Against Time for Smarter Development*.
- Unfried, M. (2024). Advancing longevity research through decentralized science. *Frontiers in Aging*, 5, 1353272. <https://doi.org/10.3389/fragi.2024.1353272>
- Van Noorden, R. (2014). Global research funding: What gets cut? *Nature News*, 505(7485), 618-619. <https://doi.org/10.1038/505618a>
- Voshmgir, S., & Zargham, M. (2020). *Foundations of Cryptoeconomic Systems*. WU Vienna University of Economics and Business. Working Paper Series / Institute for Cryptoeconomics / Interdisciplinary Research No. 1 <https://doi.org/10.57938/10d8b646-3266-4ec9-b94d-810b528c40a9>
- Walras, L. (1954). *Elements of Pure Economics*. Trans. W. Jaffé. Allen & Unwin. (Original work published 1874).
- Wang, J., Zhang, M., & Wei, C. (2021). Incentive Mechanism Design for Blockchain-based Systems. *IEEE Transactions on Network Science and Engineering*, 8(2).
- Wang, Fei-Yue & Ding, Wenwen & Wang, Xiao & Garibaldi, Jon & Teng, Siyu & Imre, Rudas & Olaverri-Monreal, Cristina. (2022). The DAO to DeSci: AI for Free, Fair, and Responsibility Sensitive Sciences. *IEEE Intelligent Systems*. 37. 16–22. 10.1109/MIS.2022.3167070.
- Wang, Fei-Yue. (2022). The Metaverse of Mind: Perspectives on DeSci for DeEco and DeSoc. *IEEE/CAA Journal of Automatica Sinica*. 9. 2043–2046. 10.1109/JAS.2022.106106.
- Weber, S. (2004). *The Success of Open Source*. Harvard University Press.
- Weiss, K. (2022). *Bitcoin grants – Quadratic funding for the world*. Bitcoin. <https://gitcoin.co>
- Wennerås, C., & Wold, A. (1997). Nepotism and sexism in peer-review. *Nature*, 387, 341–343.
- Weyl, E. G. & Buterin, V. (2019). *Liberation through radical decentralization*.
- Whitehead, Alfred North. [1929]1978. *Process and Reality*. Corrected edition, edited by David Ray Griffin and Donald W. Sherburne. New York: Free Press.
- Williams, S., Diordiiev, V., Berman, L., Raybould, I., & Uemlianin, I. (2018). *Arweave: A protocol for economically sustainable information permanence*. Arweave White Paper.
- Williamson, O. E. (1985). *The Economic Institutions of Capitalism*. Free Press.
- Wolfers, J., & Zitzewitz, E. (2004). Prediction markets. *Journal of Economic Perspectives*, 18(2), 107-126.
- Yadav, S. P., Agrawal, K. K., Bhati, B. S., Al-Turjman, F., & Mostarda, L. (2022). Blockchain-based cryptocurrency regulation: An overview. *Computational Economics*, 59(4), 1659-1675.
- Zhang, R., Xue, R., & Liu, L. (2021). Security and privacy on blockchain. *ACM Computing Surveys*, 53(1), 1-35. <https://doi.org/10.1145/3326067>