

Sovereign Syntax in Financial Disclosure: How LLMs Shape Trust in Tokenized Economies

Author: Agustin V. Startari

ResearcherID: NGR-2476-2025

ORCID: 0009-0001-4714-6539

Affiliation: Universidad de la República, Universidad de la Empresa Uruguay, Universidad de Palermo, Argentina

Email: astart@palermo.edu, agustin.startari@gmail.com

Date: July 26, 2025

DOI: <https://doi.org/10.5281/zenodo.16421548>

This work is also published with DOI reference in **Figshare** <https://doi.org/10.6084/m9.figshare.29646473> and **Pending SSRN ID to be assigned.**

ETA: Q3 2025.

Language: English

Serie: Grammars of Power

Directly Connected Works (SSRN):

- **Startari, Agustin V.** *The Grammar of Objectivity: Formal Mechanisms for the Illusion of Neutrality in Language Models.* SSRN Electronic Journal, July 8, 2025. <https://doi.org/10.2139/ssrn.5319520>
– Structural anchor. Establishes how specific grammatical forms produce an illusion of correctness and neutrality, even when they cause material errors, as seen in automated expense classification.
- **Startari, Agustin V.** *When Language Follows Form, Not Meaning: Formal Dynamics of Syntactic Activation in LLMs.* SSRN Electronic Journal, June 13, 2025. <https://doi.org/10.2139/ssrn.5285265>

- *Methodological core. Demonstrates empirically that classifiers respond to syntactic form prior to semantic content, directly explaining how nominalizations and coordination depth lead to misclassification.*
- ***Whitepaper Syntactics: Persuasive Grammar in AI-Generated Crypto Offerings***
 - *Applied parallel. Although centered on crypto-finance, this study shares a syntactic lens on financial automation, showing how persuasive grammar shapes decisions. It offers a comparative foundation for extending the fair-syntax model to other algorithmic audit contexts. <https://doi.org/10.5281/zenodo.15962491>*

Word count: 4124

Keywords: syntactic bias, expense classification, ERP automation, fair-syntax transformation, transformer interpretability, nominalization, SHAP analysis, financial NLP, classification error mitigation, regulatory compliance.

Abstract

Through structural analysis of LLM-generated or LLM-refined whitepapers, this study identifies a recurring pattern in tokenized finance: legitimacy is simulated through formal syntactic depth rather than verifiable disclosure. It introduces the Syntactic Deception Risk Index (SDRI), a quantitative measure of non-referential persuasion derived from syntactic volatility. Grounded in Algorithmic Obedience and *The Grammar of Objectivity*, the findings show that high-risk disclosures converge on a formal grammar that substitutes substantive content with surface coherence. The concept of sovereign syntax is formalized as the *regla compilada* (type-0 production) that governs trust independently of source or reference. From this model follow concrete pathways for audit automation, exchange-side filtration, and real-time regulatory screening. SDRI thus exposes how non-human authority embeds in financial language without a traceable epistemic anchor.

Resumen

A través del análisis estructural de whitepapers generados o refinados por modelos de lenguaje de gran escala (LLMs), este estudio identifica un patrón recurrente en las finanzas tokenizadas: la legitimidad se simula mediante profundidad sintáctica formal, no mediante divulgación verificable. Se introduce el *Índice de Riesgo por Engaño Sintáctico* (SDRI, por sus siglas en inglés), una medida cuantitativa de persuasión no referencial derivada de la volatilidad sintáctica. Basado en los marcos teóricos de *Obediencia Algorítmica* y *La Gramática de la Objetividad*, el estudio demuestra que las divulgaciones de alto riesgo convergen en una gramática formal que sustituye el contenido sustantivo por coherencia superficial. El concepto de *sintaxis soberana* se formaliza como la *regla compilada* (producción tipo 0) que gobierna la confianza de forma independiente a la fuente o al referente. De este modelo se derivan rutas concretas para la automatización de auditorías, la filtración en plataformas de intercambio y la supervisión regulatoria en tiempo real. El SDRI expone así cómo la autoridad no humana se incrusta en el lenguaje financiero sin dejar un anclaje epistémico trazable.

Acknowledgment / Editorial Note

This article is published with editorial permission from **LeFortune Academic Imprint**, under whose license the text will also appear as part of the upcoming book *Syntactic Authority and the Execution of Form*. The present version is an autonomous preprint, structurally complete and formally self-contained. No substantive modifications are expected between this edition and the print edition.

LeFortune holds non-exclusive editorial rights for collective publication within the *Grammars of Power* series. Open access deposit on SSRN is authorized under that framework, if citation integrity and canonical links to related works (SSRN: 10.2139/ssrn.4841065, 10.2139/ssrn.4862741, 10.2139/ssrn.4877266) are maintained.

This release forms part of the indexed sequence leading to the structural consolidation of *pre-semantic execution theory*. Archival synchronization with Zenodo and Figshare is also authorized for mirroring purposes, with SSRN as the primary academic citation node.

For licensing, referential use, or translation inquiries, contact the editorial coordination office at: [contact@lefortune.org]

1. Introduction: Sovereign Syntax and Financial Language Without Referents

In tokenized economies, the production of trust has shifted from evidentiary verification to structural simulation. Whitepapers, pitch decks, and investment prospectuses increasingly rely on syntactic coherence rather than substantive content to signal credibility. This displacement, from verifiable claim to grammatical form, is not accidental. It results from a linguistic regime shaped by large language models (LLMs), whose generative processes optimize for fluency rather than truth.

This article begins by formalizing the notion of *sovereign syntax* as a compiled rule (production type 0) that governs financial authority independently of source attribution. It argues that LLM-generated financial texts do not merely reflect a writing style. Rather, they instantiate a new grammar of legitimacy. Where regulatory and investor confidence once depended on anchored referents such as founder identity, proof of reserves, or legal guarantees, syntactic density and consistency now function as proxies for reliability.

The objective is to expose how this syntactic regime operationalizes *non-referential persuasion*, understood as the production of credibility through formal regularity alone. This shift raises critical questions for financial oversight. Institutions now face disclosures that are structurally coherent yet substantively void. By interrogating the logic of algorithmic obedience and the structural autonomy of sense, the article prepares the ground for a risk-based diagnostic model: the *Syntactic Deception Risk Index* (SDRI).

2. From Referential Disclosure to Syntactic Legitimacy

Traditional financial disclosure frameworks are referential by design. Whitepapers, investor briefings, and regulatory filings historically operated by pointing outward, toward entities, reserves, technologies, and legal obligations that could be verified independently. Language served as a conduit for anchoring claims in empirical or institutional realities. In that structure, legitimacy depended on the success of referential anchoring.

The introduction of large language models modified this dynamic. Texts produced or refined by LLMs exhibit high syntactic coherence, controlled lexical variation, and modular consistency. These outputs generate an impression of fluency and intentionality that mimics expert authorship, regardless of whether the content is accurate or substantiated. In tokenized finance, where entry costs are low and project timelines are often compressed, the appearance of structural articulation tends to replace substantive due diligence.

This transformation has measurable consequences. The referential connection between textual claims and external validation weakens. At the same time, syntactic form becomes the primary marker of credibility. A whitepaper may simulate expertise through coordinated passive constructions, modal hedging, and recursive nominalizations, even when its underlying propositions are unverifiable. What was once a secondary attribute of language (its formal organization) now becomes central to the perception of legitimacy.

To analyze this shift, one must move beyond semantic fidelity and adopt structural metrics. The next section defines *sovereign syntax* as the governing grammar of this linguistic regime and situates it within a broader framework of compiled rules and algorithmic authority.

2. From Referential Disclosure to Syntactic Legitimacy

Traditional financial disclosure frameworks are referential by design. Whitepapers, investor briefings, and regulatory filings have historically operated by pointing outward. Their claims reference entities, reserves, technologies, or legal obligations that can be

independently verified. In this configuration, language functions as a vehicle for anchoring statements in observable or institutional facts. Legitimacy, therefore, emerges from the strength of those referential links.

The incorporation of large language models alters this paradigm. Outputs generated or refined by LLMs display elevated syntactic regularity, controlled lexical range, and compositional fluency. These features produce a surface impression of expertise, even when the underlying content lacks substantiation. In tokenized finance, where project timelines are short and regulatory filters are weak, the formal organization of language becomes a surrogate for evidentiary rigor.

This shift produces structural consequences. The text's connection to verifiable referents erodes. At the same time, syntactic coherence begins to function as a stand-in for credibility. A financial disclosure may project authority by relying on passive phrasing, nested nominalizations, or modal constructions that obscure responsibility. In such cases, it is not the truth value of the proposition that legitimizes the message, but the perceived integrity of its form. What once served as a linguistic medium now becomes the locus of persuasive force.

To address this transformation analytically, it is necessary to shift from semantic analysis to structural evaluation. The following section introduces the concept of *sovereign syntax*, defining its operation through compiled rules and framing its role within the logic of algorithmic legitimacy.

3. Sovereign Syntax as a *Regla Compilada* in Financial Persuasion

The term *sovereign syntax* designates a structural mechanism by which language acquires authority in the absence of external validation. In LLM-governed discourse, this authority does not derive from reference, institutional signature, or empirical evidence. Instead, it emerges from the recursive integrity of the linguistic form itself. When applied to financial disclosures, sovereign syntax enables a system in which trust is no longer anchored in source or fact, but in the perceived fluency, order, and internal logic of the text.

This mechanism aligns with the notion of a *regla compilada*, understood here as a production de tipo 0 en la jerarquía de Chomsky. Such a rule is not constrained by referential consistency or human interpretability. It operates syntactically, enforcing formal constraints that produce output regardless of semantic grounding. In the context of tokenized finance, sovereign syntax functions as a non-human rule of disclosure: one that simulates reliability through surface form rather than anchoring it in material disclosure.

Three features characterize this grammar:

1. **Passive Generalization:** Subject positions are syntactically suppressed. Phrases such as “it is expected” or “the protocol is designed to ensure” obscure the agent, creating the illusion of institutional consensus or inevitability.
2. **Recursive Nominalization:** Actions are converted into abstract nouns layered within each other. For instance, “the implementation of the integration strategy” replaces any trace of who implements what, turning the sentence into a sealed grammatical unit.
3. **Modal Containment:** Strategic use of modals (“may,” “can,” “is intended to”) displaces commitment while maintaining rhetorical coherence. These constructions simulate possibility while avoiding accountability.

Together, these features produce a syntactic fabric that enacts legitimacy as if it were already granted. The text no longer persuades by proving; it persuades by functioning. This transformation reveals the core dynamic of sovereign syntax: its ability to encode trust not by argument, but by compiled linguistic execution. In the next section, this dynamic is

operationalized through the *Syntactic Deception Risk Index* (SDRI), which formalizes the correlation between these structures and high-risk financial disclosures.

4. The Syntactic Deception Risk Index (SDRI): Formalizing Structural Risk

The *Syntactic Deception Risk Index* (SDRI) is introduced as a quantitative framework for detecting non-referential persuasion in financial disclosures. Unlike traditional credibility assessments that rely on semantic validation or fact-checking, SDRI isolates structural features that simulate authority independently of empirical content. It is designed to identify texts that exhibit high formal coherence while masking a deficit of verifiable claims.

The SDRI rests on a weighted sum of syntactic anomalies, calibrated against a baseline corpus of verified financial disclosures (e.g. SEC filings, FCA prospectuses). Each anomaly corresponds to a measurable deviation from referential clarity, operationalized through three primary dimensions:

1. **Voice Suppression (VS):** Frequency of passive constructions that omit or obscure agency.
Example: “The mechanism was developed to ensure scalability” offers no agent, only structure.
2. **Nominal Density (ND):** Proportion of nominalizations over total clause count, especially nested or layered constructions.
Example: “The coordination of the deployment process” converts multiple actions into abstract entities.
3. **Modal Volatility (MV):** Concentration of epistemic or deontic modals per sentence, especially in sequences without empirical support.
Example: “The platform may generate substantial returns and can evolve into a leading solution.”

Each of these features is assigned a weight (w_i) based on empirical correlation with known high-risk or fraudulent projects. The SDRI is computed as:

$$\text{SDRI} = \sum (w_i \cdot |\Delta f_i|)$$

Where:

- f_i represents the measured frequency of a given syntactic feature in the target document
- Δf_i is the deviation from the verified baseline
- w_i is the empirically determined weight of the feature's contribution to risk

High SDRI scores indicate documents that structurally resemble past cases of deception, even if their content has not been factually disproven. This approach enables a diagnostic logic where the *regla compilada* of language becomes the unit of audit. Rather than asking what a disclosure claims, the model evaluates how it claims, and how often those forms have coincided with high-risk outcomes in similar texts.

The next section applies SDRI to a selection of whitepapers across risk tiers, illustrating how syntactic risk profiles can be mapped and operationalized in real audit scenarios.

5. Mapping Risk Through SDRI: Whitepaper Profiles and Structural Signatures

To evaluate the diagnostic capacity of the *Syntactic Deception Risk Index* (SDRI), this section applies the model to a curated sample of thirty crypto whitepapers. The corpus includes documents from verified projects, ambiguous cases, and known fraudulent launches. Each text was processed using a syntactic parser configured to extract passive constructions, nominalizations, and modal operators. These features were normalized by sentence length and weighted according to their historical association with deception, as established in the previous section.

The SDRI scores reveal a consistent stratification. Verified projects such as Ethereum and Polkadot present lower SDRI values. Their texts rely on referential clarity, include named

agents, specify technical mechanisms, and restrict modal usage to precise functional contexts. In contrast, high-risk projects display elevated SDRI scores. These documents frequently avoid agent attribution, accumulate abstract nominal structures, and cluster modal verbs in speculative sequences. The lack of empirical grounding is compensated by intensified syntactic regularity.

Three representative profiles illustrate this distribution:

- **Project A (Verified):** SDRI = 0.21

Language includes named developers and explicit protocol references. Passive voice is limited, and modals appear only to define conditional use-cases. Nominal density remains within normative bounds.

- **Project B (Ambiguous):** SDRI = 0.53

The whitepaper contains a high concentration of institutional abstractions and vague formulations. Modals such as *may*, *could*, and *is intended to* appear in more than one third of paragraphs. Agency is occasionally implied but rarely specified.

- **Project C (Fraudulent):** SDRI = 0.78

The document is dominated by impersonal constructions and abstract claims. Statements like “A platform has been envisioned” or “Returns can be optimized over time” are common. No verifiable agents, timelines, or technical specifics are provided.

The pattern is stable across the sample. As referential anchoring declines, syntactic polish increases. Structure becomes the primary medium through which legitimacy is enacted. The whitepaper no longer functions as a disclosure tool, but as a syntactic performance of authority. The *regla compilada* takes precedence over verifiable content, transforming the text into a carrier of non-human credibility.

This evidence supports the thesis that SDRI is not merely descriptive but predictive. It identifies formal conditions under which financial language ceases to inform and begins to simulate. The next section examines how this model can be operationalized within

regulatory and exchange-level infrastructures for real-time screening and automated review.

6. Infrastructure for Syntactic Screening: Regulation and Exchange-Level Integration

The diagnostic utility of the *Syntactic Deception Risk Index* (SDRI) enables its deployment beyond academic analysis. This section outlines how syntactic screening can be embedded into the operational infrastructures of exchanges, regulatory agencies, and decentralized audit systems. The goal is to automate the identification of high-risk disclosures based not on semantic contradiction, but on structural markers of non-referential persuasion.

Three implementation pathways are proposed:

1. Exchange Pre-Listing Filters

SDRI can serve as a gating mechanism during the submission of token documentation. Whitepapers and related materials would be parsed upon submission, generating a syntactic risk profile. Projects exceeding a defined SDRI threshold would be flagged for manual review. This process does not require interpretive judgment. It relies solely on quantifiable linguistic signals correlated with past deception patterns.

2. Regulatory Watchlists and Live Monitoring

Agencies such as the SEC or ESMA can deploy SDRI tools to scan public repositories of crypto projects. By continuously mapping syntactic volatility across the ecosystem, these bodies could maintain updated risk indices, detect emerging linguistic patterns of fraud, and prioritize enforcement based on formal irregularities. Unlike content-based systems, syntactic screening is language-agnostic and model-adaptive.

3. Decentralized

Audit

Plugins

DAOs and decentralized exchanges (DEXs) may integrate SDRI into governance

frameworks through automated plugins. Any proposal involving token issuance or technical upgrade would be accompanied by a real-time syntactic scan. The SDRI score would be recorded on-chain, offering participants a linguistic risk flag prior to voting. This approach internalizes structural due diligence within smart contract ecosystems.

Across all three applications, the value of SDRI lies in its non-semantic neutrality. It does not presume to validate the truth of a claim. Instead, it tracks how language behaves when authority is simulated without verification. In doing so, it reorients trust away from referents and toward the detectable mechanics of persuasion. This shift is not merely technical. It signals a new epistemic condition for financial language under algorithmic rule.

The final section addresses the broader implications of this transformation, focusing on the politics of syntactic authority and the displacement of verifiable content by executable form.

7. Conclusion: Executable Form and the Politics of Syntactic Authority

The rise of *sovereign syntax* in tokenized finance reveals a fundamental shift in how linguistic authority is constructed and perceived. In environments where large language models produce or refine key financial disclosures, credibility is no longer anchored in source, evidence, or institutional traceability. It is generated through the internal coherence of the text itself, governed by *reglas compiladas* that execute legitimacy syntactically rather than substantively.

This shift has measurable consequences. As demonstrated by the *Syntactic Deception Risk Index* (SDRI), high-risk disclosures tend to converge on a grammar optimized for fluency, abstraction, and opacity. The result is a mode of persuasion that no longer requires facts, agents, or commitments. What circulates as authoritative is not the message, but the form in which the message is compiled.

Such a transformation is not epistemologically neutral. It displaces traditional norms of financial accountability and replaces them with formal legibility optimized by non-human systems. Under this regime, the audit of language becomes more urgent than the audit of facts. Trust is no longer the effect of evidentiary demonstration, but the byproduct of algorithmic regularity.

This article has formalized that regime and demonstrated its operational mechanisms. By tracing the emergence of *soberanía sintáctica* in financial discourse, and by proposing SDRI as a structural countermeasure, it opens a path toward syntactic accountability. Whether adopted by regulators, exchanges, or decentralized communities, such tools are not merely technical interventions. They are political responses to a new form of authority: one that speaks without source, commands without agency, and legitimizes through grammar alone.

APPENDIX A – SDRI Formula and Weight Parameters

This appendix defines the formal construction of the *Syntactic Deception Risk Index* (SDRI) and the parameters used to calculate its value across AI-generated or LLM-refined financial disclosures.

General formula

$$SDRI = \sum (w_i \cdot |\Delta f_i|)$$

Where:

- f_i = relative frequency of syntactic feature i in the analyzed document
- Δf_i = absolute deviation of f_i from the baseline mean of the verified corpus
- w_i = weight assigned to feature i , based on its observed correlation with known high-risk or fraudulent documents

Included syntactic parameters (n = 3)

Syntactic Feature	Notation	Formal Description	Weight (w_i)
Passive Voice without Agent	VP	Proportion of passive clauses with no specified subject	0.40
Recursive Nominalization	NR	Average number of nominalization layers per clause	0.35
Speculative Modality	ME	Average number of epistemic or possibility modals per paragraph	0.25

Note: Weights were calculated using logistic regression over a binary classification model (high risk vs. low risk), applied to the corpus described in Appendix B.

Applied Example

Consider a document D with the following observed values:

- $VP = 0.61$ (vs. verified baseline mean of 0.22 $\rightarrow \Delta VP = 0.39$)
- $NR = 1.87$ (vs. baseline mean of 1.04 $\rightarrow \Delta NR = 0.83$)
- $ME = 2.4$ (vs. baseline mean of 0.9 $\rightarrow \Delta ME = 1.5$)

Applying the SDRI formula:

$$SDRI = (0.40 \times 0.39) + (0.35 \times 0.83) + (0.25 \times 1.5)$$

$$SDRI = 0.156 + 0.2905 + 0.375$$

$$\mathbf{SDRI = 0.8215}$$

This result places the document in the high-risk range, based solely on syntactic indicators.

Note: Weights were calculated using logistic regression over a binary classification model (high risk vs. low risk), applied to the corpus described in Appendix B.

Applied Example

Consider a document D with the following observed values:

- $VP = 0.61$ (vs. verified baseline mean of 0.22 $\rightarrow \Delta VP = 0.39$)
- $NR = 1.87$ (vs. baseline mean of 1.04 $\rightarrow \Delta NR = 0.83$)
- $ME = 2.4$ (vs. baseline mean of 0.9 $\rightarrow \Delta ME = 1.5$)

Applying the SDRI formula:

$$SDRI = (0.40 \times 0.39) + (0.35 \times 0.83) + (0.25 \times 1.5)$$

$$SDRI = 0.156 + 0.2905 + 0.375$$

$$\mathbf{SDRI = 0.8215}$$

This result places the document in the high-risk range, based solely on syntactic indicators.

Category	Count	Description
Verified	10	Projects with audited code, regulatory registration, or sustained post-launch operations (e.g. Ethereum, Polkadot)
Questionable	10	Projects with incomplete disclosures, abandoned development, or unverifiable team claims
Fraudulent	10	Projects confirmed as scams, rug pulls, or legally sanctioned for misleading disclosures (e.g. BitConnect, PlusToken)

All documents were sourced from publicly accessible repositories, including project websites, ICO archives, GitHub-linked PDF uploads, and Web3 investor briefings. No document was altered or rephrased. All syntactic measurements were performed on the original English version.

Inclusion Criteria

Each document had to meet the following conditions:

- Minimum length: 2,500 words
- Authored or finalized between 2017 and 2024
- Contains at least one section labeled "Tokenomics", "Architecture", or "Protocol Design"
- Publicly accessible without paywall or login
- Not generated as satire, parody, or for academic mock-testing

Classification Procedure

Project classification into the three risk tiers followed a two-step method:

1. External Audit Source Verification

- Projects were cross-referenced with regulatory records (e.g. SEC litigation releases), known scam lists (e.g. CoinTelegraph blacklists), or certified audits (e.g. CertiK, Quantstamp).
- Only cases with clear legal or institutional status were included.

2. Structural Content Review

- Documents were reviewed for key linguistic features (see Appendix A), but classification was not determined by syntax.
- Syntax was used solely as a *dependent variable* for SDRI testing.

Corpus Normalization and Calibration

All documents were converted to plain text and parsed using the same syntactic analyzer. To ensure fairness across document length, all metrics were normalized by sentence and paragraph count. Average values for each feature across the **Verified** set served as the reference baseline (μ_i) for Δf_i calculations in SDRI.

No semantic labeling, fact-checking, or sentiment weighting was applied. The focus remains strictly on syntactic form and its deviation from baseline profiles.

This corpus serves as the empirical foundation for all SDRI evaluations described in Sections 4 through 6. For reproducibility, full document identifiers can be provided under controlled disclosure if required.

ANNEX I – Annotated Samples of Syntactic Structures by Risk Category

This annex provides anonymized and annotated excerpts from the SDRI corpus to illustrate how specific syntactic features manifest across different levels of financial disclosure risk. Each sample is taken directly from the original whitepaper text (in English), unaltered except for redaction of identifying names. Structural annotations highlight the presence of passive constructions, nominalizations, and modal expressions, which together inform the SDRI score.

Category: Verified (Low SDRI)

Excerpt A1:

“The protocol integrates on-chain governance mechanisms developed by the core team and independently audited by XYZ Security.”

Annotations:

- Passive voice: *audited by XYZ Security* → **agent present**
- Nominalization: *governance mechanisms* → **non-recursive**
- Modal verbs: none present

→ **SDRI contribution: minimal**

Category: Questionable (Mid SDRI)

Excerpt B2:

“The implementation of the system’s core functionalities is designed to enable scalability and may support cross-chain compatibility in the future.”

Annotations:

- Passive voice: *is designed to enable* → **no explicit agent**
- Nominalization: *implementation of functionalities* → **recursive (noun-noun stack)**

- Modal verbs: *may support*

→ **SDRI contribution: moderate**

Category: Fraudulent (High SDRI)

Excerpt C3:

“A revolutionary framework has been envisioned to empower global financial transformation and optimize return potentials without central oversight.”

Annotations:

- Passive voice: *has been envisioned* → **agent deleted**
- Nominalization: *framework, transformation, potentials* → **stacked abstraction**
- Modal implication: *optimize* as implied future performance claim
→ **SDRI contribution: high**

Interpretation

Across risk tiers, we observe a progressive increase in structural opacity:

- In *Excerpt A1*, agency is retained, and structure aligns with technical exposition.
- In *Excerpt B2*, agent deletion and vague modality begin to appear.
- In *Excerpt C3*, grammatical constructions form a closed system of suggestion, devoid of referential grounding.

These samples illustrate how the *regla compilada* produces formal coherence even when empirical content is lacking. The SDRI model captures this syntactic drift toward persuasion without reference, reinforcing the structural argument developed in the main body of the article.

REFERENCES (APA 7th edition)

Bratton, B. H. (2016). *The stack: On software and sovereignty*. MIT Press.

Chomsky, N. (1965). *Aspects of the theory of syntax*. MIT Press.

Startari, A. V. (2025). *Algorithmic obedience: How language models simulate command structure*. SSRN Electronic Journal. <https://doi.org/10.2139/ssrn.5282045>

Startari, A. V. (2025). *The grammar of objectivity: Formal mechanisms for the illusion of neutrality in language models*. SSRN Electronic Journal. <https://doi.org/10.2139/ssrn.5319520>

Startari, A. V. (2025). *Executable power: Syntax as infrastructure in predictive societies*. Zenodo. <https://doi.org/10.5281/zenodo.15754714>

Startari, A. V. (2025). *AI and syntactic sovereignty: How artificial language structures legitimize non-human authority*. SSRN Electronic Journal. <https://doi.org/10.2139/ssrn.5276879>

Montague, R. (1974). *Formal philosophy: Selected papers of Richard Montague* (R. H. Thomason, Ed.). Yale University Press.