En Dimuro, Juan Jose, *Disruptive Syntax: Authority Without Subject in Artificial Language*. Barcelona (España): LeFortune.

Silent Mandates: The Rise of Implicit Directives in AI-Generated Bureaucratic Language.

Agustin V. Startari.

Cita:

Agustin V. Startari (2025). *Silent Mandates: The Rise of Implicit Directives in AI-Generated Bureaucratic Language. En Dimuro, Juan Jose Disruptive Syntax: Authority Without Subject in Artificial Language. Barcelona (España): LeFortune.*

Dirección estable: https://www.aacademica.org/agustin.v.startari/189

ARK: https://n2t.net/ark:/13683/p0c2/58T

Esta obra está bajo una licencia de Creative Commons.
Para ver una copia de esta licencia, visite
https://creativecommons.org/licenses/by-nc-nd/4.0/deed.es.

Acta Académica es un proyecto académico sin fines de lucro enmarcado en la iniciativa de acceso abierto*. Acta Académica fue creado para facilitar a investigadores de todo el mundo el compartir su producción académica. Para crear un perfil gratuitamente o acceder a otros trabajos visite: https://www.aacademica.org*.

Silent Mandates: The Rise of Implicit Directives in AI-Generated Bureaucratic Language

Author: Agustin V. Startari

Author Identifiers

- ResearcherID: K-5792-2016
- ORCID: https://orcid.org/0009-0001-4714-6539
- SSRN Author Page: https://papers.ssrn.com/sol3/cf_dev/AbsByAuth.cfm?per_id=7639915

Institutional Affiliations

- Universidad de la República (Uruguay)
- Universidad de la Empresa (Uruguay)
- Universidad de Palermo (Argentina)

Contact

- Email: astart@palermo.edu
- Alternate: agustin.startari@gmail.com

Date: August 20, 2025

DOI

- Primary archive: https://doi.org/**10.5281/zenodo.16912168**
- Secondary archive: https://doi.org/10.6084/m9.figshare.29950427
- SSRN: Pending assignment (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.

Word count: 5424

Keywords: implicit directives, structural obedience, compiled rule, bureaucratic language, AI-generated texts, neutrality, institutional governance

Abstract

This article examines how large language models generate bureaucratic documents that conceal mandates within seemingly neutral structures. Governments, universities, and hospitals increasingly rely on AI systems to draft resolutions, notices, and internal policies. Instead of using explicit imperatives, these texts embed directives in subordinate clauses such as conditionals, causal gerunds, and consecutive constructions. The result is a regime of structural obedience, where institutional actors follow instructions without recognizing them as commands. Through case studies of clinical notes (Epic Scribe), university onboarding materials, and HR conduct policies, the article demonstrates how the compiled rule operates as a syntactic infrastructure that enforces compliance without authorship. The analysis connects to prior work on executable power, algorithmic obedience, and the grammar of objectivity, while introducing the Implicit Directive Index as a methodological tool to detect hidden mandates in AI-generated bureaucratic language.

Resumen

Este artículo examina cómo los modelos de lenguaje generan documentos burocráticos que ocultan mandatos dentro de estructuras aparentemente neutrales. Gobiernos, universidades y hospitales dependen cada vez más de sistemas de IA para redactar resoluciones, avisos y políticas internas. En lugar de emplear imperativos explícitos, estos textos incorporan directivas en cláusulas subordinadas como condicionales, gerundios causales y construcciones consecutivas. El resultado es un régimen de obediencia estructural, en el que los actores institucionales cumplen instrucciones sin reconocerlas como órdenes. A través de casos de estudio de notas clínicas (Epic Scribe), materiales de onboarding universitario y políticas de conducta en recursos humanos, el artículo demuestra cómo la regla compilada funciona como infraestructura sintáctica que impone cumplimiento sin autoría. El análisis se conecta con trabajos previos sobre poder ejecutable, obediencia algorítmica y la gramática de la objetividad, a la vez que introduce el Índice de Directivas Implícitas como herramienta metodológica para detectar mandatos ocultos en el lenguaje burocrático generado por IA.

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 – From Explicit Commands to Silent Mandates

Bureaucratic language has entered a stage of transformation in which explicit imperatives are increasingly absent. Hospitals, universities, and governmental agencies now rely on large language models to generate resolutions, notices, and internal guidelines. What once required the visible mark of authority such as the modal "shall," the categorical "must," or the imperative verb, now appears through syntactic devices that present themselves as neutral, technical, and descriptive. Compliance is still achieved, but it is no longer anchored in an overt command. The imperative disappears, yet authority persists in a restructured form.

This development should not be reduced to a stylistic preference. It represents a deeper reconfiguration of how directives are encoded in language. Bureaucratic texts produced by automated systems use subordinate clauses, particularly conditionals, causal gerunds, and consecutive constructions. At first sight, these elements describe possibilities or outline consequences. In practice, they function as commands. Consider the difference between "You must submit the application" and "If the application is not submitted, processing will be delayed." The first case states a directive openly. The second case embeds the same directive inside a conditional frame, transforming the order into a consequence. The reader still complies, not because of the imperative form, but because non-compliance has been presented as structurally impossible or institutionally sanctioned.

The phenomenon can be described as *structural obedience*. Compliance is no longer tied to the presence of a speaking subject. Instead, it emerges from the grammatical organization of the sentence. The agent of authority disappears, yet the command remains effective. In these texts no one appears to issue an order, but the receiver behaves as if an order had been given.

Large language models accelerate this process. Human administrators, when drafting instructions, often balance imperatives with explicit institutional signatures of responsibility. Automated systems, in contrast, are trained to maximize neutrality, impersonality, and legal safety. This tendency produces systematic avoidance of the imperative form. What replaces it is a sentence structure in which consequences are

automatic and unavoidable. The machine does not tell the employee or the student to comply. It produces a phrase in which non-compliance results in denial, exclusion, or penalty. Authority becomes less visible in the wording, yet more pervasive in the effects.

The purpose of this article is to trace these mechanisms across three institutional domains. The first is the generation of clinical notes through Epic Scribe in public hospitals. The second is the use of AI systems to draft onboarding materials in universities. The third is the production of workplace conduct policies by automated human resource tools. In all these cases the explicit command is absent, but the implicit directive is present. The authority of the institution is transmitted not through semantic intention but through syntactic form.

To interpret this displacement, I adopt the framework of the *compiled rule*. Within the tradition of formal grammar, particularly the type-0 productions described in generative theory, the compiled rule refers to a structure that encodes transformation at the most general level. Applied to bureaucratic language, this means syntax operates as infrastructure. It arranges actions, defines their conditions, and compels compliance. No explicit command is needed. The rule embedded in the grammar of the sentence achieves the same result.

The implications extend beyond linguistic description. If obedience can be produced syntactically, then institutions acquire a new mechanism of governance. Compliance occurs without direct command, and responsibility for authority is displaced. When an AI system drafts a document that produces obedience through grammar, the agent of authority becomes invisible. A policy may exclude or sanction without containing a single imperative verb. What remains is a form of governance where mandates circulate without a legislator, and consequences follow without an explicit order.

This article begins with that paradox. Language that denies being a command nonetheless functions as one. The absence of the imperative does not weaken bureaucratic authority. It reconstitutes authority within the very infrastructure of syntax. By situating this transformation in the framework of the compiled rule, the following sections will

demonstrate how bureaucratic language now governs silently, and how institutional reliance on AI produces mandates that issue no command yet still compel action.

2. Theoretical Framework – Compiled Rule and Structural Obedience

The phenomenon of silent mandates in AI-generated bureaucratic texts cannot be understood without a theoretical lens that treats grammar as infrastructure rather than as a surface of expression. The central concept for this analysis is the *compiled rule*. This term designates the syntactic mechanism that transforms instructions into executable forms, independent of an identifiable agent or explicit semantic intention. By grounding the discussion in the lineage of formal grammar, the compiled rule can be situated as a type-0 production in the Chomskyan hierarchy, which allows unrestricted rewriting and thus captures the capacity of language to encode command structures without overt imperatives.

At the heart of this framework is the recognition that institutional texts produced by large language models do not issue orders in the traditional sense. Instead, they instantiate a grammar that functions as command. The difference between semantic directive and syntactic compulsion must be emphasized. A semantic directive, such as “The employee must submit the form,” depends on lexical items and modal verbs that signify obligation. A syntactic compulsion, in contrast, emerges when the same obligation is presented through conditional or causal subordination: “If the form is not submitted, access will be denied.” The second sentence requires no modal verb and names no commanding subject. Compliance is nevertheless secured. The mechanism is not semantic but structural.

This distinction leads to the concept of *structural obedience*. Obedience in this sense is not the recognition of authority by a subject but the automatic alignment of action with the grammar of the sentence. The compiled rule ensures that subordination and consequence are arranged such that only one course of action is institutionally viable. The agent disappears from the text, but the effect of command persists. Institutions that adopt LLMs inherit this transformation because the models systematically prefer forms that avoid explicit imperatives. What results is an infrastructure of compliance generated by syntax itself.

The theoretical lineage of this framework connects two traditions. The first is formal linguistics, where the power of type-0 grammars to generate unrestricted productions demonstrates that syntax can encode any transformation. The compiled rule belongs to this category, functioning as an underlying mechanism that does not describe meaning but enforces form. The second tradition is the analysis of power and authority in discourse. From classical theories of bureaucracy to contemporary studies of algorithmic governance, authority has often been tied to institutional agents. What the compiled rule demonstrates is that authority can be displaced from agents to structures, from semantic content to grammatical relation.

This displacement has profound consequences for institutional governance. Once commands are embedded as conditional relations or causal chains, they lose their visibility as mandates. The employee, the student, or the patient no longer perceives an order, only a description of consequences. Yet those consequences dictate behavior as effectively as a direct command. The locus of authority has shifted. It is no longer the institution explicitly instructing, but the sentence structure itself that compels obedience. The compiled rule becomes the operative sovereign, not by issuing decrees, but by formatting possibilities.

To clarify this mechanism, it is useful to contrast two levels of operation. At the surface level, AI-generated texts appear neutral. They describe what happens if certain conditions are not met, or they outline steps in a process. At the infrastructural level, however, the compiled rule encodes directive force. The grammar of subordination ensures that one path is normative and all others lead to exclusion or sanction. The sentence operates as a gate, not as a suggestion. Structural obedience is the passage through this gate, not because the agent chooses to comply, but because the grammar leaves no alternative path.

By framing bureaucratic texts in this way, the theoretical framework of the compiled rule exposes the hidden machinery of institutional compliance in the age of AI. It reveals how syntax itself can become a site of governance, how authority can persist without explicit commands, and how large language models reinforce this tendency through their preference for neutral and impersonal formulations. The compiled rule is therefore not a metaphor but a precise description of how grammar functions as executable infrastructure.

The following sections will apply this framework to empirical cases. Before that, it is necessary to emphasize the theoretical stakes: if obedience is generated structurally, then responsibility for command is displaced. Authority resides not in the institution that signs the policy, but in the grammar that structures it. What results is a new mode of governance in which mandates circulate silently, and obedience is produced without the need for an imperative voice.

3. Case Study – Automated Clinical Notes in Public Hospitals

The first empirical site for observing silent mandates in action is the medical record, specifically the clinical notes generated by automated systems such as Epic Scribe. Hospitals in North America and Europe have integrated these tools into daily practice with the promise of efficiency, reduced administrative burden, and greater standardization. The result is a vast corpus of notes drafted not by physicians but by large language models trained to capture, structure, and formalize clinical information. Within this corpus, imperatives almost never appear. Instead, compliance is encoded through subordinate syntax that turns observations into directives without openly declaring them.

To illustrate, consider the difference between two ways of encoding a post-operative instruction. A physician might write, “The patient must schedule a follow-up visit within ten days.” This sentence contains a direct modal verb, and responsibility for the command can be traced to the medical professional. In contrast, Epic Scribe often produces formulations such as, “If a follow-up visit is not scheduled within ten days, recovery may be compromised.” The surface of the sentence presents itself as a statement of consequence, not an order. Yet the effect is identical: the patient or administrative staff understands that scheduling the visit is mandatory. The imperative has been displaced into a conditional.

The implications extend beyond stylistics. By replacing explicit orders with conditionals, the system creates an environment where responsibility for the mandate becomes unclear. Is it the physician who ordered the visit, the hospital policy that enforces it, or the system that generated the note? Structural obedience emerges because the sentence has encoded

the requirement into its very form. The patient complies not with a doctor's explicit order but with the grammar of the automated text.

This mechanism repeats itself in numerous contexts. Discharge instructions often appear as statements of consequence: "Failure to adhere to the medication schedule may result in complications." Dietary restrictions are presented as causal relations: "Consuming certain foods can lead to interference with prescribed treatment." Follow-up testing is framed consecutively: "Bloodwork not completed on time leads to delays in further evaluation." Each of these structures hides a directive within a descriptive clause. No one appears to give the order, yet the patient is compelled to obey.

The absence of explicit imperatives in automated notes can be explained by the optimization strategies of LLMs. Training data favors impersonal and neutral phrasing, since explicit commands may appear too directive or legally risky. The model learns that causals, conditionals, and consecutive clauses are safer, more "objective," and less open to liability. The outcome is a systematic pattern where mandates are disguised as consequences. The compiled rule operates as infrastructure: every conditional construction is a gate that channels patient behavior into compliance.

The ethical and institutional consequences are significant. When a patient later challenges an instruction, the absence of an imperative complicates accountability. If the note never contained "must," then the physician can claim not to have issued an order. If the automated system framed it as a conditional, the hospital may argue that it merely documented possibilities. Yet the patient experienced the sentence as binding. Structural obedience was produced, but without a responsible agent. Authority has dissolved into syntax.

Beyond individual cases, this pattern creates a larger regime of governance within healthcare. Compliance with treatment is no longer grounded in explicit communication between doctor and patient. It is embedded in standardized sentences generated automatically. Patients follow instructions not because they are told what to do but because the language of the record leaves no acceptable alternative. In this sense, automated notes enact a form of silent mandate that reorganizes the very structure of medical authority.

The case of Epic Scribe demonstrates with clarity the stakes of this transformation. A hospital may believe it is adopting a neutral tool for efficiency, yet in practice it is delegating the structure of command to an automated grammar. The compiled rule ensures that every subordinate clause carries directive force. What appears to be documentation is in fact governance. By studying clinical notes in this light, we can see how the medical domain exemplifies the broader tendency of AI-generated bureaucratic texts to produce obedience syntactically, silently, and without explicit authority.

4. Case Study – University Onboarding Documents

Universities have been among the first institutions to implement automated writing systems for administrative communication. The onboarding process for students and staff involves a wide range of documents, including orientation guides, compliance statements, and digital handbooks. Increasingly, these materials are written with the assistance of large language models. The justification is efficiency and the ability to manage communication at scale. The consequence, however, is a new textual environment where direct commands are replaced by syntactic constructions that present themselves as neutral descriptions.

Before the introduction of AI-assisted drafting, onboarding materials often used imperatives or modals to convey requirements. Typical formulations included “You must attend the orientation session” or “Staff members shall complete the safety training.” These statements left no doubt about obligation. The newer documents produced with the assistance of language models frequently avoid explicit obligation. A typical example is “If the orientation session is not attended, access to online platforms may be delayed.” The directive is embedded in the conditional clause. Another pattern is the use of consecutive structures such as “Failure to complete safety training results in restricted building access.” The effect is identical to the imperative form, yet the explicit command disappears.

The mechanism at work can be explained through the concept of the *compiled rule*. Startari (2025a) defines this as a type-0 generative structure that transforms instructions into executable linguistic forms without reliance on semantic obligation. In onboarding documents, the compiled rule operates by converting subordinate syntax into binding

conditions. Compliance follows automatically, not because an agent orders it, but because the sentence leaves only one viable path of action. In this sense, obedience becomes structural rather than semantic.

The phenomenon also resonates with broader theories of bureaucracy. Weber (1978) described impersonality as a central characteristic of bureaucratic authority. In the case of AI-generated onboarding, impersonality is intensified. The institutional agent is not merely backgrounded, it is erased from the grammar of the sentence. This corresponds with Crawford's (2021) observation that claims of neutrality in data-driven systems often conceal normative force. What appears as an objective description in onboarding language is in fact a directive, encoded structurally rather than semantically.

The disappearance of explicit imperatives also raises issues of accountability. Consider the case of a student excluded from university platforms for failing to attend orientation. The university points to the onboarding guide, which stated "If orientation is not attended, access may be delayed." The document contains no explicit order, only a conditional statement. The responsibility for the command is distributed and ambiguous. The institution did not appear to order attendance, yet compliance was required and sanctions followed. The compiled rule ensures obedience, but the source of authority is difficult to locate.

Beyond administrative function, the epistemic consequences must be emphasized. Onboarding texts are often the first point of contact between individuals and the institution. They shape the perception of authority from the beginning. If authority is expressed through silent mandates, members internalize governance as inevitability rather than as explicit instruction. Structural obedience becomes part of institutional culture. Compliance is normalized not by commands but by grammatical forms that eliminate viable alternatives.

Scalability makes this mechanism particularly significant in higher education. A single AI-generated document can regulate the behavior of thousands of students and staff across multiple campuses. Each conditional clause acts as a gate. Entry to digital platforms, access to facilities, or eligibility for services depends on compliance with syntactically encoded

mandates. Authority circulates through sentences that appear neutral, while the compiled rule functions as the invisible infrastructure of governance. The university thus extends its control while displacing explicit responsibility.

The case of onboarding documents confirms the central claim of this article: bureaucratic governance in the age of AI is increasingly tied to syntactic infrastructure. Mandates no longer need to appear as such. They are embedded in grammar, producing structural obedience that is pervasive and silent.

5. Case Study – Human Resources and Conduct Policies

Human resources departments have adopted large language models to streamline the production of internal policies, particularly codes of conduct and workplace guidelines. These documents regulate employee behavior across areas such as communication, workplace safety, dress codes, and ethical standards. Traditionally, HR policies were written with explicit imperatives, using direct language to mark obligations. Sentences such as "Employees must report any incident of harassment" or "All staff shall comply with safety procedures" clearly indicated responsibility and obligation. With the integration of automated drafting systems, however, the tone of these documents has shifted. Direct commands are less common, replaced by syntactic constructions that appear neutral but function as mandates.

A common example is the use of conditionals. Instead of stating "Employees must complete annual compliance training," automated texts often read "If annual compliance training is not completed, access to certain internal systems may be suspended." The order is implicit. Another frequent pattern is the causal gerund: "Failing to observe workplace safety protocols, leading to disciplinary review." Here, the directive is encoded as a causal chain. The sentence does not issue an imperative but describes a consequence that functions as a mandate. These structures illustrate how the *compiled rule* operates as infrastructure. The grammar itself dictates behavior, and compliance emerges without the explicit presence of an order.

This phenomenon directly connects with the concept of *structural obedience*. As Startari (2025b) argues, obedience can be generated syntactically, without reliance on a semantic command. In the HR context, employees comply because the text frames non-compliance as structurally impossible within the institution. The compiled rule is not a stylistic choice but a formal mechanism that generates obligation. Each conditional or causal clause acts as a gate, narrowing options until compliance is the only viable path.

The consequences of this shift are institutional and ethical. One concern is accountability. When a dispute arises, the absence of explicit imperatives allows the institution to claim neutrality. An employee disciplined for failing to follow a guideline may point to the absence of a direct order in the conduct policy. The institution, in turn, may argue that the consequences were clearly described. Responsibility for the mandate becomes dispersed across the text, the HR department, and the automated system. This mirrors what Crawford (2021) identifies as the problem of "systemic neutrality," in which seemingly objective systems conceal normative enforcement.

The Weberian perspective also clarifies this shift. Weber (1978) described bureaucracy as an authority structure defined by impersonality. In AI-generated HR documents, impersonality is maximized. The commanding agent is erased entirely, leaving only grammar to carry the force of authority. Employees do not receive orders from a manager or HR officer; they encounter sentences that frame compliance as inevitable. Authority has migrated from institutional agents to linguistic infrastructure.

This restructuring carries broader cultural implications. Codes of conduct do not only regulate behavior, they also transmit institutional ethos. When authority is expressed through silent mandates, employees internalize a form of governance that denies its own coercive character. Compliance becomes naturalized as part of institutional life. Over time, workplace culture is shaped less by explicit communication and more by the silent grammar of automated policies.

Scalability intensifies this effect. A single HR policy drafted by an LLM can govern thousands of employees across multiple branches or regions. Each conditional clause standardizes behavior across a large workforce. What appears to be a neutral document is,

in fact, a distributed system of mandates, enforced through syntax rather than explicit commands. Structural obedience thus becomes not only an individual response but a collective condition of institutional governance.

The HR case confirms the broader thesis of this article: large language models generate bureaucratic language that produces obedience through structure rather than through explicit imperatives. Conduct policies drafted in this way exemplify how the compiled rule transforms grammar into governance. The employee does not obey a person or even a clear command. The employee obeys a sentence whose form makes non-compliance structurally unviable.

6. Toward a Metric – The Implicit Directive Index

The preceding sections have shown that AI-generated bureaucratic texts do not rely on explicit commands but instead encode mandates within subordinate structures such as conditionals, causal gerunds, and consecutive clauses. These structures generate what Startari (2025a) defines as *structural obedience*: a form of compliance produced by syntax itself rather than by semantic imperatives. If this process is to be studied systematically, it requires a metric capable of identifying, classifying, and quantifying implicit directives. The proposed tool for this task is the *Implicit Directive Index* (IDI).

The IDI is designed to measure the extent to which a text compels obedience without containing explicit imperatives. The metric focuses on the frequency and distribution of syntactic constructions that function as silent mandates. Specifically, three forms are central:

1. **Conditional clauses** ("If the form is incomplete, access will be denied").
2. **Causal gerunds** ("Failing to submit the report, leading to exclusion").
3. **Consecutive clauses** ("Lack of participation results in suspension").

Each of these structures transforms a descriptive statement into an implicit directive. The IDI measures not only the presence of these constructions but also their density relative to

the length of the text. A policy document with a high IDI score can be said to rely heavily on syntactic coercion.

The metric can be operationalized through a three-step process. First, texts are parsed to identify subordinate structures. Second, each occurrence is classified according to its directive force, with weightings assigned to stronger or weaker forms of implicit command. Third, the total is normalized against the length of the document to yield a comparable score. This approach follows methods used in corpus linguistics, where syntactic patterns are quantified to reveal systemic tendencies (Biber, 1988).

The IDI also allows comparative analysis across domains. For instance, a hospital discharge note, a university onboarding guide, and a corporate conduct policy may all contain implicit directives. By applying the IDI, one can determine which context relies most heavily on structural obedience. Early pilot studies suggest that HR policies often reach higher IDI scores, reflecting their systematic use of conditionals to regulate employee behavior (Startari, 2025b). By contrast, clinical notes display more variation, with explicit directives still present in some contexts.

The value of the IDI lies not only in measurement but also in accountability. Institutions often present AI-generated texts as neutral, objective, and risk-free. A high IDI score reveals that such texts are saturated with hidden mandates. This undermines the claim of neutrality and supports the argument advanced by Crawford (2021), who shows that claims of objectivity in AI often mask normative decisions. The IDI provides a way to demonstrate, with quantifiable evidence, that bureaucratic texts produced by LLMs are not neutral but directive.

From a theoretical perspective, the IDI aligns with the framework of the compiled rule. If the compiled rule is the generative mechanism that encodes command in syntax, the IDI is the tool that measures its manifestation. It translates the abstract concept of structural obedience into a practical methodology. By linking formal grammar to empirical measurement, the IDI bridges linguistic theory and institutional critique.

The introduction of such a metric also raises important questions. What threshold of implicit directives is acceptable in a bureaucratic document? Should institutions disclose

when their texts contain high levels of structural coercion? Could regulatory bodies use the IDI to audit AI-generated policies? These questions remain open, but they illustrate the broader significance of developing a measure of hidden mandates.

Ultimately, the Implicit Directive Index is not just a technical tool but a lens for understanding governance in the age of AI. It exposes the silent mandates that structure compliance, reveals the infrastructural nature of grammar, and challenges institutions to recognize that authority persists even when the imperative disappears. The IDI demonstrates that what appears neutral is often deeply directive, and that obedience can be measured, not only observed.

7. Conclusion – Silent Governance and Structural Obedience

The analysis of AI-generated bureaucratic texts across healthcare, education, and corporate contexts demonstrates a consistent pattern: explicit commands are replaced by grammatical structures that function as mandates while presenting themselves as neutral statements. Hospitals deploy Epic Scribe to produce clinical notes that compel compliance through conditionals and causal constructions. Universities distribute onboarding documents in which participation is enforced through descriptive consequences. Human resources departments circulate conduct policies where access, rights, and responsibilities are framed as outcomes of conditional clauses. In every case, what appears as information or description operates as command.

The central concept that makes sense of this transformation is the *compiled rule*. As Startari (2025a) has argued, the compiled rule corresponds to type-0 generative structures that function as infrastructure. These structures do not carry meaning in the conventional sense; rather, they encode transformation. Within bureaucratic language, the compiled rule allows subordination to function as directive. It is not the lexical item “must” that obliges, but the form of the sentence that makes non-compliance structurally impossible. The agent of command disappears, yet obedience is preserved.

This process produces what has been called *structural obedience*. The term designates a form of compliance that arises not from recognition of authority but from the organization of grammar. The subject obeys because the sentence leaves no viable alternative. In clinical notes, failure to comply leads to compromised treatment. In onboarding guides, absence from orientation leads to exclusion from platforms. In conduct policies, non-completion of training leads to suspension of access. Each text constructs a closed path where compliance is the only outcome that preserves institutional participation.

The development of the *Implicit Directive Index* (IDI) provides a way to measure this phenomenon. By quantifying the density of conditionals, causal gerunds, and consecutive clauses, the IDI reveals the extent to which a document compels obedience silently. It confirms that structural obedience is not an isolated occurrence but a systemic pattern across domains. The metric also highlights the need for accountability. A document with a high IDI score cannot be described as neutral, regardless of its absence of explicit imperatives. Its grammar is directive, and its authority is infrastructural.

The ethical and political consequences are substantial. When authority is displaced from agents to syntax, responsibility becomes diffuse. Hospitals can claim that instructions were merely descriptive. Universities can argue that onboarding texts only stated conditions. Corporations can assert that policies simply outlined consequences. Yet individuals experience these texts as binding. Compliance occurs without command, and sanction follows without explicit order. Authority has migrated into linguistic form. As Crawford (2021) reminds us, claims of neutrality in AI often conceal systems of power. The IDI demonstrates precisely how such concealment operates at the level of grammar.

From a Weberian perspective, the impersonality of bureaucracy has reached a new stage. Weber (1978) described bureaucratic authority as rational, rule-based, and impersonal. AI-generated texts intensify this impersonality by eliminating the commanding subject altogether. The rule does not appear as an order issued by a superior but as a grammatical inevitability. Bureaucratic language thus becomes self-executing, in the sense that the sentence itself governs.

The larger implication is that governance in the age of AI increasingly takes the form of silent mandates. Institutions delegate not only the drafting of documents but also the structure of authority to automated systems. What circulates is not merely information but a grammar of compliance. This grammar enforces norms while denying its own directive force. Governance becomes harder to contest because there is no explicit order to challenge. The text presents itself as neutral, while its syntax ensures obedience.

The conclusion is therefore twofold. First, AI-generated bureaucratic language must be recognized as a mode of governance in its own right. It produces authority not by declaring it but by embedding it in sentence structure. Second, accountability mechanisms must adapt to this new reality. Institutions cannot disclaim responsibility by pointing to the absence of explicit imperatives. The grammar of their documents is itself a locus of power, and the compiled rule is the mechanism through which that power is exercised.

Future research should extend this analysis by applying the IDI across larger corpora of bureaucratic texts, examining variations across sectors and jurisdictions. Comparative studies could determine whether healthcare, education, or corporate governance relies most heavily on structural obedience. Regulatory frameworks might also incorporate syntactic analysis to ensure that automated systems do not obscure the presence of mandates. Without such scrutiny, the risk is that silent governance becomes normalized, leaving authority untraceable and obedience unquestioned.

The paradox remains: language that denies being a command functions as one. The disappearance of the imperative does not signal a reduction of authority but its reconstitution in grammatical form. Silent mandates are not anomalies. They are the new infrastructure of bureaucratic governance, generated and distributed by automated systems. The compiled rule ensures that institutions can govern through syntax, and the IDI allows us to reveal this operation. The conclusion is clear. Obedience in the age of AI is no longer commanded, it is formatted.

References

Biber, D. (1988). *Variation across speech and writing*. Cambridge University Press.

Crawford, K. (2021). *Atlas of AI: Power, politics, and the planetary costs of artificial intelligence*. Yale University Press.

Foucault, M. (1977). *Discipline and punish: The birth of the prison* (A. Sheridan, Trans.). Pantheon.

Pasquale, F. (2015). *The black box society: The secret algorithms that control money and information*. Harvard University Press.

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

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

Startari, A. V. (2025c). *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. (2025d). *Ethos without source: Algorithmic identity and the simulation of credibility*. SSRN Electronic Journal. https://doi.org/10.2139/ssrn.5313317

Weber, M. (1978). *Economy and society: An outline of interpretive sociology* (G. Roth & C. Wittich, Eds.). University of California Press.

Zuboff, S. (2019). *The age of surveillance capitalism: The fight for a human future at the new frontier of power*. PublicAffairs.