International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 11, CONFLUENCE - 2026

Beyond Conventional AI: Generative Intelligence, Adaptive Systems, and Digital Economies

A Context-Aware Multimodal Legal Information Framework and a Case Study of Adaptive AI Bosses and Blockchain Tokenomics in MIR5

Part A: Generative, context-aware, multimodal legal information support

Part B: Adaptive AI boss (Asterion) and WEMIX-based tokenomics of MIR5

Author: BHARATH V & Co-Author: SYAMESH

Department: Computer Applications

Institution: SVR College of Commerce and Management Studies

Location: Bengaluru, India

Corresponding Author E-mail: bharathv2006@gmail.com

Co-Author E-mail: 20.syamesh.06@gmail.com

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 11, CONFLUENCE - 2026

Abstract - Intelligent systems are increasingly shaping decisions in very different areas — from people looking up legal information to the digital economies that power online games. In this paper we ask how the outputs of AI systems can stay grounded in real context, stay transparent, stay properly bounded, and stay open to human checking. We explore the question through two complementary cases: a generative, context-aware, multimodal framework for legal information support, and a desk-based analysis of adaptive AI and tokenomics in the game MIR5.

Part A looks at an implemented framework built around Indian constitutional and criminal-law material — the Constitution of India, the Indian Penal Code (IPC) and the Bharatiya Nyaya Sanhita (BNS). We describe how information flows through the system, how it handles contextual queries, how it supports multimodal interaction, and the safeguards we put in place for responsible use. The evaluation so far is only preliminary and rests on informal functional testing; we make no claims based on formal benchmarks or user studies.

Part B examines MIR5 through the features that have been described publicly: the announced MMORPG, the adaptive AI boss Asterion, and a planned blockchain-based economy. We label every technical claim as a confirmed fact, a vendor claim, a reasoned inference, or something still undisclosed. An illustrative simulation that uses invented parameters shows how boss difficulty, player success, reward distribution and token emission might be linked — without claiming to represent any live-game measurements.

Looking across both cases, a shared design pattern becomes clear: keep adaptation constrained, make the status of the evidence visible, document uncertainty, and keep meaningful human oversight in place. The main risk is not the capability of the models themselves, but the chance that readers will treat unmeasured, inferred or vendor-described behaviour as if it were established fact. Our findings remain provisional — legal-system benchmarking is still incomplete and the implementation details of MIR5 are not public. What the paper offers, therefore, is a comparative framework for analysing responsible AI, not a claim that performance has already been demonstrated.

Keywords: Legal Information Support; Generative AI; Natural Language Processing; Multimodal AI; Adaptive Game AI; NVIDIA ACE; Blockchain Tokenomics; Responsible AI; Human–AI Collaboration.

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 11, CONFLUENCE - 2026

PART A — Generative, Context-Aware Multimodal Legal Information Support

A.1 Introduction

Legal information is formally public, yet in practice it is hard to use. Citizens still have to find the right provision, decode specialised wording, and connect it to their own situation. Generative models can restate that text in plain language, but they also produce fluent errors. General-purpose models have been shown to hallucinate in 58%–88% of verifiable questions about US federal cases [1], and even retrieval-based commercial legal tools still do so in 17%–33% of queries [2]. People who receive AI assistance also often accept wrong suggestions, even when explanations are provided [3]. The law itself also changes. In India, the Bharatiya Nyaya Sanhita (BNS) replaced the Indian Penal Code (IPC) from 1 July 2024 [4], which means a system built on older material may end up quoting provisions that no longer apply.

This part asks how generative AI and natural language processing (NLP), combined with structured legal content and conversational context, can make legal information more understandable while still keeping the answers grounded and verifiable. We describe an implemented framework, we state clearly what is and is not yet evidenced, and we propose an evaluation protocol. Throughout, the system is intended to give information, not legal advice.

A.2 Related Work and Research Gap

This targeted review is not systematic, so any conclusions we draw about research gaps are necessarily bounded (Table 1). Retrieval-augmented generation (RAG) joins a generative model to an external memory so that the outputs can draw on retrieved, updatable evidence [5]. In the legal domain, however, models often accept a user's false premise [1] and even retrieval-based tools still make errors [2]. Grounding is therefore necessary, but it is not sufficient on its own. Human reliance is a second concern. In one study that forced 199 participants to engage with AI explanations, over-reliance was reduced, yet the system earned the lowest ratings [3].

Research gap. Most existing work on AI legal assistance concentrates on text-based systems outside India. Within the literature we reviewed, none integrates Indian constitutional and criminal-law information, the IPC–BNS transition, conversational context, speech and image interaction, verification prompts, and an evaluation that focuses on comprehension. That does not prove that no such system exists; the contribution of this work lies in the proposed integration and in the evaluation framework we put forward.

Table 1. Selected related work and its relevance to this study

Work	Reported finding	Limit relevant here
Dahl et al. [1]	58%–88% hallucination rates	Evaluation only
Magesh et al. [2]	17%–33% in commercial RAG tools	US proprietary tools
Buçinca et al. [3]	Forcing functions cut over-reliance	Non-legal tasks

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 11, CONFLUENCE - 2026

Work	Reported finding	Limit relevant here
This framework	Conversational GenAI over legal content	Informal testing only

A.3 Research Methodology

The study follows a design-science approach: we identify the problem, formulate objectives, develop the framework, carry out informal functional testing, and propose an evaluation protocol. Formal benchmarking and user studies remain future work. To keep the evidence honest, we separate findings into four classes: literature-supported (A), implementation-supported (B), observational (C) and proposed (D). Class A rests on cited sources, B on the implemented system, C on unrecorded informal testing, and D on features that are proposed but not yet implemented.

A.4 Objectives and Research Questions

Objectives. O1: design and implement a generative legal information-support system over structured Indian legal content. O2: support conversational, context-aware access, with speech and image input as accessibility aids. O3: build transparency, stated limitations and verification prompts into the interaction. O4: define a defensible evaluation framework and report the available evidence accurately.

Research questions. RQ1: how relevant, grounded and factually consistent are the responses? RQ2: can non-specialist users understand and use the information? RQ3: do the interface and responses communicate limitations and encourage verification? With no formal evaluation data yet, these questions will be addressed through the protocol in Section A.6; no hypotheses are tested and no empirical outcomes are claimed.

A.5 System Design and Implementation

Role of generative AI and NLP. The goal is explanation, not retrieval alone. A generative model can restate a provision in plain language and keep a conversation going, while NLP interprets free-text or spoken queries and matches them to a supported legal topic. Because fluent output can still be wrong, generation is deliberately coupled to structured legal content and to human verification.

A.5.1 Architecture

The framework uses a three-tier web architecture (Fig. 1). A React/Vite client captures text, voice and image input; a Node.js/Express server exposes REST APIs; MongoDB Atlas, accessed through Mongoose, stores both application data and the structured legal content; and generative models are reached through the OpenRouter API. Of the eight components, four are fully implemented, one is implemented and has been informally tested, and three remain proposed (Fig. 4(b)). Fig. 1 is conceptual; it is not a measured result.

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 11, CONFLUENCE - 2026

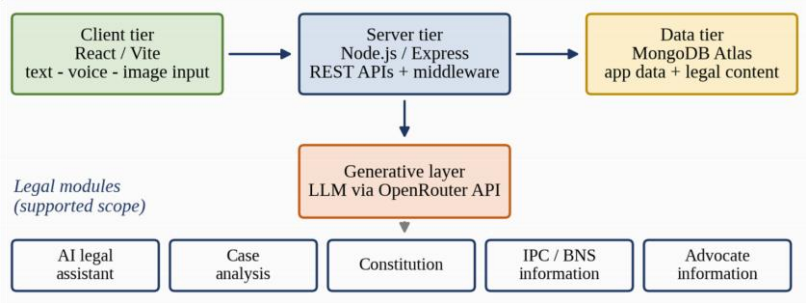


Fig. 1. Conceptual three-tier architecture and supported legal modules.

A.5.2 Query-processing workflow and logic

A query passes through seven steps (Fig. 2): the user submits text, voice or an image; speech is converted to text; the system interprets the query and checks its scope; context is assembled on the server; a response is generated; the answer is presented together with an information notice; and finally the user verifies the result. Step 7 lies outside the software by design — the system can prompt verification, but it cannot perform that verification itself.

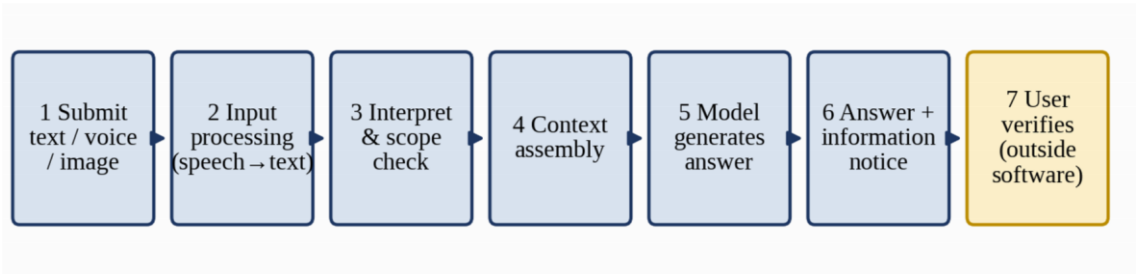


Fig. 2. Conceptual query-processing workflow (step 7 lies outside the software).

A.5.3 Contextual processing and multimodal interaction

Context assembly brings together the current query, the preceding turns and the relevant legal content, so that a follow-up such as “and what is the punishment?” is read against the provision already under discussion. The model instructions are written to ground answers in the supplied content, to say clearly when the material does not cover the question, and to avoid presenting uncertain material as settled. These prompt-level controls reduce unsupported generation, but they cannot eliminate it. Spoken and typed queries follow the same path, and responses can be read aloud; these features are accessibility aids rather than research contributions. A misrecognised section number would silently change the legal question, so the idea of showing the recognised text before submission is proposed, not claimed. Every response carries a notice that the output is information, not advice, together with a pointer to the official text or a qualified advocate (Fig. 3).

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 11, CONFLUENCE - 2026

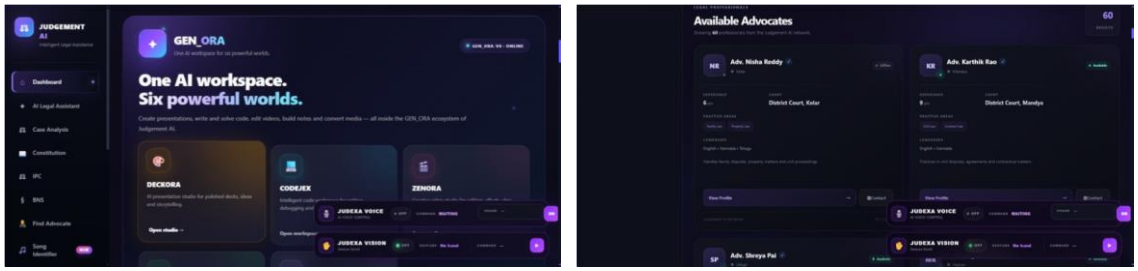


Fig. 3. Screenshots of the implemented prototype. Voice-control and gesture-scroll panels appear at lower right; gesture control is not evaluated here.

A.6 Evaluation Framework and Current Evidence

So far the framework has been examined only through informal testing of function, relevance and usability. No test set, scores or user data were recorded, so we report no accuracy, F1, usability or comprehension figures, and this part must not be read as a validation of response quality. Fig. 4(a) therefore shows literature values for context only, while Fig. 4(b) shows the status of the framework's own components. Table 2 lists the proposed criteria. Because model output is non-deterministic, each test should record the model identifier, the date and the settings used.

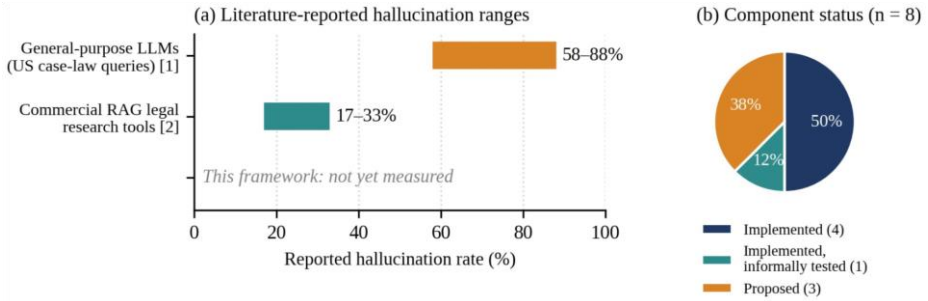


Fig. 4. (a) Hallucination ranges reported in [1] and [2], not results of this framework; (b) status of the framework's eight components.

Table 2. Evaluation criteria (all proposed; none yet measured)

Dimension	Question	Measure / instrument
Response relevance	Does the answer address the query?	Rater score 0–2
Retrieval quality	Is the correct provision identified?	Provision present (yes/no)
Contextual grounding	Is the answer based on supplied content?	Rater score 0–2; baseline
Factual consistency	Does it match the official text?	Rater score 0–2
Comprehension	Can lay users explain it afterwards?	Quiz score; Likert items

A.7 Advantages, Logic and Disadvantages

Advantages. The framework lowers the effort needed to reach and read legal text by explaining constitutional articles, IPC sections and BNS sections in plain language. Answers are read against earlier turns, speech input

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 11, CONFLUENCE - 2026

and output widen access, and scope stays bounded to defined legal content, with users routed to official text and advocate information. Typical uses are a student reading a basic provision, a citizen asking about a process and a reader comparing IPC and BNS provisions; none of these uses should be treated as a substitute for a lawyer.

Disadvantages and threats to validity. Hallucination remains possible even with supplied context [1], [2], and no rate has been measured here. Content can become outdated, as the IPC-to-BNS change shows [4], and the external generative layer can change behaviour. Simplification may omit nuance and procedures vary, so users should check the official text and consult an advocate. A small convenience sample would not generalise beyond Indian constitutional and criminal law.

A.8 Responsible AI, Privacy and Security

Legal queries can reveal sensitive personal circumstances, so responsible use covers privacy and security as well as accuracy; Table 3 lists the main risks and safeguards.

Table 3. Responsible-AI risks, safeguards and status

Risk	Safeguard	Status
Unsupported or outdated answer	Scope check; pointer to official text	Implemented
Uncritical acceptance	“Information, not advice” notice	Implemented
Misheard section number	Show recognised text before submission	Proposed
Unverified citation	Citation checking; forced verification	Proposed
Query privacy and key security	Minimise stored queries; disclose external model use	Proposed

A.9 Future Applications and Work

The immediate task is the evaluation protocol: a reference set of provisions, rater scoring on the Table 2 dimensions, an ungrounded baseline and a small comprehension study. Extensions include a retrieval index, citation checking and forced verification; wider uses include legal-literacy teaching, help-desk triage and regional-language access.

PART B — Adaptive AI Bosses and Blockchain Tokenomics: MIR5 Case Study

B.1 Introduction and Release Status

MIR5 is a PC open-world MMORPG developed by Wemade Next as the successor to MIR4. Wemade Next has announced Asterion, an AI boss built with NVIDIA ACE, and tokenomics designed for the WEMIX PLAY platform [6], [7]. This part examines how the two may interact, focusing on adaptive difficulty, rewards, economic stability and fairness. As of 8 October 2025 no official release date or beta schedule had

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 11, CONFLUENCE - 2026

been identified, and plans may have changed since. Evidence is classified as confirmed fact (F), vendor claim (V), reasoned inference (I) or undisclosed information (U).

B.2 Background and Related Work

Dynamic difficulty adjustment (DDA) adapts the challenge of a game to the performance of the player, while reinforcement learning allows agents to optimise their decisions. Retrieval-augmented generation [5] could, in principle, support the use of earlier battle records. Research on virtual economies stresses the need to balance currency sources and sinks and to allow for market volatility [8].

B.3 System Analysis: AI Boss Architecture and Tokenomics

Wemade Next announced Asterion on 7 January 2025, after six months of research and development with NVIDIA, and named NVIDIA ACE and NIM microservices as the supporting technologies (F) [6]. The company states that Asterion can perceive its environment, record previous battles and learn using machine learning and small language models (V) [6]. Model specifications, the scope of its memory, the action space, decision frequency, hosting arrangements, inference costs and any safeguards remain undisclosed (U). Fig. 5 shows a hypothetical architecture only.

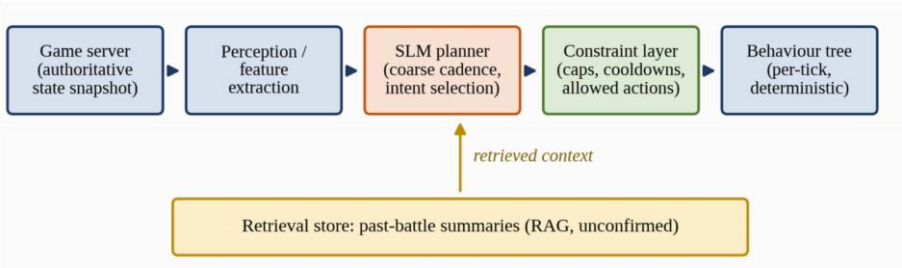


Fig. 5. Hypothesised Asterion pipeline (inference; not disclosed).

B.3.1 Tokenomics design

Earlier Wemade titles have gradually moved from multi-token stacks towards direct use of the platform token together with capped in-game coins (Table 4, compiled from Wemade statements and trade press). If MIR5 follows the same path, its economy would resemble Legend of YMIR more closely than MIR4 (I). That is a hypothesis, not a disclosed design (Fig. 6).

Table 4. Documented Wemade token models versus MIR5 (U = undisclosed; V = vendor claim)

Dimension	MIR4	MIR M	Legend of YMIR	MIR5
Core tokens	Dark steel → DRACO; later HYDRA	DRONE; governance token DOGMA	WEMIX used directly; capped coins	(U)
Sinks	Shop, staking, auctions	Not detailed	Capped crafting coins; auctions	(U)
Inter-game link	HYDRA with other WEMIX games	Explicit, with MIR4	Direct WEMIX integration	“Under discussion” (V)

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 11, CONFLUENCE - 2026

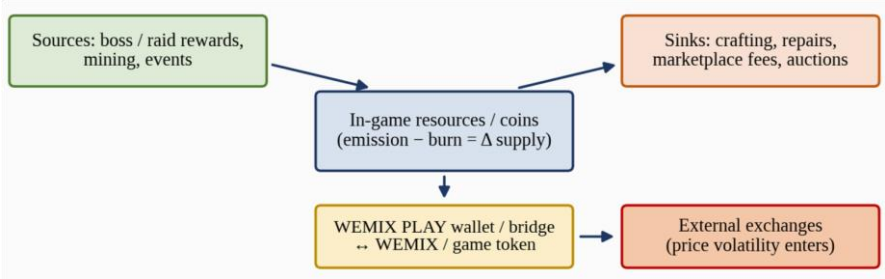


Fig. 6. Hypothesised MIR5 token flow (inference; not disclosed). The bridge is where external volatility and security risk enter.

WEMIX integration risk: In February 2025 a cross-chain bridge incident led to the theft of roughly 8.65 million WEMIX tokens (reported figures differ). Korean exchanges then delisted WEMIX for a second time, and the token fell by more than 60% [9]. Any MIR5 economy that is tied to WEMIX would inherit this platform-level risk (I).

B.4 Methodology and Results

The analysis is desk-based. Claims about MIR5, Asterion and tokenomics were extracted from public sources, tagged F, V, I or U, and rated ordinally across two panels — AI boss (14 cells) and tokenomics (28 cells). Fig. 7(a) shows the 42 ratings; inference-based ratings are counted in their stated bin and "Inherited" is counted as Unknown. Eleven ratings (about 26%) fall into the Unknown category, which shows how much remains undisclosed. Fig. 7(b) groups the 25 claims that we flagged for independent verification; tokenomics claims form the largest share (8 of 25).

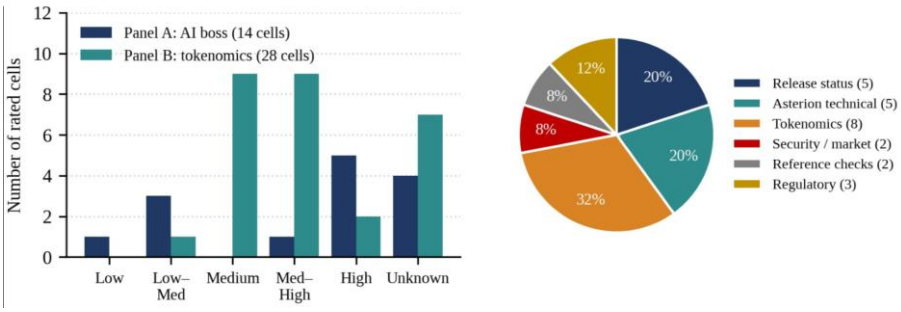


Fig. 7. (a) Histogram of the 42 ordinal ratings (Panel A: AI boss, 14 cells; Panel B: tokenomics, 28 cells); (b) the 25 claims flagged for independent verification, by category.

B.5 Illustrative Simulation

To demonstrate the coupling (Fig. 8), the authors built a browser simulation with invented values, not Asterion's logic. Five players (tank, healer, three damage dealers) fight a 3,000-HP boss for 120 seconds across 300 raids with identical random seeds; the adaptive boss targets top contributors (strength 1.5, 60% attack cap, 50% memory, 45 damage per second, seed 7). The scripted boss cleared 36% of raids (mean 40 s); the adaptive boss cleared none. Emission E is computed from players, attempts, clear rate and tokens per clear, and removal R from sink spending and burned fees, so the change in supply is $\Delta S = E - R$. With no adaptive-

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 11, CONFLUENCE - 2026

boss clears, emission falls to zero while 6,300 tokens are removed daily, driving supply to zero and the toy price index towards 28 million by day 90. This is an artefact of invented parameters, not a forecast; it shows only that emission depends on boss difficulty.

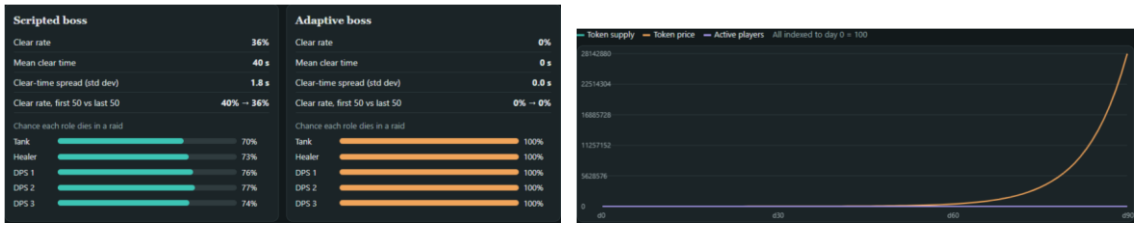


Fig. 8. Screenshots of the authors' illustrative demonstration (invented parameters).

B.6 Advantages, Logic and Disadvantages

Table 5. Advantages and disadvantages of the announced MIR5 design

Feature	Advantages (mostly vendor-claimed)	Disadvantages and risks
Adaptive boss	Fewer memorised fights; replay variety	Moving target for guilds; cost and exploit resistance undisclosed
Token economy	Rewards play value; asset ownership	Price volatility; bridge incidents [9]; regulatory exposure
Coupling (I)	Boss rewards can feed a sustained economy	Boss tuning changes emission; exploits become cash-outs; opaque adaptation may feel like gambling

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 11, CONFLUENCE - 2026

B.7 Conclusion and Cross-Case Synthesis

Both cases examine AI systems whose outputs carry real consequences — legal understanding in Part A, and token-based economic outcomes in Part B. The shared principles that emerge are the same: keep model behaviour bounded, classify the status of evidence transparently, retain human oversight, and measure before making performance claims. Part A relies on structured legal content and scope checks; Part B considers constraint layers and reward caps. Evidence is classified A–D in Part A and F/V/I/U in Part B. Part A has a proposed evaluation protocol but no formal results yet, while Part B relies on public claims and hypothetical analysis without access to the actual implementation. The conclusions therefore remain provisional until systematic evaluation can be carried out.

REFERENCES

[1] M. Dahl, V. Magesh, M. Suzgun and D. E. Ho, "Large legal fictions: Profiling legal hallucinations in large language models," *Journal of Legal Analysis*, vol. 16, no. 1, pp. 64–93, 2024.

[2] V. Magesh *et al.*, "Hallucination-free? Assessing the reliability of leading AI legal research tools," *Journal of Empirical Legal Studies*, vol. 22, no. 2, pp. 216–242, 2025.

[3] Z. Buçinca, M. B. Malaya and K. Z. Gajos, "To trust or to think: Cognitive forcing functions can reduce overreliance on AI in AI-assisted decision-making," *Proceedings of the ACM on Human-Computer Interaction*, vol. 5, CSCW1, 2021, doi: 10.1145/3449287.

[4] Government of India, *The Bharatiya Nyaya Sanhita, 2023*, Act No. 45 of 2023, effective 1 July 2024.

[5] P. Lewis *et al.*, "Retrieval-augmented generation for knowledge-intensive NLP tasks," in *Advances in Neural Information Processing Systems 33*, 2020, pp. 9459–9474.

[6] Wemade Next Co., Ltd., "Wemade Next to develop an AI boss in MIR5 in collaboration with NVIDIA," *Business Wire*, Jan. 7, 2025. [Online]. Available: https://www.businesswire.com/news/home/20250106300797/en

[7] "Wemade unveils new title MIR5 for the first time" (in Korean), *Byline Network*, Mar. 27, 2024. [Online]. Available: https://byline.network/2024/03/27-335/

[8] V. Lehdonvirta and E. Castronova, *Virtual Economies: Design and Analysis*. Cambridge, MA, USA: MIT Press, 2014.

[9] "Wemix becomes 1st cryptocurrency to be delisted twice in Korea," *The Korea Times*, May 2, 2025.

(This work is licensed under a Creative Commons Attribution 4.0 International License.)