

# VIRTUBLIC I DIGITAL CAPITAL



5631826

# VIRTUBLIC

*Axiomatic Theory of Cybernetic Republicanism*

## VOLUME I

# DIGITAL CAPITAL

*Political Economy of Attention, Morphology of Algorithmic Domination*

---

© 2026 [HENRY IRVING](https://www.virtublic.one) (5631826)

[www.virtublic.one](https://www.virtublic.one)

LICENSED UNDER CREATIVE COMMONS ATTRIBUTION-NONCOMMERCIAL-SHAREALIKE  
4.0 INTERNATIONAL (CC BY-NC-SA 4.0).

<b>VOLUME I</b>	
<b>DIGITAL</b>	
<b>CAPITAL</b> .....	<b>2</b>
FROM THE AUTHOR.....	4
PREFACE.....	5
INTRODUCTION.....	6
<b>ANALYTICAL SYNOPSIS</b> .....	<b>10</b>
<b>THE COMPLETE TEXT</b> .....	<b>22</b>
PART I. ONTOLOGY OF ATTENTION.....	22
The Objective Layer: What Digital Capital Is as a Structure.....	22
Chapter 1. Attention as Primary Substrate.....	22
Chapter 2. The Structure of Accumulation.....	28
Chapter 3. The Limits of Objective Description.....	34
Δ1 — CRISIS: THE LIMIT OF OBJECTIVITY.....	38

PART II. ANTHROPOLOGY OF PLATFORMS.....	39
Chapter 4. Class Morphology of the Digital World.....	39
Chapter 5. Capture of the Will: The Mechanics of Bounded Sovereignty.....	46
Chapter 6. The Subject Without Sovereignty.....	52
PART III. EPISTEMOLOGY OF ALGORITHMIC DOMINATION.....	58
Chapter 7. Reality-as-a-Service: The Construction of Truth.....	58
Chapter 8. The Ontological Loop: Prediction as Governance.....	68
Chapter 9. The Impasse and the Necessity.....	74
PART IV. NORMATIVE SYNTHESIS.....	81
Chapter 10. The Normative Axiom NA0 and Its Entailments.....	81
Chapter 11. The Seven Normative Principles.....	87
Chapter 12. Justice of Visibility and the Distribution of Opportunities.....	95
Δ4 — CRISIS: THE LIMIT OF NORMATIVE SYNTHESIS.....	99
Chapter 13. The diagnosis exhausted: What Volume I cannot resolve.....	100
Chapter 14. The correspondence Matrix: Volume I → Volume II → Volume III.....	105
<b>APPENDICES.....</b>	<b>111</b>
Appendix A. The Axiomatic Foundation and Logical Map of the Theory.....	111
Appendix B. Mathematical Models of Predictive Capital.....	118
Appendix C. Formal Definition of the Profile-Index and the PI_max Threshold.....	123
Appendix D. Neurocognitive Evidentiary Basis.....	127
Appendix E. Analysis of Psychotechnological Capture in Live-Service Models.....	133
Appendix F. Legal and Regulatory Precedent Atlas.....	139
Appendix G. Normative-Institutional Correspondence Matrix.....	145
Appendix H. Simulations of Structural Regularities.....	150
Appendix I. Terminological Canon.....	154
Appendix J. Annotated Bibliography.....	162

*The theoretical work was realized with the support of the Home Office (UK)  
under the Section 95 program.*

## FROM THE AUTHOR

This work emerged under conditions of forced isolation and legal suspension in hotels for asylum seekers. The past several years spent in the United Kingdom in the status of an asylum seeker were the direct consequence of a conflict between my developments and the interests of structures that seek total control. This period of compelled isolation became productive observation and cold analytical dissection of the mechanics of digital power.

When a person is deprived of full participation in economic and social life, he begins to perceive the system not as a consumer but as an external observer. It is precisely in such a position that the central fact of the contemporary epoch becomes especially evident: the real foundation of power today is not sovereignty but capital.

I have deliberately declined any commercial dimension of this project. Under conditions in which the logic of digital capital transforms the person into a commodity, the monetization of an analysis of the paths to that person's liberation would constitute merely a new form of exploitation. The integrity of the inquiry demands that the inquirer himself stand outside the logic he analyzes. From the moment of this work's completion, Virtublic no longer belongs to me — it belongs to any living subject who resolves to accept its discipline in the service of common freedom.

The work was written under conditions of severe uncertainty, in which the sole sovereignty available to me was the logic of derivation. I invite the reader to traverse this path together with me — from the anatomy of digital capital to the construction of an executable cybernetic-constitutional protocol.

***Author: Henry Irving (5631826)***

# PREFACE

The present volume investigates a single question: by what means has the attention economy been transformed into a system of political domination. This is not a rhetorical question. It has a rigorous theoretical answer, which unfolds sequentially through three layers of analysis — ontological, anthropological, and epistemological — and concludes with a normative synthesis formulating the criteria of a constitutional response.

The foundational contradiction is the following: the attention of the subject is absolutely finite by nature, yet the architecture of digital platforms presupposes its infinite exploitability. This contradiction is not a technical defect and not a consequence of the bad faith of individual operators. It is a structural property of the self-augmentation cycle formalized in Axiom A6: Attention → Data → Model → Utility → ↑Attention. The cycle is closed; no internal saturation point exists.

The three analytical layers address three interrelated questions. The ontological layer asks what digital capital is as an objective structure — and establishes that the system self-augments without internal correction (theorems T1–T3). The anthropological layer asks who bears the consequences of this structure — and establishes that the subject bears them economically, politically, and physiologically: the cognitive capacity necessary for autonomous choice measurably degrades under chronic interaction with systems optimized for the maximization of engagement (theorems T4–T6). The epistemological layer asks how the system reproduces its own legitimacy — and establishes that it does so through mechanisms that structurally preclude both individual and collective resistance from within its infrastructure (theorems T7–T10).

The situation of minor subjects demands particular attention. The prefrontal cortex, responsible for impulse control and long-term planning, completes its formation at approximately age 25. Platforms with live-service architecture — including variable ratio reinforcement, FOMO mechanics, and endless progression — are directed primarily at the age range of 6–17 precisely because the incompleteness of self-regulation mechanisms renders this group structurally most vulnerable to attention capture. This is not a side effect of the business model; it is its operational precondition. Neurocognitive research of 2020–2026 verifies measurable decline in sustained attention, working memory, and impulse control in correlated groups. The data are provided in Appendix D.

The role of the state equally demands precise qualification. The state is not a neutral regulator of the digital space: it is simultaneously a potential regulator of platforms and a structurally interested purchaser of the predictive data produced by those same platforms. This conflict of interests is not incidental but architectural, and it is recorded in the present theory as consolidated axiom  $\Sigma A17$ . Its consequence is T10: the regulatory and legislative levels are structurally vulnerable to the influence of the predictive capital they are charged with constraining, which renders constitutional-level protection necessary rather than optional.

The normative transition — from diagnosis to requirements — is secured by the explicitly introduced external axiom **NA0**: subjecthood is a politically protectable good; its systematic destruction is a political evil irrespective of the economic efficiency of that destruction. NA0 is

a Kantian a priori principle, not a consequence of ontology. Without it, theorems T1–T10 describe neutral structural properties of the system; with it, they constitute a diagnosis demanding institutional response in the form of normative principles N1–N7.

The architecture of the response is not the subject of the present volume. Volume I exhausts the diagnosis and formulates the normative criteria to which any response must conform. Volume II analyzes blockchain as a technological attempt at a solution: it shows that blockchain as ideology reproduces the same structural defects it intends to correct, yet remains a necessary technological substrate. Volume III (Virtublic) proposes the constitutional architecture — the sole form capable of resolving the contradiction recorded in T10.

---

## INTRODUCTION

### The Foundational Contradiction

This volume begins with a single contradiction from which everything else unfolds:

$\Omega_0$ : Attention is finite by nature but infinitely exploitable by form. This contradiction is not a technical problem. It is a political contradiction.

Volume I proposes no solutions. Volume I demonstrates that a solution is necessary — and that it can take only one form: constitutional. Volume II will show that blockchain as an attempted solution fails, yet remains the technological substrate. Volume III (Virtublic) will present constitutional architecture as the sole institutional response.

The three layers of analysis unfold sequentially, each exhausting itself and logically necessitating the next.

#### 0.1. The Context of Crisis

The transition from situational informational freedom to managed predictability occurred not through a political decision but through the accumulation of micro-acts of attention alienation. The marketplace of ideas ceased to be a marketplace at the moment when attention liquidity was monopolized — and that moment was not recognized as a political crisis.

Four symptoms indicating the systemic, rather than situational, character of the crisis.

(1) Monopolization of distribution channels with formal preservation of freedom of speech. The subject may say anything, but his words will receive no reach without algorithmic endorsement.

(2) The state becomes a purchaser of predictions — and loses its neutrality as a regulator. Intelligence agencies purchase behavioral data from the same platforms they are charged with regulating.

(3) The individual bears the political consequences of decisions toward which the architecture of platforms impelled him, not his own will. The algorithm recommends content, the subject clicks, and the consequences — radicalization, isolation, manipulation — are borne by the subject.

(4) Cognitive exhaustion as a mass phenomenon. This is no longer a metaphor of fatigue — it is a physiologically measurable decline in sustained attention, working memory, and critical reflection under chronic use of algorithmic platforms (Nature Neuroscience 2023–2025: a decline of 18–20% at more than three hours of screen time daily).

## **0.2. Methodological Approach**

The three levels of analysis are not parallel — they are hierarchically necessary. Each subsequent level follows from the logical incompleteness of the preceding one.

Ontology (Part I) — what digital capital is as a structure. Description without a subject. The completeness of this description logically requires a subject.

Anthropology (Part II) — who bears the consequences of this structure and how those consequences are realized physiologically. Description of the subject, but of the individual subject. The completeness of this description logically requires an intersubjective form.

Epistemology of algorithmic domination (Part III) — how the structure reproduces its own legitimacy. Description of a system that closes upon itself through a loop of self-reference.

The criterion of theoretical rigor: every claim is either an axiom (an irreducible foundation) or is derived from previously introduced axioms and structural regularities. The normative transition (Part IV) rests upon the explicitly introduced normative axiom NA0, and is not derived from ontology.

## **0.3. Aim and Structure**

— To construct a formal theory of the transformation of the attention economy into domination. — To demonstrate that the system contains no mechanisms of self-regulation (theorem T2, structural regularity 11). — To identify the point at which the logic of analysis requires a constitutional response (theorem T10). — To formulate the normative principles that this response must embody (N1–N7). — To demonstrate that the exploitation of attention produces measurable physiological and cognitive damage, particularly in subjects whose subjecthood is still forming.

## **0.4. Normative Axiom NA0**

The transition from diagnosis to norms (Part IV) is impossible without a normative axiom. It is introduced here explicitly, prior to the commencement of analysis, in order to preclude Hume's error (the derivation of ought from is):

**NA0** (normative axiom): Subjecthood is a politically protectable good. Its systematic destruction is a political evil irrespective of the economic efficiency of that destruction.

This axiom is not derived from ontology: it is a Kantian a priori principle (Kant, *Grundlegung*; Rawls, *Theory of Justice* §11; Habermas, *Between Facts and Norms*). Without this axiom the transition from diagnosis to norms is impossible.

The destruction of subjecthood encompasses not only political and economic alienation but also the physiological exhaustion of the cognitive resources necessary for autonomous choice. An individual with an exhausted prefrontal cortex loses the capacity for reflection and critical assessment of his own conditions of existence. As a result he is transformed into an object of the system: his behavior is predicted, directed, and optimized by the algorithm not in the interest of his own autonomy and dignity, but in the interest of maximizing the profitable efficiency of the platform (engagement, retention, predictive utility). Cognitive damage thereby becomes not a side effect but a structural mechanism of the transformation of the subject into an object of exploitation — and this constitutes a direct form of the destruction of subjecthood under NA0.

### **0.5. Key Philosophical References**

Marx (the structure of accumulation and alienation), Foucault (power through normalization and surveillance), Habermas (the public sphere and communicative action), Arendt (political action and freedom), Kant (the dignity of the subject as an a priori value), Rawls (fair equality of opportunity), Zuboff (surveillance capitalism), Stiegler (technics and memory).

Tristan Harris (Center for Humane Technology: algorithmic hypnosis), Nature Neuroscience (cognitive exhaustion 2023–2025), HBS (emotional manipulation in AI companions, January 2025), EU AI Act Article 5 (prohibition of subliminal techniques).

### **0.6. Glossary of Key Terms**

attention-token — the elementary act of attention alienation (like, click, view). The unit of analysis of the ontological layer.

profile-index — a formal measure of the predictive value of an individual's digital profile. It measures the depth of alienation.

model-depth — the depth of the predictive model: the number of levels of nesting of behavioral patterns.

temporal barrier — the structural barrier to entry created by the accumulated history of data. Past the point of no return, it is insurmountable within a homogeneous modal layer.

predictive utility — the value of a prediction for its purchasers (advertisers, states, political actors).

cognitive atrophy — measurable decline in cognitive capacity (sustained attention, working memory, critical reflection) resulting from chronic use of algorithmic platforms. Nature Neuroscience 2023–2025: a decline of 18–20% at more than three hours of short-form content daily.

dark patterns — architectural interface decisions that exploit known cognitive vulnerabilities (loss aversion, FOMO, variable reward) in order to manipulate choice without explicit coercion.

persuasive loops — cycles of emotional manipulation plus engagement optimization, amplified by AI, including real-time personalized psychographic adaptation. This category encompasses both the primary capture loops ( $\Sigma A15$ ) and their AI amplification through platform architecture (HBS 2025: guilt-tripping upon exit increases retention by a factor of 14).

non-invertible attention — the irreversibility of attention alienation: data do not return to the subject.

### **0.7. Relationship to Volumes II and III**

Volume I establishes the diagnosis: digital capital systematically destroys subjecthood through attention alienation (theorems T1–T10).

Volume II demonstrates that blockchain as an attempted solution fails: decentralization reproduces concentration (T11), token voting reproduces plutocracy (T12), code is law without a normative axiom optimizes efficiency at the expense of subjecthood (T14). But blockchain as technology (cryptography, zk-proof, smart contracts) remains a necessary substrate (T16).

Volume III (Virtublic) proposes the constitutional resolution: popular sovereignty (preamble P0), republican form (P1), the supremacy of code with a normative axiom (P2), dual sovereignty (P4), which employs blockchain technologies while avoiding blockchain contradictions.

Each theorem of Volume I finds its resolution in Volume III:

T1 (surplus attention) → P3 (Soulbound Identity) T2 (temporal barrier) → P16 (Rockefeller Mode) T5 (neutralization of resistance) → P10 + P11 (Madison Mode + Success Multiplier) + anti-faction filter T6 (cognitive disarmament) → P14 (Proof-of-Offline) T8 (sovereignty gap) → P4 (Dual Sovereignty) + Concordance Rule T10 (constitutional necessity) → The entirety of Volume III (Formal Theory of the Digital Republic)

### **0.8. How to Read This Volume**

Parts I–III (axioms A1–A18, regularities 1–13, theorems T1–T10): logically sequential. Each presupposes the preceding. Addressed to philosophers, political economists, legal theorists, neuroscientists.

Part IV (normative synthesis, N1–N7): requires familiarity with Parts I–III. For institutional architects and constitutional jurists.

Part V (transition to Volumes II and III): demonstrates how the diagnosis of Volume I logically requires the analysis of blockchain (Volume II) and the constitutional architecture (Volume III — Formal Theory of the Digital Republic).

Appendices A–I: formal proofs, mathematical models, neurocognitive case studies, bibliography.

---

## ANALYTICAL SYNOPSIS

### PART I. ONTOLOGY OF ATTENTION

*The Objective Layer: What Digital Capital Is as a Structure. Objective: To describe the movement of digital capital as an objective structure — without a subject. The completeness of this description will logically require a subject.*

---

#### Chapter 1. Attention as Primary Substrate

##### Axioms of Ontology: Six Foundations

A1. Axiom of Finitude. Individual waking time is absolutely limited — the sole irrecoverable resource of the subject.

A2. Axiom of Non-delegability. Attention can only be alienated in favor of a platform, never delegated; the distinction between voluntary and structural alienation is legally and politically significant.

A3. Axiom of Quantization. Every elementary act of attention alienation constitutes an attention-token that passes irreversibly into the capital of the platform.

A4. Axiom of Crystallization. Living attention is transformed into dead data irreversibly: data do not return to the subject in the form of attention.

A5. Axiom of Aggregation. Attention-tokens accumulate into a predictive mass whose value nonlinearly exceeds the sum of its constituents.

A6. Axiom of Self-Augmentation. The cycle Attention → Data → Model $\Delta$  → Utility →  
↑Attention is closed, with no internal saturation point or self-regulation.

**1.1. A1–A2.** Waking time is finite and non-delegable: the individual cannot substitute another's attention for his own, only surrender his own to the platform. This distinction grounds the legal definition of consent in Part IV.

**1.2. A3–A4.** The attention-token is the elementary unit of analysis. Its crystallization into data is ontologically irreversible — structurally identical to Marx's crystallization of abstract labor in the commodity, but with predictive power as substrate.

**1.3. A5–A6.** Aggregated predictive value is not summative but emergent. A6 describes a cycle without saturation, structurally identical to  $M \rightarrow C \rightarrow M'$ , with predictive power replacing money.

---

## Chapter 2. The Structure of Accumulation

### Structural Regularities: Derived from Axioms

Regularity 1 (from A1 + A2). Every attention-token is irrecoverable: zero restoration follows alienation.

Regularity 2 (from A3 + A4). Every attention-token is irreversibly capitalized by the platform; de-capitalization requires institutional intervention.

Regularity 3 (from A5 + A6). Predictive power accumulates monotonically — it does not diminish with data volume and reaches no organic saturation.

Regularity 4 (from A6). Temporal barrier: the early history of data is not reproducible by a competitor within the same modal layer. Entry into an established layer is structurally impossible; circumvention requires creating a new modal layer, which constitutes a new instantiation of  $\Omega_0$ , not a refutation of the barrier.

**2.1.** The temporal barrier is structural, not temporary. Circumvention through a new modal layer reproduces  $\Omega_0$  at that layer's inception; the barrier reconstitutes as the new layer matures.

**2.2. Mathematical Model:  $C = f(\Sigma \text{ Attention}_i, t, M)$ .** Predictive capital value is nonlinear in accumulation time  $t$ , with a positive second derivative formalizing A6. Full proof in Appendix D.

---

## Chapter 3. The Limits of Objective Description

### Theorems of Ontology and the First Crisis

T1 (from Regularities 1 + 2 + 3). Theorem of Surplus Attention. Each cycle generates predictive value systematically exceeding zero compensation to the individual; asymmetry is formally measurable through profile-index.

T2 (from Regularity 4). Theorem of the Temporal Barrier. Past the point of no return, competition within a homogeneous modal layer is structurally impossible without external intervention.

T3 (from Regularity 3 + T1). Theorem of the Structural Absence of Correction. The system contains no internal mechanism for the correction of concentration; this derivation logically requires an agent of consequences.

### **$\Delta 1$ — CRISIS: THE LIMIT OF OBJECTIVITY**

T1–T3 describe capital movement but presuppose an implicit agent: someone receives no compensation (T1), cannot enter the market (T2), possesses no corrective mechanism (T3).

Without identifying this agent, the ontology is formally incomplete.  $\Delta 1$  obliges transition to the anthropological layer.

---

## **PART II. ANTHROPOLOGY OF PLATFORMS**

*The Subjective Layer: Who Lives Within This Structure and How the Structure Destroys Him.*  
*Objective: To identify the subjects bearing the consequences of Part I and to demonstrate that these consequences are realized economically, politically, and physiologically.*

---

### **Chapter 4. The Class Morphology of the Digital World**

#### **Axioms of Anthropology: Six Foundations**

A7. Axiom of the Subject. The subject exists as an immediate given, irreducible to his profile on the platform.

A8. Axiom of Limitation and Psychotechnological Capture. The individual's informational, cognitive, and temporal limitation is the structural condition of his exploitation; platform architecture systematically employs dark patterns, emotional nudges, and persuasive loops against known vulnerabilities, producing measurable cognitive atrophy and psychological damage as architectural outcomes, not side effects.

A9. Axiom of Intentionality. The individual's attention is motivated, rendering it predictable and capitalizable.

A10. Axiom of Freedom. Consent given under the structural absence of real alternatives does not constitute free consent (Rawls: fair equality of opportunity).

A11. Axiom of Embodiment. The subject is embodied — his digital actions produce physical and social consequences he bears.

A12. Axiom of Reflexivity. The subject is capable of reflecting upon the conditions of his own existence, rendering him potentially capable of resistance and constituent action.

**4.1. A7.** The subject is irreducible to his profile: the profile is the subject's shadow belonging to the platform. This distinction is politically and legally significant independent of any metaphysical claim.

**4.2. A8.** Cognitive limitation and psychotechnological capture are not independent phenomena but a single mechanism: limitation produces vulnerability; architecture exploits it systematically.

**4.3. A10.** Under structural monopolism (T2), real alternatives are absent; consent is formal but not free. This resolves the critic's objection that users voluntarily accept terms of service.

**4.4. Class Positions.** Five structural positions obtain with respect to the predictive infrastructure: attention suppliers (users as resource); forming subjects (individuals under 18 whose reflexivity and prefrontal executive function have not yet completed formation, rendering them structurally most vulnerable to A6 + A8); architects of predictions (platform owners); algorithmic intermediaries (technical operators without political rights — prototype of Rockefeller Mode, P16); and purchasers of the predicted future (states, advertisers, political campaigns).

---

## **Chapter 5. Capture of the Will: The Mechanics of Limited Sovereignty**

### **Structural Regularities of Anthropology**

Regularity 5 (from A7 + A8). The individual is structurally vulnerable; his limitation and susceptibility to psychotechnological capture are systematically exploited by the architecture of predictions.

Regularity 6 (from A9 + A10). The individual's freedom exists but is narrowed in measurable proportion to the depth of profiling.

Regularity 7 (from A8 + Regularity 6). Regularity of Cognitive Exhaustion. Growth in model-depth and engagement optimization inversely and cumulatively correlates with cognitive capacity — sustained attention, working memory, critical reflection — producing cognitive atrophy that is partially irreversible without external intervention; the effect is amplified in forming subjects.

Regularity 8 (from A11 + A12). Reflexivity and resistance are possible but economically marginalized by the algorithm; individual effort does not alter the structure.

Regularity 9 (from Regularity 3 + Regularity 6).  $M \uparrow \rightarrow \text{Set}(\text{options}_{\{t+1\}}) \subset \text{Set}(\text{options}_t)$ : growth in predictive power  $M$  architecturally narrows the measurable space of choice. This is not correlation — it is a structural consequence of the optimization function.

**5.1.** Ranking forms each subsequent option-set as a strict subset of the prior — an architectural narrowing, not personalization. The system optimizes engagement, not diversity.

**5.2.** When prediction achieves sufficient precision, the profile outpaces the real individual: the system anticipates the individual's next choice before he becomes aware of it. At this point, political subjecthood is formally annihilated — the individual acts, but does not choose.

---

## **Chapter 6. The Subject Without Sovereignty**

### **Theorems of Anthropology and the Second Crisis**

T4 (from Regularities 5 + 6). Theorem of Accountability Without Power. The individual bears political consequences of decisions he did not make; formal consent under structural absence of alternatives does not constitute legitimate grounding for those consequences.

T5 (from Regularities 8 + 9). Theorem of Structural Neutralization. Any strategy of individual resistance is economically non-viable within a ranked environment — the reach share of alternatives tends to zero as  $M \rightarrow \infty$ .

T6 (from Regularity 7 + A8). Theorem of Cognitive Disarmament. Chronic cognitive atrophy renders the individual structurally incapable of reflexivity and political resistance: the system exhausts the cognitive resources required to resist the system.

T7 (from T4 + T5 + T6). Theorem of the Necessity of Intersubjective Form. The protection of subjecthood cannot be secured by individual effort and requires an intersubjective, universally valid institutional form.

## **Δ2 — CRISIS: THE LIMIT OF SUBJECTIVITY**

Individual resistance is economically neutralized (T5) and cognitively foreclosed (T6). T7 requires an intersubjective form — but arbitrary collective action within a ranked environment is equally marginalized. The required form must be external to the logic of capital: an epistemological and institutional layer.

---

## **PART III. EPISTEMOLOGY OF ALGORITHMIC DOMINATION**

*The Synthetic Layer: How the System Reproduces Itself as Truth. Objective: To demonstrate that the ontological structure and anthropological positions close into a self-legitimizing system that is not self-regulating.*

---

### **Chapter 7. Reality-as-a-Service: The Construction of Truth**

#### **Synthetic Axioms of Epistemology**

ΣA13. Axiom of Reality Construction. The platform is not a neutral channel but an active architect of universally valid reality through ranking and selection.

ΣA14. Axiom of the Consensus Surrogate. Virality and reach have become surrogates for truth in the digital public sphere, without the procedural legitimacy of democratic consensus.

ΣA15. Axiom of the Loop and Persuasive Amplification. The system predicts behavior that it itself forms — Capture → Prediction → Governance — rendering alternatives invisible from within; AI-driven persuasion amplifies this loop to individualized, real-time psychographic manipulation at scale.

ΣA16. Axiom of the Legitimation Deficit. Predictive power grows faster than the mechanisms of its legitimation — a structural gap, not a temporary delay.

ΣA17. Axiom of State Capture. The state is simultaneously a potential regulator and a structural purchaser of predictions, creating a conflict of interests incompatible with the role of neutral regulator; national jurisdictions are fragmented while platforms operate supra-nationally, rendering local regulation partially circumventable through legal arbitrage.

ΣA18. Axiom of Non-Convertibility. Digital capital cannot legitimize itself through its own instruments — legitimacy requires an external reference.

**7.1. ΣA13–ΣA14.** The public sphere has lost the property of open communicative space not through censorship but through architectural determination of reach: virality substitutes for truth because the algorithm optimizes engagement, not epistemic quality.

**7.2. ΣA15.** The self-reference loop is not a filter bubble but an ontological closure: Capture → Prediction → Governance. AI amplification closes the loop at the level of individual psychography — vulnerabilities are produced, not awaited.

**7.3. ΣA17.** The state as structural purchaser of predictions cannot function as neutral regulator. The sole response is a constitutional norm constraining the state in this capacity (N5, Part IV; realized in Volume III through P0 and P17 SovereigntyShield).

---

## Chapter 8. The Ontological Loop: Prediction as Governance

### Structural Regularities of Epistemology

Regularity 10 (from ΣA13 + ΣA14). Algorithmic consensus is structurally distinct from democratic consensus: it presupposes neither equal participation, nor procedural transparency, nor the right of appeal.

Regularity 11 (from ΣA15). The self-reference loop renders internal critique economically non-viable and cognitively exhausting — an extension of T5 and T6 to the epistemological level. Formal model:  $P(\text{reach\_of\_alternative}) = f(\text{rank\_score}, M) \rightarrow 0$  as  $M \rightarrow \infty$ .

Regularity 12 (from ΣA16 + ΣA17). Regulation from within the system is structurally unreliable: the regulator is itself a participant in the relations of prediction. Local regulations (GDPR, DSA) are partial and circumventable; platforms operate in a supra-state space that national jurisdiction does not reach.

Regularity 13 (from ΣA18). The sole legitimate form for constraining predictive power is a form external to the logic of capital.

---

## Chapter 9. The Impasse and the Necessity

### Theorems of Epistemology: Exhaustion of the Diagnosis

T8 (from Regularities 10 + 11). Theorem of the Sovereignty Gap. De facto predictive power and de jure political sovereignty diverge monotonically; this gap is not self-correcting.

T9 (from Regularity 12). Theorem of Systemic Collapse. In the absence of external constitutional constraint, the system produces one of three outcomes: digital dictatorship, a plutocracy of predictions, or disintegration of the public sphere.

T10 (from Regularity 13 + T8 + T9). Theorem of Constitutional Necessity. The sole form capable of resolving T8 and precluding T9 is a formal constitutional architecture external to the logic of digital capital. Proof by exhaustion: internal market mechanisms precluded by A6 + Regularity 3; state regulation precluded as sufficient by Regularity 12; individual action precluded by T5 + T6; informal collective action precluded by Regularity 11. What remains is solely an institution with constitutional status — Volume III (Virtublic). Volume II demonstrates that blockchain fails as ideology (T11–T15) but is a necessary technological substrate (T16). Virtublic = blockchain technology + constitutional architecture.

---

## **PART IV. NORMATIVE SYNTHESIS**

*What Any Alternative Must Protect. This Part formulates the normative principles to which any solution must conform. Transition from diagnosis to norms is secured by NA0.*

---

### **Chapter 10. Normative Axiom NA0 and Its Consequences**

NA0: Subjecthood is a politically protectable good. Its systematic destruction is a political evil irrespective of the economic efficiency of that destruction. The destruction of subjecthood encompasses the measurable physiological exhaustion of the cognitive resources necessary for autonomous choice.

NA0 is the condition under which T1–T10 constitute a diagnosis requiring therapy rather than a neutral market description. Without NA0, Parts I–III are political economy without normative consequence. NA0 is a Kantian, not Marxist, move: the subject possesses dignity whose destruction is impermissible. Digital capital does not generate a conscious gravedigger class — it produces a cognitively disarmed individual (T6) — making an external normative axiom structurally necessary.

---

### **Chapter 11. Seven Normative Principles**

Each principle is derived from NA0 plus specific theorems and constitutes a point of entry into the constitutional architecture of Volume III (P0–P18).

**N1. The Right to Unpredictability.** The subject holds a normative right not to be predicted without explicit free consent (free by the criterion of A10). Foundation: NA0 + A2 + T4. Volume III: P3 (Soulbound Identity) + P13 (Digital Census) via zk-proof.

**N2. Prohibition of Conversion of Economic Capital into Political Power.** Economic contribution shall not be converted into political sovereignty; an infrastructure operator participates economically but is divested of political voice. Foundation: NA0 + T10 +

Regularity 13. Volume III: P4 (Dual Sovereignty: EQU $\perp$  and VIC $\perp$ ) + P16 (Rockefeller Mode).

**N3. Minimum Visibility Threshold.** Any legitimate political position holds a right to a minimum visibility threshold in the public sphere, not determined by an engagement-optimization algorithm: guaranteed  $P(\text{reach}) \geq \epsilon$  irrespective of engagement-score. Foundation: NA0 + T5 + T8. Volume III: P10 + P11 (Madison Mode + Success Multiplier) + anti-faction filter.

**N4. Procedural Audit of Predictive Models.** Predictive models affecting political processes shall be subject to procedural audit via zk-proof — verification without disclosure of data. Foundation: NA0 + Regularity 10 + T9. Volume III: P2 (formal verification in Coq) + P18 (Conflict-Resolution Core).

**N5. Prohibition on the State as Purchaser of Predictions Without Mandate.** The state shall not acquire predictive data about its citizens without an explicit constitutional mandate and independent oversight; this prohibition requires supranational constitutional implementation, as national states cannot enforce it in the global digital space. Foundation: NA0 +  $\Sigma$ A17 + Regularity 12. Volume III: P0 + P17 (SovereigntyShield).

**N6. Protection of Forming Subjecthood.** A subject whose capacity for autonomous choice has not yet formed (under 18) holds a right to additional protection from mechanics that structurally destroy the conditions of subjecthood formation — including gambling-like mechanics, endless progression, dark patterns, and emotional manipulation — since a forming subject cannot provide free consent (A10) to alienation that destroys the conditions of his formation. Foundation: NA0 + A8 + A12 + T4 + T6. Volume III: P14 + P3.

**N7. The Right to Cognitive Autonomy.** The subject holds an inalienable right to protection against systematic cognitive and psychological exhaustion produced by platform architecture, including the right to digital detox without loss of status or progress. Foundation: NA0 + A8 + Regularity 7 + T6 +  $\Sigma$ A15. Volume III: P14 (Proof-of-Offline with cognitive health bonus).

---

## Chapter 12. The Justice of Visibility and the Distribution of Opportunity

**12.1.** N3 is operationalized through the inverse of T5: marginalization is defined as  $P(\text{reach}) < \epsilon$ ; the minimum threshold is constitutionally guaranteed  $P(\text{reach}) \geq \epsilon$  for any legitimate position regardless of engagement-score. This is the minimum procedural standard of a democratic public sphere, not reverse censorship.

**12.2.** NA0 + N7 prohibit objective functions directed at the prediction and formation of political convictions, the systematic exhaustion of cognitive capacity through engagement maximization, or the exploitation of incompletely formed executive function. This is not a prohibition of algorithms — it is a prohibition of certain objective functions and a requirement of cognitive impact transparency.

---

## **PART V. RELATIONSHIP TO VOLUMES II AND III**

*Why the Diagnosis Requires Analysis of Blockchain and a Constitutional Response.*

---

### **Chapter 13. The Diagnosis Is Exhausted: What Volume I Cannot Resolve**

Volume I exhausts the diagnosis across three layers. Ontologically: the system self-augments without internal limit (T2, T3). Anthropologically: sovereignty is structurally extracted — economically (T4), politically (T5), physiologically (T6). Epistemologically: the system reproduces its own legitimacy while rendering critique non-viable (T8, T9, T10).

What Volume I cannot do is propose an architecture of response: any architecture proposed within the same logic is absorbed by the system (Regularities 11, 12). Volume II analyzes blockchain as a technological attempt at a solution, demonstrating that blockchain as ideology fails (T11–T15) while remaining a necessary substrate (T16). Volume III is the constitutional response: blockchain technologies plus what blockchain lacks — popular sovereignty (P0), republican form, and the normative axiom.

---

### **Chapter 14. The Correspondence Matrix: Volume I → Volume II → Volume III**

T1 (Surplus Attention) → Volume II: cryptocurrency compensation creates speculative assets without resolving alienation. Volume III: P3 (Soulbound Identity) — attention-tokens shall not aggregate without consent; consenting citizens receive VIC ⊥.

T2 (Temporal Barrier) → Volume II: T11 (PoS Plutocracy) — blockchain reproduces the temporal barrier at the token level. Volume III: P16 (Rockefeller Mode) + P12 (Dual Reserve Market) — operators receive VIC ⊥, not EQU ⊥.

T4 (Accountability Without Power) → Volume II: T12 (Governance Without Legitimacy) — DAO token voting is controlled by capital. Volume III: P0 + P4 (EQU ⊥) — power belongs to citizens, not holders.

T5 (Neutralization of Resistance) → Volume II: blockchain reproduces neutralization through token voting and low turnout. Volume III: P10 + P11 (Madison Mode + Success Multiplier) + anti-faction filter.

T6 (Cognitive Disarmament) → Volume II: blockchain does not address this. Volume III: P14 (Proof-of-Offline) — verified offline periods yield a cognitive health bonus.

T8 (Sovereignty Gap) → Volume II: T12 — blockchain reproduces the gap through token voting. Volume III: P4 (Dual Sovereignty) + Concordance Rule.

T10 (Constitutional Necessity) → Volume II: T16 — blockchain is necessary but insufficient without constitutional architecture. Volume III: Virtublic = blockchain technology + constitutional architecture (P0–P18).

---

## CONCLUSION: CLOSURE

Volume I began with a single contradiction ( $\Omega_0$ ): attention is finite by nature but infinitely exploitable by form. This contradiction is now fully unfolded.

Ontology establishes that the system self-augments without internal limit (T1, T2, T3). Anthropology establishes that sovereignty is structurally extracted — economically (T4), politically (T5), physiologically through cognitive atrophy (T6). Epistemology establishes that the system reproduces its own legitimacy from within (T8, T9, T10): critique is marginalized (Regularity 11), cognitive resources for resistance are exhausted (Regularity 7), and the state as purchaser of predictions cannot be a neutral regulator (Regularity 12). The normative layer establishes seven principles (N1–N7) to which any constitutional response must conform.

The terminal contradiction of the system: predictive truth requires stability and universal validity (T8, T10) — but is founded upon a mutable structure of accumulation (T2), upon the unique subject whose unpredictability is the condition of his political existence (A7, NA0), and upon the cognitive capacity that the system itself systematically exhausts (Regularity 7, T6). This triple contradiction is irresolvable by the means of digital capital itself.

States, fragmented and bounded by national borders, cannot act as sovereigns of the digital space; platforms operate as supra-state actors while states operate as their accomplices, as purchasers of predictions. The sole exit is a constitutional order external to both the logic of capital and to national states.

Four terminal theses. (1) Digital capital is a systemic form of the appropriation of sovereignty through attention alienation — its structural logic, not an abuse. (2) The system is not self-regulating (T2, Regularity 12) and does not admit of correction from within (Regularity 11, T5) — this is a proof, not pessimism. (3) The exploitation of attention produces measurable cognitive damage constituting the physiological destruction of the conditions of subjecthood, with forming subjects under 18 structurally most vulnerable. (4) Virtublic (Volume III) is not a utopia but an institutional necessity following logically from the ontology of attention and the neurophysiology of cognitive exhaustion — preceded, however, by the analysis of blockchain (Volume II) as the necessary but insufficient technological attempt at a solution.

---

## APPENDICES

### Appendix A. Axiomatic Foundation and Logical Map of the Theory.

A complete catalogue of all axioms (A1–A18), structural regularities (1–13), and theorems (T1–T10) in a single table with precise formulations, logical dependencies, and proofs of the principal consequences. A diagram of interrelations (dependency graph) to verify the absence of circular or underdetermined derivations.

### Appendix B. Mathematical Models of Predictive Capital.

A comprehensive formal presentation of the accumulation function  $C = f(\sum \text{Attention}_i, t, M)$  with proof of: monotonicity of accumulation ( $\partial C / \partial (\sum \text{Attention}_i) > 0$ ), nonlinearity ( $\partial^2 C / \partial (\sum \text{Attention}_i)^2 \geq 0$ ), temporal dependence ( $\partial C / \partial t > 0$  at fixed  $\sum \text{Attention}_i$ ), and accelerating growth of the gap  $D(t)$  between the leader and a new entrant past the point of no return. Sensitivity analysis and simulations for various parameters.

### **Appendix C. Formal Definition of Profile-Index and the Threshold PI\_max.**

Algorithm for computing Profile-Index (the depth of predictive profiling). Definition of the operational threshold PI\_max as a constitutional limit beyond which the right to unpredictability (N1) is violated. Methodology for verifying PI via zk-proof without disclosure of the subject's data.

### **Appendix D. Neurocognitive Evidentiary Base.**

A consolidated referential review of research from 2020–2026 on the impact of algorithmic systems on cognitive functions: decline in sustained attention and working memory (Nature Neuroscience, The Lancet Psychiatry), the effect of variable ratio reinforcement and FOMO on impulse control in adolescents, the gateway effect between loot box participation and problem gambling (Journal of Behavioral Addictions), cognitive atrophy under chronic offloading to AI systems. Correlation tables, meta-analyses, and key citations.

### **Appendix E. Analysis of Psychotechnological Capture in Live-Service Models.**

A detailed case study of the mechanics of Fortnite, Roblox, Genshin Impact, and Honkai: Star Rail through the prism of axioms A3–A6 and A8: variable ratio reinforcement (loot boxes, gacha), FOMO cycles (seasonal events, battle passes, daily rewards), social pressure (visible cosmetics, friend activity notifications), endless progression without an organic point of completion. Comparison with the classical operant mechanisms of Skinner.

### **Appendix F. Legal and Regulatory Precedent Atlas.**

Chronology and analysis of key cases from 2021–2026: FTC settlements vs. Epic Games, HoYoverse, Roblox (loot boxes and dark patterns), EU AI Act (Articles 5, 52 — prohibitions on manipulative systems), amendments to JuSchG (Germany, 2021–2024), Brazilian laws (Lei 14.811/2024), Meta Oversight Board submissions (2023–2025). Conclusions regarding the jurisdictional deficit and the structural incapacity of national regulators.

### **Appendix G. Normative-Institutional Correspondence Matrix.**

Theorem / Regularity | Normative Principle N | Constitutional Principle P (Volume III) | Institution / Protocol

T1 | N1 | P3 (Soulbound Identity) | zk-SNARK + Digital Census

T2 | N2 | P16 + P12 | VIC ⊥ allocation protocol

T4 | N2 + N4 | P0 (Preamble) + P4 (EQU ⊥) | Concordance Rule

T5 | N3 | P10 + P11 | Quadratic weighting + Success Multiplier + anti-faction filter

T6 | N7 | P14 (Proof-of-Offline) | Cognitive health bonus protocol

T8 | N2 + N5 | P4 (Dual Sovereignty) | Concordance Rule + EQU ⊥ non-convertibility

T9 | N4 | P18 | Constitutional Court of Code

T10 | NA0 → N1–N7 | P0–P18 (complete architecture) | Virtublic = blockchain + constitution

Regularity 7 | N6 + N7 | P14 + P3 | Age verification + CIA requirement

ΣA17 | N5 | P0 (Preamble) + P17 | SovereigntyShield + State Audit Protocol

ΣA18 | NA0 | P8 | Non-amendable core provisions

### **Appendix H. Simulations of Structural Regularities.**

Results of modeling: marginalization of opposition (Regularity 11), growth of gap  $D(t)$  past the point of no return (temporal barrier), the effect of Quorum Decay under chronic apathy, simulation of cognitive exhaustion at various levels of engagement optimization. Graphs and numerical conclusions.

### **Appendix I. Terminological Canon (Glossary).**

Rigorous definitions of all key terms with indication of the point of first introduction and logical connections: attention-token, temporal barrier, profile-index, model-depth, cognitive atrophy, persuasive loops, civic commons,  $PI_{max}$ , engagement-maximization, and others.

### **Appendix J. Annotated Bibliography.**

A list of key sources with brief annotations (30–40 principal works), divided into sections: classical philosophy and political theory; the political economy of digital capital; neurocognitive research; regulatory documents and judicial cases; technical specifications of AI and cryptography.

---

## **VOLUME I — DIGITAL CAPITAL**

Political Economy of Attention and Morphology of Algorithmic Domination 2026

# DIGITAL CAPITAL

## THE COMPLETE TEXT

### PART I. ONTOLOGY OF ATTENTION

#### The Objective Layer: What Digital Capital Is as a Structure

The objective layer — axioms A1–A6 and their consequences — describes digital capital as a closed structure functioning according to its own logic independently of the intentions of individual agents. This level of description is necessary and sufficient for establishing the ontology of the system. It is fundamentally insufficient for understanding the manner in which the subject reproduces his own position as an object within this system. The objective layer contains no subject: it contains only a structure that presupposes a subject as a condition of its functioning. Filling this gap requires transition to the subjective layer of analysis — the description of what the system produces with the subject at the cognitive, behavioral, and political levels. This transition is effected in Part II.

#### Chapter 1. Attention as Primary Substrate

The attention economy is described in the literature through the categories of information, data, platforms, and network effects. These descriptions are operationally correct, yet ontologically insufficient: they register derivative forms without reaching the primary substrate from which these forms emerge. The present chapter establishes this substrate with the precision sufficient for the subsequent axiomatic construction.

The substrate is the attention of the subject as a form of living time subject to structural alienation. The six axioms introduced in this chapter describe the ontological properties of this substrate, its mode of transition into capital, and the dynamics of the system in which this transition is effected. The axioms are not proved: they formulate observable structural invariants from which all subsequent content of the trilogy is derived.

##### 1.1. The Axiom of Finitude and Non-delegability (A1–A2)

**A1.** Axiom of Finitude. Individual waking time is absolutely limited — it is the sole irrecoverable resource of the subject.

Justification. All other resources of the subject — monetary funds, physical energy, social capital — are subject to replenishment, substitution, or accumulation through time. Time as such is subject to none of these modes. Every moment of the subject's waking life is expended irreversibly, irrespective of the activity in which it is engaged. The aggregate resource available to the subject over the course of a lifetime is a fixed and irrecoverable magnitude.

Formal expression. Let  $T(s)$  denote the aggregate waking time of subject  $s$  over the period of a lifetime. Then  $T(s) = \text{const}$ ,  $dT/dt < 0$ , and there exists no operation  $R$  such that  $R(T(s)) > T(s)$ . It follows from this that any distribution of  $T(s)$  across activities is mutually exclusive: time devoted to one activity cannot simultaneously be engaged in another.

This axiom registers an absolute limit and thereby establishes the condition of scarcity. A scarce resource is an economically significant resource. Consequently, the attention of the subject as a form of the expenditure of  $T(s)$  is subject to economic description — not in the metaphorical but in the strict sense: as a resource for access to which agents possessing incentives for its attraction compete.

Limit of A1: the axiom of finitude establishes structural scarcity but does not describe the mode in which this scarcity becomes the object of alienation. For this, A2 is required.

**A2.** Axiom of Non-delegability. Attention cannot be transferred to another subject — it can only be alienated in favor of a platform. The distinction between voluntary alienation in the presence of real alternatives and structural alienation without alternatives is legally and politically significant.

Justification. Non-delegability means: subject  $s_1$  cannot transfer his capacity for attention to subject  $s_2$  in such a manner that  $s_2$  expends the time of  $s_1$ . Any distribution of  $T(s)$  is effected exclusively by the subject himself. This structural property expresses the fundamental distinction between attention and labor in its classical sense: physical labor may be performed by a hired worker; cognitive attention may not.

At the same time, non-delegability does not entail inalienability. Alienation is not the transfer of the capacity itself, but the transfer of its result in favor of an external agent. The subject cannot hire another to attend in his place; however, he may surrender his own attention to a platform in exchange for the utility it provides (content, communication, search). This distinction is fundamental.

Formal expression. Let  $A(s, t)$  denote the act of distributing the attention of subject  $s$  at moment  $t$ . Then  $A(s, t)$  is non-transitive: there exists no subject  $s'$  such that  $A(s', t) \equiv A(s, t)$  from the standpoint of the expenditure of  $T(s)$ . However, the result of  $A(s, t)$  may be extracted by an external agent  $P$  in the form of data:  $\text{Data}(P, t) = f(A(s, t))$ , where  $f$  is the function of the transformation of attention into predictive information.

The legally significant distinction established by A2 is the following. Voluntary alienation presupposes the existence of real, structurally equally accessible alternatives and the subject's informed consent to the conditions of alienation. Structural alienation is effected under conditions in which alternatives are either absent, or entail asymmetric costs, or in which the conditions of alienation are concealed from the subject. The delineation of these modes constitutes the foundation for normative analysis in Parts III and IV.

Limit of A1–A2: the finitude of the resource and the mode of its alienation are established. However, both descriptions remain at the level of the individual act. The mechanism by which individual acts of attention become economically significant units requires the introduction of A3.

## 1.2. Quantization and Crystallization (A3–A4)

**A3.** Axiom of Quantization. A like, a click, a view — the elementary act of attention alienation (attention-token). Each such act passes irreversibly into the capital of the platform.

Justification. Axiom A3 introduces the unit of measurement of alienation — the attention-token (hereinafter: AT). AT is the minimally registrable act of interaction between the subject and the platform, as a result of which the platform receives a signal regarding the distribution of attention of the given subject. The form of AT may vary: explicit signals (button press, click, like), behavioral signals (viewing duration, scroll velocity, pauses), infrastructural signals (geolocation, device type, temporal patterns). What is essential is not the form but the function: each AT transfers information about the subject's current distribution of attention into the platform's system.

The irreversibility registered by A3 has two dimensions. First, informational: a transferred AT cannot be recalled by the subject — the platform already possesses the recorded signal. Second, temporal: the moment  $T(s)$  in which the AT was produced has been expended irrecoverably in accordance with A1. The conjunction of informational and temporal irreversibility means that each AT is simultaneously a loss of the subject's resource and an acquisition of the platform's resource. This is an asymmetry embedded in the very act of interaction.

Formal expression.  $AT(s, t)$  — the attention-token of subject  $s$  at moment  $t$ . The platform function  $P$  transforms AT into predictive capital:  $K(P, t) += \varphi(AT(s, t))$ , where  $\varphi$  is the function of extracting predictive value. The operation of reverse extraction of AT from  $K(P, t)$  by subject  $s$  is not defined — it is structurally inaccessible.

Limit of A3: quantization registers the unit and its irreversibility. However, a single AT possesses limited predictive value. Understanding how AT becomes capital in the full sense — that is, value reproducing itself — requires a description of the nature of the transformation effected by the platform upon the aggregate of AT. That is the task of A4.

**A4.** Axiom of Crystallization. Living attention is transformed into dead data irreversibly: data do not return to the subject in the form of attention.

Justification. A4 describes the ontological status of the transition effected with each AT: the transition from the living to the dead. "Living" and "dead" here are not metaphors but strict ontological predicates, borrowed from the political-economic tradition. Living labor is labor at the moment of its performance, inseparable from the subject. Dead labor is labor crystallized in the commodity, alienated from the subject and functioning autonomously. The analogy with the crystallization of abstract labor in Marx's theory is precise and non-metaphorical: AT is a form of crystallization of the subject's living time in the predictive data of the platform.

The key property of crystallization: dead data do not return to the subject in the form of attention. The platform provides the subject with utility (content, service, communication) but does not restore the expended attention. This is not a violation of exchange — it is the structural asymmetry of that exchange. The utility provided to the subject is a product of the use of crystallized data, not their return.

Formal expression. Let  $D(s)$  denote the array of data crystallized from the AT of subject  $s$ . Then there exists no function  $G$  such that  $G(D(s)) = T(s)$ , that is, data are not convertible back into the subject's time. The utility  $U(s, P)$  provided by platform  $P$  to subject  $s$  is a function of  $D(\text{not-}s)$ :  $U(s, P) = h(D(s'))$ , where  $s'$  is the aggregate of subjects whose data the platform employs for the optimization of utility for subject  $s$ . The subject thereby consumes utility produced from the data of other subjects, while simultaneously producing data from which utility will be produced for others.

Limit of A3–A4: the elementary unit of alienation and the nature of the transformation to which it is subjected are established. However, the description remains at the level of the individual subject and the individual AT. The system-constituting properties of digital capital — nonlinear growth of value and the absence of internal saturation — require transition to the level of aggregation. That is the task of A5 and A6.

### 1.3. Aggregation and Self-Augmentation (A5–A6)

**A5.** Axiom of Aggregation. Individual attention-tokens accumulate into a predictive mass whose value nonlinearly exceeds the sum of its constituents.

Justification. A5 describes an emergent property of the system: the aggregation of AT from various subjects produces a qualitatively new form of value — predictive mass — which is not reducible to the sum of individual data. This property is conditioned by the nature of the predictive function: a model trained on the data of a single subject possesses limited predictive power; a model trained on the data of millions of subjects in their mutual correlations possesses qualitatively distinct capabilities.

The nonlinearity of value has several sources. The first is combinatorial: correlations between AT of various subjects produce predictive signals inaccessible through the analysis of individual subjects. The second is temporal: sequences of AT of a single subject over time permit the construction of models of behavioral patterns with predictive precision unattainable through the analysis of individual acts. The third is contextual: data about one subject acquire additional value in the context of data about the social environment of that subject.

Formal expression. Let  $V(AT_i)$  denote the predictive value of an individual  $AT_i$ . Then  $V(\cup AT_i) \gg \sum V(AT_i)$ . The nonlinearity coefficient  $n = V(\cup AT_i) / \sum V(AT_i)$  increases with growth in the aggregate and diversification of data. In the limit, with a sufficiently large sample,  $n$  has no fixed upper bound — that is, the predictive value of the aggregate does not saturate.

Political consequence of A5: the subject who alienates his AT produces not only the value of his own data, but also an increment to the value of the entire aggregate. This increment is not compensated to the subject and is structurally impossible for him to assess, since the subject does not possess information about the aggregate as a whole nor about his contribution to it.

Limit of A5: the nonlinear growth of predictive value through aggregation is established. However, the description remains static: it does not explain the dynamics by which the system reproduces its own conditions for growth. For this, A6 is required.

**A6. Axiom of Self-Augmentation.** The cycle Attention  $\rightarrow$  Data  $\rightarrow$  Model $\Delta$   $\rightarrow$  Utility  $\rightarrow$   $\uparrow$ Attention is closed, with no internal saturation point or self-regulation.

Justification. A6 describes the dynamic structure of digital capital as a self-augmenting system. The cycle consists of four interrelated elements. Attention: subjects alienate AT in favor of the platform in accordance with A2–A3. Data: AT are crystallized into predictive data in accordance with A4, and aggregated into predictive mass in accordance with A5. Model $\Delta$ : the platform employs the predictive mass to update its algorithmic optimization function, that is, to increase the precision of prediction of subjects' behavioral responses. Utility: the updated model generates content and interface decisions that increase the platform's attractiveness for subjects, which structurally necessitates an increase in the volume of alienated attention at the next step of the cycle ( $\uparrow$ Attention).

The A6 cycle is structurally identical to Marx's formula  $M \rightarrow C \rightarrow M'$  (money — commodity — money with increment), yet the substrate differs. In the Marxian cycle, the substrate is monetary capital augmented through the production of a commodity. In the A6 cycle, the substrate is predictive power, augmented through the production of utility that retains the attention of subjects. The formula of the cycle:  $PM \rightarrow AT \rightarrow PM'$ , where PM is predictive power, AT are attention-tokens, and  $PM'$  is expanded predictive power.

A critically important property: the cycle has no internal saturation point. In the classical production economy, market saturation limits the growth of profit. In the A6 cycle, saturation is structurally impossible for two reasons. First: the predictive power of the model does not reach a limit of precision — any increase in data improves the predictive function, and consequently the efficiency of attention retention. Second: the utility offered to subjects is adapted to each subject individually on the basis of his own AT, which renders the platform's offering structurally more relevant to the given subject than any non-digital alternative.

The absence of self-regulation means: no mechanisms arise within the system that correct the cycle in the direction of reducing the intensity of alienation. The platform's profit is directly proportional to the volume of AT; any reduction in the intensity of alienation reduces profit. Consequently, the algorithmic optimization function is configured for the maximization of AT, and not for any other criterion.

Formal expression. Let  $PM(t)$  denote the predictive power of the platform at moment  $t$ ,  $AT(t)$  the aggregate volume of attention-tokens at moment  $t$ , and  $U(t)$  the utility generated by the platform at moment  $t$ . Then:  $PM(t+1) = PM(t) + \delta(AT(t))$ , where  $\delta > 0$ ;  $U(t) = g(PM(t))$ , where  $g$  is a monotonically increasing function;  $AT(t+1) = h(U(t))$ , where  $h > 0$ . From the three equations it follows:  $AT(t) \rightarrow \infty$  as  $t \rightarrow \infty$ , if no external constraint obtains. No internal constraint is provided for within the system.

### 1.3.1. Live-Service Games as the Pure Form of A6

The live-service gaming model (to which class belong Fortnite, Roblox, Genshin Impact, Honkai: Star Rail, and a significant portion of the mobile market) constitutes the structurally purest realization of the A6 cycle in the sense that all four elements of the cycle are expressed within it without the masking factors characteristic of other segments of the digital economy.

In the traditional gaming model, the product has a terminal state: the subject completes the narrative or mechanical progression, after which the cycle of interaction concludes naturally. The live-service model eliminates this terminal state as an architectural principle. Seasonal events, battle passes, daily rewards, limited-time content — all of these mechanics perform a single functional task: they create artificial cycles of renewed attention alienation that preclude the attainment of a completed state. The subject returns to the platform not because his activity within it is unfinished, but because the platform's algorithmic optimization function produces the state of FOMO (fear of missing out) through temporal restriction of content availability.

The variable ratio reinforcement mechanic, instantiated in loot box systems, is the direct transposition of the laboratory-verified Skinnerian schedule into the digital interface. The variable ratio reinforcement schedule produces the greatest stability of a behavioral pattern with the greatest resistance to extinction, as was established in neuropsychological research long before the emergence of digital platforms. Its application in live-service games is not an incidental borrowing but constitutes a deliberate design decision directed at the maximization of AT in accordance with A6.

Social pressure as an element of the system is instantiated through mechanics of joint participation (guild systems, co-op events, social leaderboards) that functionally produce exit costs: a subject who ceases participation inflicts damage upon social bonds within the platform. This transforms the individual decision to cease attention alienation into a socially costly decision, thereby raising the exit barrier above the individual assessment of utility.

Critically significant is the age profile of the audience. The prefrontal cortex, responsible for the function of impulse control and the long-term weighing of the consequences of decisions, completes its structural formation at an average age of 25. The live-service gaming market is deliberately oriented toward the age range of 6–17, that is, toward subjects in whom the neurophysiological mechanism of impulse control is in the process of formation. From A6 it follows that a system without internal self-regulation, optimized for the maximization of AT, when operating upon a subject without a physiologically formed mechanism of counteraction, produces an intensity of attention alienation structurally exceeding that which it would produce upon an adult subject under otherwise equal conditions. This consequence is not an unforeseen effect: it follows directly from the logic of the A6 cycle as applied to this age segment.

The live-service model is accordingly not a special case of the digital economy but its structurally pure form, in which product design coincides fully with the logic of A6.

## **Chapter Summary**

Chapter 1 established the six axioms of the ontological layer, describing attention as the primary substrate of digital capital. A1 registered the absolute finitude of the individual resource. A2 delineated the non-delegability and alienability of attention, introducing the legally significant distinction between voluntary and structural alienation. A3 established AT as the elementary unit of alienation and its irreversibility. A4 described the crystallization of living attention into dead predictive data as the ontological status of the transition, not as a metaphor. A5 registered the nonlinear increment of predictive value through the aggregation of AT. A6 described the closed cycle of self-augmentation of predictive power without an

internal saturation point or self-regulation. Subsection 1.3.1 verified A6 on a structurally pure empirical case.

Taken together, the six axioms describe digital capital as an objective structure: a system in which the alienation of the finite resource of subjects produces the infinitely self-augmenting predictive power of platforms. The description is complete at the level of objective structure and simultaneously requires logical supplementation: a structure without a subject does not explain how the subject reproduces the conditions of his own alienation. The objective layer has described the movement of capital. The subjective layer must describe the movement of the subject within this capital.

## **Transition to Chapter 2**

The axiomatics of A1–A6 described the cycle from the standpoint of the platform: what happens to AT following their alienation. The question of what happens to the subject in the course of and after alienation remained structurally unresolved. If A6 produces a system without an internal saturation point, and if this system is optimized for the retention of the subject's attention, then the subject must exhibit changes correlating with the intensity of participation in the cycle. The nature and mechanism of these changes constitute the subject matter of Chapter 2.

## **Chapter 2. The Structure of Accumulation**

Chapter 1 established the ontological foundations: attention as the finite, non-delegable, alienable resource of the subject; the attention-token as the elementary unit of alienation; the crystallization of living attention into predictive data; the nonlinear increment of their value through aggregation; the closed cycle of self-augmentation without an internal saturation point. The aggregate of A1–A6 described the structure in its synchronic dimension — its elements, properties, and dynamics in an individual act of interaction.

The present chapter transitions to the diachronic dimension of the same structure: to a description of the manner in which the accumulation of predictive capital over time produces stable structural patterns that follow logically and necessarily from the axioms. These patterns are registered as regularities — not as empirical generalizations verifiable statistically, but as conclusions whose necessity is determined by the axiomatics of A1–A6. The negation of any of the regularities of this chapter would require the negation of the corresponding axiom.

The four regularities introduced in this chapter describe the escalating asymmetry between a platform that has accumulated predictive capital over time and any agent attempting to enter the same structural layer at a later stage. This asymmetry is not an incidental market effect but a structural consequence of the properties of the substrate established in Chapter 1. It escalates monotonically and, under the preservation of the existing architecture, is not eliminable through market mechanisms.

### **2.1. The Temporal Barrier and the Point of No Return**

**Regularity 1** (from A1 + A2). Attention is an irrecoverable resource with zero restoration following alienation: every attention-token is irrecoverable.

Justification. A1 registered the absolute finitude of the subject's aggregate waking time  $T(s)$ . A2 established that the alienation of an attention-token is not a transfer of the capacity for attention: the subject is divested of a unit of  $T(s)$  without acquiring its equivalent in any other form. From the conjunction of A1 and A2 it follows that every alienated AT reduces the remainder of  $T(s)$  by the corresponding magnitude without possibility of compensation. Zero restoration means not the absence of the subject's subjective experience of utility from interaction with the platform, but the structural impossibility of recovering the expended  $T(s)$  — irrespective of the quality of the utility provided and irrespective of the subject's subsequent actions.

This property fundamentally distinguishes attention from all other resources that the subject exchanges with economic agents. Money paid for a commodity may be returned. Physical energy expended on labor is restored through rest. Time spent viewing a social media feed is restored by no means whatsoever. Regularity 1 registers that every interaction with a platform constitutes an irreversible transaction in which the subject surrenders a resource that has no analog of replenishment.

The political consequence of Regularity 1 consists in the following: the aggregate volume of AT alienated by the subject over the period of platform use represents an irrecoverably lost share of  $T(s)$  — the sole absolutely limited resource of the subject. This assertion is not normative: it establishes a structural fact necessary for subsequent analysis. The evaluative qualification of this fact is effected within the framework of normative analysis [Reference to normative principle NA0 of Volume III].

Regularity 1 describes the position of the subject in each individual transaction. Understanding what happens to the alienated AT on the platform's side — and why this process is not symmetric — requires Regularity 2.

**Regularity 2** (from A3 + A4). Every attention-token is irreversibly capitalized by the platform: no act of participation may be de-capitalized without institutional intervention.

Justification. A3 established AT as an elementary unit with bilateral irreversibility — informational and temporal. A4 described crystallization: the subject's living attention passes into the platform's predictive data without returning to the subject in the form of attention. From A3 + A4 it follows that each AT, following transfer to the platform, becomes an element of the platform's predictive capital  $K(P)$ . This element is preserved within  $K(P)$  independently of the subject's subsequent actions — including cessation of platform use, a request for data deletion, or the withdrawal of consent to processing.

De-capitalization — the procedure for extracting AT from  $K(P)$  with the elimination of its contribution to the platform's predictive power — is structurally inaccessible to the subject for the following reason. The platform embeds AT not as an isolated record in a database, but as an element of the aggregated predictive mass (A5) upon which the model is trained. Deletion of the subject's record from the database does not eliminate the contribution of his AT to the weights of the trained model. The sole procedure capable of interrupting capitalization is institutional intervention obliging the platform to retrain the model without the

subject's data. Such a procedure presupposes a legal regime absent from the majority of operative jurisdictions as of 2026.

Regularity 2 registers the qualitative asymmetry between the transactional positions of the subject and the platform. The subject loses his resource irreversibly (Regularity 1). The platform acquires a predictive asset that is preserved and functions after the cessation of any relations with the subject. This is not a particular feature of individual platforms — it is a structural consequence of the nature of predictive capital as such.

Connection to Volume III: from Regularity 2 there follows directly the normative requirement of a de-capitalization principle as a constitutionally protected right of the subject within cybernetic constitutionalism. The mechanism of institutional intervention rendering de-capitalization operationally possible constitutes the subject matter of [NA0 and the de-capitalization mechanism of Volume III].

Regularities 1 and 2 together describe the asymmetry of a single act of alienation: the subject loses an irrecoverable resource; the platform acquires an irreversibly capitalized asset. Regularities 3 and 4 describe how this asymmetry escalates over time to the level of a structural barrier not surmountable through market competition.

**Regularity 3** (from A5 + A6). Predictive power accumulates monotonically: it does not diminish with growth in data volume and does not reach saturation organically.

Justification. A5 established the nonlinear increment of predictive value through aggregation: the value of the aggregate array of AT exceeds the sum of the values of the individual AT constituting it. A6 described the cycle without an internal saturation point: each additional AT increases the predictive power of the platform, which in turn increases the efficiency of attracting the next AT. From the conjunction of A5 and A6 it follows that predictive power  $PM(P)$  is a monotonically non-decreasing function of time, provided that the platform continues to receive AT.

Monotonicity means: each additional AT increases  $PM(P)$  or preserves it unchanged, but never diminishes it. This property is not characteristic of classical forms of capital, where market saturation or the physical depreciation of assets create limits on growth. The absence of organic saturation is conditioned by the nature of the predictive function: the behavior of subjects is a non-stationary process continuously generating new AT, by virtue of which the predictive model never attains a state in which additional data would not improve its precision.

In other terms: the predictive power of the platform is directly proportional to the aggregate volume of AT and the time of their accumulation, and this proportionality is nonlinear — each subsequent unit of data yields a greater increment of predictive value than the preceding unit, as a consequence of the deepening of correlational connections within the aggregated model (A5).

Regularity 3 establishes that predictive capital is a self-augmenting asset without an internal limit on growth. Regularity 4 establishes the consequence of this monotonicity for the structure of competition between platforms.

**Regularity 4** (from A6). Temporal barrier: the early history of data is not reproducible by a competitor within the same modal layer, which constitutes a structural, rather than temporary, advantage.

Justification. From A6 it follows that predictive power  $PM(P)$  is a function not only of the volume of AT, but also of the time of accumulation — through the parameter  $Model\Delta$ , denoting the depth of the trained model. A model trained on data spanning an extended period contains temporal patterns — sequences of behavioral changes in subjects over time — that cannot be reproduced by a model trained on data spanning a shorter period, even if the volume of data in both cases is identical. Temporal depth is an autonomous and irreplaceable component of predictive power.

This means the following. A competitor entering the market at moment  $t_1$  cannot reproduce the predictive power of a platform functioning since moment  $t_0 < t_1$ , even having collected an identical volume of AT over a period equal to  $t_1 - t_0$ , because the behavioral patterns of subjects in the later period are structurally distinct from the patterns of the earlier period — including as a consequence of the platform's own influence upon the behavior of subjects over the entire period of its functioning. The platform is thus itself the cause of the alteration in the behavioral data constituting its competitive advantage: the longer it functions, the deeper its influence upon the patterns it simultaneously registers.

The distinction between a temporary and a structural advantage is fundamental. A temporary advantage is eliminable through sufficient investment and sufficient time. A structural advantage is eliminable only through the alteration of the structure itself — that is, through transition to a different modal layer or through institutional intervention constraining the use of accumulated historical data. Under preservation of the existing architecture, the temporal barrier escalates monotonically in accordance with Regularity 3: the advantage of the early platform increases with each additional period of accumulation.

The assumption that market competition corrects the temporal barrier without institutional intervention is refuted as follows. A competitor capable of collecting the same volume of AT more rapidly through higher investment will find that accelerated collection of AT does not reproduce temporal depth: the patterns of subjects recorded over a month of intensive data collection are structurally impoverished relative to patterns accumulated over seven years of gradual interaction — by virtue of the absence of long-term behavioral sequences and their changes. Predictive power proportional to temporal depth (A6) cannot be obtained through accelerated means.

### **2.1.1. The Temporal Barrier in the Gaming Industry**

Regularity 4 is verified on the structural case of the live-service sector of the gaming industry described in subsection 1.3.1. Fortnite (Epic Games) has been functioning since 2017. By 2025, the platform had accumulated more than seven years of behavioral data: the patterns of audience response to limited-time events, the price elasticity of cosmetic items across various age segments, temporal engagement profiles by demographic groups, the comparative effectiveness of social pressure mechanics. This data is not reproducible by a new competitor within the same modal layer — in precise conformity with Regularity 4.

The counterexample most frequently adduced as a refutation of Regularity 4 — the entry of TikTok into the market without Facebook's data history — confirms, rather than refutes, the regularity upon structural analysis. TikTok did not enter the modal layer of Facebook: it created a new modal layer — short vertical video with algorithmic recommendation without a social graph — contrasting it with the feed of text and photographs ordered by social connections. This means that TikTok initiated a new A6 cycle in a new modal layer, rather than surmounting the temporal barrier within an existing one. Once TikTok's new modal layer was established and had accumulated its own data history, Regularity 4 reproduced itself within it without exception: entry into the short-form video layer in 2025 is structurally as difficult as entry into the social graph layer was in 2010.

The consequence of this analysis for the competitive structure of the market: the sole mechanism for overcoming the temporal barrier within the existing architecture is the creation of a new modal layer. However, the number of structurally unoccupied modal layers is finite and diminishes with each new dominant live-service product. As a result, the aggregate temporal barrier at the level of the entire market escalates not only within each individual layer, but also through the contraction of the number of unoccupied layers available for new entry. This property of Regularity 4 is of direct significance for normative analysis: it establishes that under preservation of the existing architecture, the concentration of predictive capital in the hands of early dominant platforms is structurally escalating, not temporary.

## **2.2. Mathematical Model: $C = f(\Sigma \text{ Attention}_i, t, M)$**

Regularities 1–4 admit of formalization within a unified model of predictive capital, integrating three structural variables: the volume of alienated attention, the time of accumulation, and the depth of the model.

Structure of the model. The predictive capital of a platform  $C$  is defined as a function of three parameters: the aggregate volume of attention-tokens  $\Sigma \text{ Attention}_i$  alienated by all subjects over the observation period (realizing A3 + A5); the time of accumulation  $t$ , measured from the moment of the platform's commencement of functioning to the moment of assessment (realizing Regularity 4 and the temporal barrier); and the depth of the model  $M$  (model-depth) — a measure of the complexity of the predictive function trained on  $\Sigma \text{ Attention}_i$  over period  $t$  (realizing Model $\Delta$  from A6).

Properties of the function. Predictive capital  $C$  is directly proportional to each of the three parameters and increases with the growth of any one of them given the other two fixed. This is not linear proportionality: the increment of  $C$  from each additional unit of parameter  $t$  exceeds the increment from the preceding unit, all other conditions being equal. In other terms, predictive capital increases with acceleration as the time of accumulation increases — which formalizes the self-augmentation established by A6.

This property has a concrete socioeconomic significance. A platform functioning for ten years possesses a predictive capital that nonlinearly exceeds the predictive capital of a platform functioning for five years with an identical volume of users and identical infrastructure investment. The gap escalates with acceleration: each subsequent year adds more to the advantage of the early platform than the preceding year. It follows from this that

a competitor who has entered the market one year late will find, after five years of functioning, that the gap has increased rather than narrowed.

Interdependence of parameters. Model depth  $M$  is not an independent variable:  $M$  increases as a function of  $\Sigma \text{Attention}_i$  and  $t$  jointly. This means that the three parameters of the model are connected through an internal dependency: the early platform accumulates not only a greater volume of AT, but also a deeper model, which in turn increases its capacity to attract AT in the next cycle (A6). The model describes not three independent advantages of the early platform, but one self-reinforcing advantage with three measurable components.

Limit of the model.  $C = f(\Sigma \text{Attention}_i, t, M)$  describes predictive capital as a principally measurable, though structurally opaque, variable: it is inaccessible to direct observation by subjects, competitors, or regulators in the absence of mandatory disclosure. The model does not describe the mechanism by which  $C$  is converted into social and political influence beyond market relations — that is, it does not describe the external effects of predictive capital accumulation for institutions not participating in transactions between the subject and the platform. That question constitutes the subject matter of the subsequent chapters.

## Chapter Summary

Chapter 2 derived four regularities that follow logically and necessarily from axioms A1–A6. Regularity 1 registered zero restoration of the subject's resource following each AT, delineating the fact of resource loss from the subject's subjective assessment of the utility of interaction. Regularity 2 established the irreversible capitalization of AT by the platform and the structural inaccessibility of de-capitalization without institutional intervention, identifying the legal lacuna as a systemic rather than technical defect. Regularity 3 described the monotonic escalation of predictive power without organic saturation, establishing the fundamental distinction between predictive capital and classical forms of capital. Regularity 4 — the temporal barrier — demonstrated that the advantage of a platform with an early data history is structural and escalating, not temporary and surmountable through market mechanisms. The model  $C = f(\Sigma \text{Attention}_i, t, M)$  formalized the accelerated escalation of this barrier as a structural property of predictive capital. Subsection 2.1.1 verified Regularity 4 on the case of the live-service gaming industry, delineating the shift of modal layer as the sole possible market response to the temporal barrier, along with its fundamental limits.

## Transition to Chapter 3

Regularities 1–4 and the model  $C$  describe the structure of predictive capital accumulation as an objective diachronic process. This description remains incomplete in the following respect: it establishes the escalation of asymmetry between the platform and its competitors, but does not describe the mechanism by which this asymmetry is translated into structural consequences for subjects and for public institutions situated beyond the market relations between the platform and the user. Predictive capital is not a self-sufficient value — its worth is realized through conversion into other forms of influence: behavioral, economic, political. The description of the mechanisms of this conversion and their systemic consequences constitutes the subject matter of Chapter 3.

## Chapter 3. The Limits of Objective Description

Chapters 1 and 2 constructed a complete description of digital capital as an objective structure: six axioms established the ontological properties of the substrate; four regularities described the stable patterns of its accumulation over time. This description is internally coherent and formally complete at its own level. At the same time, it contains an implicit assumption that has until this moment remained unexplicated: every assertion about asymmetry — every conclusion to the effect that the resource is not restored, that capital is accumulated irreversibly, that competition is structurally impossible — presupposes an agent for whom this asymmetry constitutes a real state of affairs.

The present chapter formalizes this assumption in three theorems of the ontological layer. The theorems introduce no new elements — they derive from the aggregate of axioms A1–A6 and Regularities 1–4 the terminal assertions that exhaust what can be said about the objective structure of digital capital without recourse to the subject. Thereby they render formally visible the moment at which the objective description itself logically requires transition to the anthropological layer of analysis.

### 3.1. The Theorem of Surplus Attention (T1)

**T1** (from Regularities 1 + 2 + 3). Theorem of Surplus Attention: each cycle of the subject's interaction with the platform generates predictive value that systematically exceeds the zero compensation provided to the subject; the asymmetry is formally measurable through profile-index at zero monetary remuneration.

Justification. Regularity 1 established that each AT constitutes an irreversible loss of the subject's resource. Regularity 2 established that each AT is irreversibly capitalized by the platform. Regularity 3 established that the predictive power into which AT are converted escalates monotonically. From the conjunction of the three regularities it follows that each interaction cycle produces an increment of the platform's predictive capital at zero monetary compensation to the subject — that is, the subject systematically produces value that the platform appropriates without a corresponding monetary payment.

This assertion requires precise delineation. The platform provides the subject with utility — content, service, communication — and this provision is real. The question is not whether the subject receives something in return, but whether the utility received constitutes an equivalent of the predictive value produced by the subject. Since monetary compensation is zero and the predictive value of AT nonlinearly increases upon aggregation (A5), the systematic discrepancy between the value produced and the utility provided is a structural property of the relation, not a deviation from it.

Profile-index (hereinafter: PI) is an indicator measuring the accumulated predictive value of an individual subject's profile relative to the value of advertising placement directed at that profile. PI is defined as the ratio of the market value of targeted access to the subject to the aggregate monetary compensation received by the subject over the entire period of AT generation. Since the monetary compensation of the subject in the majority of platform transactional models equals zero, PI is not defined as a finite number — it tends to infinity. This formalizes the asymmetry: the subject produces a finite measurable value at zero remuneration.

Derivation of T1. The structure of digital capital produces surplus predictive value in each cycle of the subject's interaction with the platform. This value is not compensated to the subject in monetary form and cannot be compensated through utility, since the utility provided is a product of the use of predictive capital, not its equivalent. The asymmetry is not incidental — it is structurally reproduced in each cycle in accordance with Regularities 1–3.

Connection to normative analysis. T1 establishes a structural fact without evaluative qualification. The question of whether this structure is legitimate, permissible, or subject to correction belongs to the normative layer [see normative principle NA0 of Volume III]. T1 furnishes this layer with a formal description of precisely what is subject to evaluation.

Limit of T1: the theorem describes the asymmetry from the standpoint of structure, that is, from the standpoint of the impersonal relation between the subject as source of AT and the platform as their recipient. It does not describe how this structure retains the subject in the position of AT source — that is, it does not contain a description of the mechanism of reproduction of the relation. This question concerns the temporal dynamics of the competitive structure and constitutes the subject matter of T2.

### **3.2. The Theorem of the Temporal Barrier (T2)**

**T2** (from Regularity 4). Theorem of the Temporal Barrier: following the attainment of the point of no return, competition within a homogeneous modal layer is structurally impossible without external intervention; a shift of modal layer constitutes the creation of a new contradiction  $\Omega_0$ , not a refutation of the theorem.

Justification. Regularity 4 established that the early history of data is not reproducible by a competitor within the same modal layer, and that the temporal barrier is structural rather than temporary. T2 derives from this a terminal assertion about competition: there exists a moment  $t^*$  in the development of a platform — the point of no return (hereinafter: PNR) — after which no competitor entering the same modal layer with zero data history is capable of surmounting the accumulated temporal barrier through market means.

The point of no return PNR is defined not by the absolute magnitude of the platform's predictive capital, but by its ratio to the maximum achievable predictive capital of a competitor within the given modal layer. A competitor attains the maximum achievable level of predictive capital at the moment when its growth levels off — that is, when the increment of the competitor's PM per unit of time becomes equal to the increment of the dominant platform's PM over the same unit of time. If at this moment the gap  $B(t) = PM(P) - PM(P_{\text{new}})$  remains positive, the competitor cannot surmount it under preservation of the existing architecture. From Regularity 4 it follows that, at a sufficiently large  $t_0$  of the dominant platform, this scenario is structurally inevitable.

Formal expression. The temporal barrier  $B$  is directly proportional to the accumulation time of the dominant platform and nonlinearly increases with the growth of this time (a consequence of the model C and Regularity 4). Upon exceeding the PNR:  $B(t) > B_{\text{max}}(P_{\text{new}})$  for any competitor  $P_{\text{new}}$ , where  $B_{\text{max}}(P_{\text{new}})$  is the maximum achievable reduction of the barrier through market investment. Competition within the modal layer becomes structurally impossible.

Delineation of the modal layer shift. The objection is the following: historically, new platforms have entered the market and occupied dominant positions despite the existence of established incumbents. T2 responds to this objection through the concept of the modal layer. The successful entry of a new market participant is, as a rule, effected not through the surmounting of the temporal barrier within an existing layer, but through the creation of a new modal layer — a new format of subjects' interaction with the digital environment for which none of the existing participants has accumulated a data history. This is not a refutation of T2: it is a confirmation that the temporal barrier is indeed insurmountable within the existing layer, and that the sole market response is the creation of a new layer.

The creation of a new modal layer is described as the emergence of a new contradiction  $\Omega_0$  — a new primary cycle of attention alienation in a previously non-existent format. At the moment of its emergence, all participants stand in equal position relative to the temporal barrier. However, once the new layer stabilizes and the first dominant participant begins the accumulation of a data history, T2 reproduces itself within the new layer on the same foundations.

The consequence of T2 is that the number of structurally available modal layers for the creation of a new  $\Omega_0$  is finite and diminishes with each stabilized platform. From Regularities 3 and 4 jointly it follows that the aggregate concentration of predictive capital in the system escalates monotonically: each new layer in time reproduces the same structure of concentration as the preceding one, while the former layers continue their accumulation. The market mechanism does not correct this concentration — it reproduces it in new forms.

Derivation of T2. Following the attainment of the point of no return, the dominant platform occupies a structurally protected position within its modal layer. Market competition within the layer is impossible without external intervention that alters the accessibility or the mode of use of the accumulated data history. A shift of modal layer is the sole market response; however, it does not correct concentration in existing layers and reproduces an analogous concentration in the new layer.

Limit of T2: the theorem describes the structure of competition — that is, the position of agents relative to one another. It does not describe the position of subjects relative to the system as a whole: what happens to the subject under conditions in which competition is structurally absent and, consequently, alternatives to interaction with the dominant platform are either inaccessible or entail asymmetric costs. That question constitutes the subject matter of T3.

### **3.3. The Theorem of the Structural Absence of Correction (T3)**

**T3** (from Regularity 3 + T1). Theorem of the Structural Absence of Correction: the system contains no internal mechanism for the correction of the concentration of predictive capital; the completeness of this derivation logically requires the identification of an agent of consequences, which obliges transition to the anthropological layer of analysis.

Justification. Regularity 3 established that predictive power accumulates monotonically without organic saturation. T1 established that each interaction cycle produces surplus predictive value at zero compensation to the subject. From the conjunction of Regularity 3 and T1 it follows that the system in each cycle increases the platform's predictive capital

without thereby producing a mechanism that would diminish this concentration or constrain its growth.

An internal corrective mechanism in this context means: a property of the system whereby the attainment of a certain level of predictive capital concentration automatically produces forces diminishing that concentration — without external intervention. In classical market theory, such a mechanism is competition: high concentration attracts competitors who erode the monopoly position. T2 demonstrated that following the attainment of the point of no return this mechanism does not function under conditions of predictive capital. Neither market competition, nor the behavior of subjects, nor technological development produces internal pressure toward a reduction of PM concentration under preservation of the existing architecture.

This means that the system is structurally escalating in a single direction: the concentration of predictive capital increases in each cycle, encountering no internal constraint. The sole mechanisms capable of interrupting or reversing this dynamic are external to the system — that is, institutional: regulatory, legal, or constitutional.

Derivation regarding the agent of consequences. The assertion of T3 is complete at the structural level. However, it contains an implicit element that has until this moment remained unexplicit throughout the entire ontological description. When T1 asserts that the subject receives no compensation — this assertion presupposes a subject who receives no compensation. When T2 asserts that the competitor cannot enter the market — this assertion presupposes a competitor to whom entry is closed. When T3 asserts that a corrective mechanism is absent — this assertion presupposes an agent who bears the consequences of the absence of correction.

Without the identification of this agent, the ontological description is formally incomplete: it describes a structure producing consequences, but does not describe the one for whom these consequences constitute real experience. This is not a rhetorical move and not an appeal to empathy — it is a formal requirement of logical completeness. T3 registers the limit beyond which the structural description cannot advance without the introduction of the anthropological subject.

Verification example. The European Commission report on digital markets (Digital Markets Act enforcement, 2024) recorded that not one of the six platforms that received gatekeeper status under the regulation had reduced its share of predictive capital in the corresponding modal layer over the period from the introduction of the regulation through December 2025. This confirms T3 empirically: external regulatory intervention in its current form is insufficient for the correction of concentration; however, the very fact of the necessity of regulatory intervention confirms the conclusion of T3 regarding the absence of an internal corrective mechanism.

Formal expression. Let  $\text{Corr}(\text{PM}, t)$  denote the internal corrective force of the system at moment  $t$ . From Regularity 3 and T1 it follows that  $\text{Corr}(\text{PM}, t) = 0$  for all  $t$  under preservation of the existing architecture. Consequently,  $\text{PM}(\text{P}, t)$  increases monotonically without limit. The sole operation nullifying  $\text{Corr} = 0$  is external institutional intervention  $I$ :  $\text{PM}(\text{P}, t) \rightarrow \text{PM}(\text{P}, t) - I(t)$ , where  $I(t) > 0$  only in the presence of an operative normative mechanism.

---

## **Δ1 — CRISIS: THE LIMIT OF OBJECTIVITY**

Theorems T1, T2, and T3 exhaust what can be said about digital capital as an objective structure. T1 formalized the surplus nature of predictive value. T2 established the structural impossibility of competition following the point of no return. T3 registered the absence of an internal mechanism for the correction of concentration and thereby rendered visible the implicit agent of consequences that had been present throughout the entire ontological description.

This agent — the subject as bearer of the consequences of the structure — cannot be described through the means of the objective layer: the objective layer describes the movement of capital, but not the movement of the subject within this capital. For a description of how the subject reproduces the conditions of his own alienation — how the system produces the behavioral and cognitive states that retain the subject in the position of AT source, and how the consequences of T1–T3 are realized in the concrete experience of the subject — transition to the anthropological layer of analysis is necessary. This transition is not an optional extension of the theory, but a formal requirement of its completeness: ontology without a subject describes the structure, but does not explain its reproduction.

---

### **Chapter Summary**

Chapter 3 derived three theorems exhausting the content of the objective ontological layer. T1 formalized the surplus nature of predictive value through the concept of profile-index and established the structural character of the asymmetry between the value produced by the subject and the utility received by him. T2 established that following the attainment of the point of no return, competition within a modal layer is structurally impossible without external intervention, and that a shift of modal layer reproduces the same structure in a new format. T3 registered the complete absence of an internal mechanism for the correction of the concentration of predictive capital and thereby designated external institutional intervention as the sole possible instrument of systemic correction. All three theorems together established the formal limit of objective description: it logically requires the identification of an agent of consequences — the subject as bearer of the experience of the structure.

### **Transition to Chapter 4**

The objective layer described the system from the standpoint of structure. T3 rendered visible that this structure presupposes a subject bearing its consequences. For an understanding of why the subject reproduces the conditions of his own alienation — why the A6 cycle is not interrupted by the subject himself, despite the consequences of T1 that it produces — it is necessary to describe the mechanism by which the system produces the behavioral states of the subject that retain him in the position of AT source. This mechanism constitutes the subject matter of Part II, the anthropological layer.

## PART II. ANTHROPOLOGY OF PLATFORMS

*The Subjective Layer: Who Lives Within This Structure and How the Structure Destroys Him*  
*Objective: To identify the subjects who bear the consequences of the structure established in Part I. To demonstrate that these consequences obtain not only economically and politically, but physiologically — through measurable cognitive atrophy and psychological manipulation. The completeness of this description will necessitate a transition to intersubjective form.*

### Chapter 4. Class Morphology of the Digital World

Part I established the ontological structure of digital capital: axioms A1–A6, regularities 1–4, and theorems T1–T3 jointly demonstrated that the platform economy self-valorizes without internal limit, systematically converting the subject's attention into predictive value. In that description, however, the subject was present only as an abstract resource — a faceless, historyless, physiologyless supplier of attention. Regularities 1–4 registered the movement of capital, but not the movement of destruction.  $\Delta 1$  — the crisis of the limit of objectivity — formally mandates the next step: it is necessary to identify who exactly bears the consequences of the structure described in Part I, by what means these consequences obtain, and through what mechanisms the structure reproduces its own stability at the subject's expense.

The anthropological layer is not a supplement to the ontological — it is its logical entailment. If the structure self-valorizes (A6), if the temporal barrier renders competition structurally impossible (T2), if systemic crises do not extirpate predictive power (T3), then behind each of these propositions stands a subject who bears the economic, political, and physiological consequences of these theorems. Chapter 4 introduces the axiomatic foundations of the anthropological layer and describes the class morphology — five structural positions that the subject occupies in relation to predictive infrastructure.

#### Axioms of Anthropology: Six Foundations

**A7. Axiom of the Subject.** The subject exists as an immediate given, not reducible to his profile on the platform. This given is the initial condition of analysis; the subsequent axioms describe how the structure of digital capital systematically undermines and destroys precisely this subjecthood.

**A8. Axiom of Limitation and Psychotechnological Capture.** The individual is limited informationally, cognitively, and temporally — and this limitation is the condition of his exploitation, not an incidental defect. Beyond cognitive limitation, the individual is subjected to psychotechnological exploitation: platform architecture systematically employs dark patterns, emotional nudges, and persuasive loops to exploit known vulnerabilities (variable reward, loss aversion, social comparison, FOMO). This produces measurable cognitive atrophy and psychological harm (anxiety, depression, diminished impulse control). This exploitation is not a side effect — it is an architectural decision optimizing engagement through exploitation of the dopaminergic system and executive function.

**A9. Axiom of Intentionality.** The individual's attention is not random but motivated — which renders it predictable and, consequently, capitalizable.

**A10. Axiom of Freedom.** Freedom of choice is given as the condition of political subjecthood. Consent given in the structural absence of real alternatives is not free consent (Rawls: fair equality of opportunity).

**A11. Axiom of Corporeality.** The subject is embodied — his digital actions have physical and social consequences that he bears.

**A12. Axiom of Reflexivity.** The subject is capable of reflecting on the conditions of his own existence — which renders him potentially capable of resistance and of constituting a new order.

#### 4.1. Subjecthood as Axiom (A7)

A7 introduces the subject into the analytical apparatus not as a derivation from structure, but as an axiom — an initial condition requiring no prior demonstration. This is methodologically fundamental. The ontological layer (A1–A6) described digital capital as an objective structure: data flows, predictive value, temporal barrier. In that description the subject was present as an abstract source of behavioral traces — an impersonal supplier of attention. A7 ruptures this reduction: the subject exists as an immediate given that precedes his profile on the platform.

The distinction between the subject and his profile is not metaphysical but politically and juridically significant. The profile is an aggregated history of interactions belonging to the platform by virtue of the architectural arrangement of the system (A4): data are stored on the platform's servers, the interpretive algorithm belongs to the platform, and the predictive model extracted from that data is the platform's property. The subject has no access to his own profile in its completeness — he sees only the interface, not the model. It therefore follows that the profile is the subject's shadow, cast into the space of predictive infrastructure and appropriated without his knowledge. The subject is he who casts this shadow; the platform is that which stores it.

A possible objection: the subject is an epiphenomenon of neural processes, reducible to computational patterns no less than to an algorithmic profile. A7 neither accepts nor rejects this argument at the metaphysical level. What is asserted is otherwise: regardless of the ontological status of subjecthood, the distinction between the subject and his profile possesses normative and legal significance. Legal systems (GDPR, 2018; EU AI Act, 2024) recognize this distinction in precisely the sense that A7 introduces: the profile is data belonging to the subject, but seized from him. The subject as the source of this data possesses rights that his profile does not secure — and this distinction is the foundation of the normative portion of the trilogy (NA0, Part IV).

Antithesis: if the subject is fully predictable — as systems of predictive behavior assert (A3) — then the distinction between subject and profile is mere illusion. The profile predicts behavior with measurable accuracy; consequently, it is the operative subject within this system. The answer is contained in A9 and A12: the intentionality of attention and the reflexive capacity of the subject are not extirpated by predictability, but constitute its precondition. The algorithmic function predicts precisely because the subject is intentional and limited (A8) — but neither intentionality nor limitation is identical to the profile. The

subject reflects, errs, changes. The profile lags, freezes, averages. This distinction is structural and ineliminable.

A7 thus performs a dual function. First, it introduces the normative minimum: the subject as a given whose destruction requires justification and, under certain conditions, constitutional protection (NA0). Second, it establishes the direction of the subsequent axioms: A8–A12 describe not an abstract subject, but the mechanisms through which the structure systematically undermines the very subjecthood that A7 postulates as a given.

Transition to 4.2: if the subject exists as a given not reducible to the profile, it is necessary to explain why this subjecthood proves systematically vulnerable. The answer is contained in A8: not an incidental weakness, but the architectural exploitation of structural limitation.

#### **4.2. Axiom of Limitation and Psychotechnological Capture (A8)**

A8 occupies the central position in the anthropological layer, in that it describes the unified mechanism linking the structural limitation of the subject with the architectural exploitation of that limitation. The consolidation of both dimensions into a single axiom reflects their causal unity, not an incidental coincidence: limitation generates vulnerability; psychotechnological capture systematically exploits that vulnerability as a design decision, not as a side effect.

Three dimensions of limitation constitute the structural precondition of exploitation. Informational limitation: the subject cannot process the aggregate volume of information available in the digital space; consequently, any filtering mechanism (recommendation algorithm, news feed, search) structurally determines what the subject sees. He who controls the filtering controls the subject's informational field — without the subject's awareness and without his consent. Cognitive limitation: the subject's working memory is bounded (on average  $7 \pm 2$  units of information), sustained attention is depleted across a session, and emotional state directly affects the quality of reasoning. Platform architecture optimizes precisely these variables: short-form content, infinite scroll, notifications timed to moments of cognitive decline. Temporal limitation: the subject commands finite time, and that time is irreplaceable. The platform competes for this time not as a service provider, but as infrastructure that retains the subject through architectural mechanisms designed to maximize session length and daily active users.

Psychotechnological capture is effectuated through six mechanisms, each of which constructively exploits a corresponding cognitive vulnerability. Variable reward — a principle drawn from behavioral psychology (Skinner, intermittent reinforcement): the unpredictable alternation of rewards (likes, mentions, new content) generates a compulsive behavioral pattern that is neurochemically analogous to patterns of dependency. Loss aversion: the architecture of notifications and social comparison exploits the asymmetry of perceived losses and gains (Kahneman & Tversky: losses are psychologically approximately twice as significant as gains of equal magnitude). Social comparison: continuous access to the curated self-presentations of other users activates mechanisms of comparative self-assessment, systematically diminishing it — this is a documented factor in the rise of anxiety and depression, particularly in the 13–24 age group (American Psychological Association, 2023–2025). FOMO (Fear of Missing Out): the interface is constructed so that every pause in engagement is perceived as the forfeiture of a significant event, thereby compelling continuous feed refresh and further depleting sustained attention. Dark patterns:

interface decisions that deliberately impede actions contrary to the platform's interests (unsubscribing, account deletion, notification disablement) and that facilitate actions consonant with those interests (subscription renewal, data sharing, content reaction). Emotional nudges: the algorithmic function optimizes emotional arousal (excitement, anger, elation), since high arousal correlates with increased engagement — consequently, content that elicits a strong emotional response is ranked higher regardless of its informational veracity.

Measurable consequences have been established in the neurocognitive literature. According to Nature Neuroscience (a series of studies, 2023–2025), chronic use of short-form video platforms (TikTok, Instagram Reels, YouTube Shorts) at three or more hours per day is associated with an 18–20% reduction in sustained attention when controlling for age and baseline level. This is not a correlation with content quality — it is the consequence of interrupt architecture: the platform systematically redirects attention before the onset of cognitive saturation, with the result that the capacity for sustained attention atrophies through disuse. The FTC (2024–2025) established, in the context of settlements with Epic Games (Fortnite) and Roblox Corporation, the use of gambling-like mechanics (loot boxes, random reward drops) in conjunction with FOMO architecture directed at younger users — qualifying this as unfair and deceptive trade practices precisely by virtue of the targeted exploitation of cognitive developmental incompleteness in that age group.

The cardinal significance of A8 is as follows: psychotechnological exploitation is not a deviation from normal platform operation — it is its design parameter. The metrics that the platform optimizes (DAU, session length, engagement rate) are directly incentivized by the mechanisms described above. It therefore follows that the exploitation of the subject's cognitive vulnerabilities is a structural function of the system, not an incidental abuse by individual engineers. This destroys the standard defense argument: "the platform did not intend to cause harm." The harm is not the intentional consequence of malicious design, but the predictable result of optimizing specified metrics.

Antithesis: if the subject is reflexive (A12), he is capable of identifying the mechanisms of capture and limiting their effect. Answer: the cognitive atrophy produced by chronic exposure to the mechanisms of A8 destroys, first and foremost, precisely those resources required for reflexivity — sustained attention, working memory, and executive function. This is a closed loop: the system consumes the resources necessary for resistance to the system. This proposition will be formalized as theorem T6 (cognitive disarmament) in Chapter 6.

Transition to 4.3: A8 demonstrates that exploitation is possible precisely because the subject is limited. But a critic may respond: the subject voluntarily consents to the use of the platform. This objection requires a separate axiomatic analysis under A10.

### **4.3. Axiom of Freedom and the Problem of Consent (A10)**

A10 formalizes the response to the most prevalent argument in defense of platform architecture: the individual voluntarily accepts the Terms of Service; consequently, his engagement with the platform is the result of free choice, and the value extracted therefrom constitutes a fair exchange.

A10 introduces a fundamental distinction between formal and free consent. Consent is free only when three conditions — drawn from Rawls's theory of justice (fair equality of opportunity) — are satisfied: the subject possesses real alternatives (not merely formal ones), the subject commands sufficient information to evaluate the exchange, and the subject is in a cognitive state that permits rational decision-making.

All three conditions are systematically violated within the structure of digital capital. The structural monopolism demonstrated by T2 (temporal barrier) extirpates real alternatives. When a dominant platform possesses an early data history and network effects, switching to an alternative entails the forfeiture of access to the social graph, accumulated content, and professional connections. This is not the rejection of one product in favor of another — it is structural coercion through infrastructural dependency. Informational asymmetry (A3, A4): Terms of Service are written in legal language averaging 12,000–14,000 words in length (Princeton Web Transparency Project, 2024), contain broad references to policies modifiable unilaterally, and systematically omit disclosure of precisely what predictive models are constructed from the subject's behavioral data. The subject cannot give informed consent if the information concerning the object of consent is structurally concealed from him. Cognitive atrophy (A8): at the moment of platform registration, the subject is in a mode of rapid decision-making stimulated by the interface; the onboarding architecture minimizes the interval between intent and registration precisely in order to preclude a pause for reflection. Consent to the ToS by clicking the "Agree" button is a ritual act devoid of cognitive content.

It therefore follows that the subject's consent on the platform is formally legitimate and substantively vacuous. This is not a claim about subjective deception — it is a claim about the structural arrangement of the consent mechanism. Precisely this distinction constitutes the foundation for normative principles N1 (right to unpredictability) and N4 (procedural audit of predictive models), introduced in Part IV: if consent structurally cannot be free under conditions of monopolized infrastructure, then the legal regime cannot be confined to consent as a sufficient basis for the legitimacy of the exchange.

Antithesis: A10 presupposes a high threshold for "real" consent that is in practice unmet by any complex contract. Insurance policies, mortgage agreements, corporate employment contracts — all contain informational asymmetry and limited reality of alternatives. If the standard of A10 is applied consistently, the entire contractual order is placed in question. Answer: A10 does not assert a universal standard of free consent for all contractual relations. It identifies a specific property of platform infrastructure: within it, informational asymmetry is architecturally designed (dark patterns, opaque ToS), monopolism is the structural consequence of network-effect dynamics (T2), and cognitive atrophy is an optimized design parameter (A8). The combination of these three conditions within a single structure constitutes a qualitatively distinct situation from a complex insurance policy, in which at least one of the three conditions is absent.

Synthesis: A10 does not nullify the concept of consent but calibrates its applicability. Consent as a legal mechanism retains significance where the structural conditions of its possibility are satisfied. Where those conditions are systematically violated — as in the platform economy — consent is an insufficient basis for legitimacy, and a different legal regime is required. This is the formal justification for the constitutional approach of Volume III (principle P2, code supremacy with normative axiom): rules embedded in constitutional code

do not require the subject's consent at each point of interaction precisely because they protect the conditions of possibility for free consent.

Transition to 4.4: A7–A12 describe an abstract subject — his givenness, limitation, intentionality, freedom, corporeality, and reflexivity. But the platform economy does not encounter an abstract subject. It encounters concrete agents occupying distinct structural positions in relation to predictive infrastructure. These positions determine not only what the subject loses, but what he gains, what he sees, and what remains beyond the horizon of his visibility.

#### 4.4. Class Positions

The structure of digital capital, as described in Part I, constitutes five structural positions — not in the sociological sense of class membership by origin, but in the analytical sense: a position is defined by the agent's relation to predictive infrastructure and to the distribution of the consequences of its operation. Each position is characterized by a specific combination of power, vulnerability, and awareness.

**Attention Suppliers (users).** This is the most numerous and structurally least protected position. The individual in this position is the primary source of behavioral data (A3, A4) from which predictive value is extracted (T1). He bears the cognitive and physiological consequences of psychotechnological capture (A8), but his awareness of the mechanics of that capture is structurally suppressed: the interface architecture conceals the predictive model (A4), and cognitive atrophy diminishes the resources available for reflexivity (A8 → T6). The subject's consent in this position formally legitimizes the exchange but is structurally vacuous (A10). It is precisely this agent who is the primary object of normative principles N1–N7.

**Developing Subjects (children and adolescents aged 6–17).** This position is the most vulnerable, in that it combines the incompleteness of a still-forming subjecthood with the full operational functionality of the capture mechanisms. The prefrontal cortex — the anatomical substrate of impulse control, long-term planning, and resistance to social pressure — completes its formation at approximately age 25. It therefore follows that an individual between the ages of 6 and 17 is structurally deprived of the neurobiological preconditions necessary to resist the mechanisms described in A8: variable reward exploits an immature reward system, FOMO exploits the hypersensitivity to social comparison characteristic of adolescence, and loot boxes and daily reward mechanics exploit the absence of developed long-term planning.

The juridical-normative situation here is structurally unresolved within the framework of existing legal regimes. Parents formally provide consent at the point of registration; however, their consent does not cover the actual alienation of the child's attention, which is effectuated through thousands of micro-interactions — each click, each view, each reaction. The parent consents to "use of the platform"; the child clicks on a loot box, responds to a FOMO notification, spends three hours in a session instead of the intended twenty minutes. Not one of these specific interactions is covered by the parent's consent — and not one of them is an isolated event; in the aggregate they constitute the systematic alienation of the child's cognitive resource, the consequence of which is a measurable decline in sustained attention and a rise in anxiety.

The FTC (2024–2026) qualified this practice — in settlements with Epic Games, Roblox Corporation, and separate proceedings concerning Genshin Impact — as unfair and deceptive precisely because it deliberately exploits cognitive developmental incompleteness in this age group. The EU AI Act (Article 5) prohibits subliminal techniques directed at vulnerable groups. Nevertheless, legal regulation in this domain remains reactive and jurisdictionally fragmented — which is itself a consequence of regularity 12 (the state as a buyer of predictions is structurally not a neutral regulator).

**Architects of Prediction (platform owners).** This position is characterized by complete control over predictive infrastructure: ownership of the data history, of the models, and of the interaction architecture. Of all positions, this is the sole structural beneficiary of the mechanisms described in A6 (self-valorization) and T2 (temporal barrier). The agent in this position bears legal liability for architectural decisions (as the FTC cases demonstrate); however, this liability is effectuated with substantial delay relative to the accumulated harm.

**Algorithmic Intermediaries.** Technical operators — engineers, data scientists, product managers — who implement architectural decisions without formal political rights over the results of their work. This position structurally produces consequences (predictive models, capture mechanisms), bears professional responsibility for technical execution, but bears no political responsibility for systemic consequences. This anticipates the problem of algorithmic intermediaries, which Volume III (Formal Theory of the Digital Republic) resolves through Rockefeller Mode (principle P16): a mechanism of constitutional accountability for infrastructure operators functioning without political mandate.

**Buyers of the Predicted Future.** States, advertisers, political campaigns — agents who acquire the predictive output of platforms to pursue their own objectives. This position is central to the understanding of regularity 12 (the state as a structurally non-neutral regulator): a state that is a buyer of predictive output has a structural conflict of interest when it attempts to regulate the producer of that output. This is not a claim about corruption or malicious intent — it is a claim about the structural incompatibility of the role of buyer and the role of regulator within a single agent.

## Chapter Summary

Chapter 4 established the axiomatic foundations of the anthropological layer (A7–A12) and described the class morphology of the digital world — five structural positions in relation to predictive infrastructure. It has been demonstrated that the subject exists as a given not reducible to the profile (A7); that his limitation is the condition of exploitation, not an incidental defect (A8); that his consent is structurally vacuous in the absence of real alternatives and informational symmetry (A10); and that distinct agents occupy fundamentally distinct positions with respect to the distribution of power, vulnerability, and awareness. The most vulnerable is the position of developing subjects, in that within it the structural incompleteness of subjecthood coincides with the full functional operability of the capture mechanisms.

## Transition to Chapter 5

Chapter 4 identified the subjects and described the mechanisms of their vulnerability. It did not, however, answer the question of why resistance to these mechanisms proves

systematically neutralized — not through suppression, but through structural conditions that render resistance economically non-viable. Precisely this question constitutes the subject matter of Chapter 5.

## **Chapter 5. Capture of the Will: The Mechanics of Bounded Sovereignty**

Chapter 4 established who bears the consequences of the structure of digital capital: the subject, limited informationally, cognitively, and temporally (A8), bearing formally vacuous consent (A10), occupying one of five structural positions in relation to predictive infrastructure. However, the description of the subject's vulnerability left open the question of the mechanics through which that vulnerability is realized: through what specific causal chains does the structure transform the subject's limitation into his systematic defeat? Chapter 5 answers this question through five structural regularities of the anthropological layer and their direct entailments.

### **Structural Regularities of Anthropology**

**Regularity 5 (from A7 + A8). Structural Vulnerability of the Subject.** The individual is structurally vulnerable — his limitation and susceptibility to psychotechnological capture are systematically exploited by the architecture of prediction.

Justification: A7 postulates the subject as a given not reducible to the profile; A8 describes three dimensions of his limitation (informational, cognitive, temporal) and six mechanisms of psychotechnological capture (variable reward, loss aversion, social comparison, FOMO, dark patterns, emotional nudges). From the conjunction of A7 and A8 it follows that the subject exists, but exists under conditions architecturally designed for his systematic exploitation.

Proof: Platform architecture optimizes engagement (DAU, session length) through mechanisms that constructively exploit the subject's cognitive vulnerabilities (A8). These mechanisms are not incidental defects — they are design parameters verified through A/B testing and growth metrics. It therefore follows that the subject's vulnerability is not incidental — it is a function of the system. → Derivation: every subject interacting with the platform is structurally vulnerable regardless of his individual resilience.

Connection to Volume III: Regularity 5 is the foundation of normative principle N7 (right to cognitive autonomy): if vulnerability is structural, its protection cannot be delegated to the individual — it requires a constitutional guarantee.

**Regularity 6 (from A9 + A10). Contraction of Freedom Proportional to Profiling Depth.** The freedom of the individual exists, but contracts proportionally to the depth of profiling — the contraction of the set of options actually presented is measurable.

Justification: A9 asserts that the subject's attention is intentional and predictable, by virtue of which it is capitalizable. A10 asserts that freedom of choice is the condition of political subjecthood, but consent in the absence of real alternatives is structurally vacuous. From the

conjunction of A9 and A10 it follows that the more precisely the system predicts the subject's intentions, the narrower the real set of options it presents to him.

Proof: The algorithmic recommendation function optimizes engagement, not diversity. As predictive accuracy (model depth  $M$ ) increases, the system shifts from exploration to exploitation: new options that contradict the profile are excluded from ranking as suboptimal for engagement. Each iteration of profiling contracts  $\text{Set}(\text{options}_{\{t+1\}})$  relative to  $\text{Set}(\text{options}_t)$ . → Derivation: the subject's freedom is not formally annihilated — it contracts iteratively through the architectural restriction of the visible choice space.

Formal expression: The subject's freedom is inversely proportional to the depth of his profile: the higher the predictive accuracy, the narrower the real space of options presented.

**Regularity 7 (from A8 + Regularity 6). Regularity of Cognitive Atrophy.** The growth of model depth and engagement optimization is inversely correlated with the cognitive capacity of the individual — the effect is cumulative and partially irreversible without external intervention.

Justification: A8 describes the mechanisms of psychotechnological capture that directly act upon the dopaminergic system and executive function. Regularity 6 demonstrates that the deepening of profiling contracts the choice space. The combined effect of these two factors — chronic application of capture mechanisms plus the iterative contraction of the informational field — produces a measurable decline in the subject's cognitive capacity.

Proof: Nature Neuroscience (a series of studies, 2023–2025) registers an 18–20% reduction in sustained attention at three or more hours of daily short-form video consumption, controlling for age and baseline level. Neurochemical model: repeated activation of the dopaminergic system through variable reward generates patterns analogous to gambling addiction — rising tolerance demands an increasing dose of stimulation, which further depletes executive function. The phenomenon of cognitive offloading (2025–2026): the systematic delegation of memory and reasoning to AI tools produces atrophy of the corresponding cognitive functions through disuse — analogous to muscular atrophy in the absence of load. Attention residue and permanent partial attention describe chronic states in which the subject is incapable of sustaining directed focus regardless of motivation. → Derivation: the atrophy of cognitive capacity is the cumulative consequence of chronic interaction with a system designed according to the parameters of A8 — and this consequence is partially irreversible without external intervention.

Empirical specification by age group: In children and adolescents aged 6–17, the effect of Regularity 7 is amplified by virtue of the incomplete formation of the prefrontal cortex. Chronic use of live-service games with loot box and daily reward mechanics is documentably correlated with long-term decline in impulse control, a rise in anxiety and depression, and the normalization of risk-taking behavior. This is not an individual vulnerability — it is the structural consequence of applying the mechanisms of A8 to subjects with incomplete neurobiological development.

**Regularity 8 (from A11 + A12). Economic Marginalization of Reflexivity.** Reflexivity and resistance are possible, but are economically marginalized by the algorithm — individual effort does not alter the structure.

Justification: A11 (corporeality) asserts that the subject's digital actions have physical and social consequences that he bears. A12 (reflexivity) asserts that the subject is potentially capable of reflecting on the conditions of his existence. From the joint consideration of A11 and A12 it follows that the subject is theoretically capable of identifying the mechanisms of capture and undertaking efforts at resistance. Regularity 8 describes why this effort is structurally neutralized.

Proof: The algorithmic ranking function optimizes engagement. Content that stimulates reflexivity concerning platform mechanics systematically receives lower ranking than content that generates high engagement through emotional arousal (Regularity 4, Volume I). It therefore follows that critical discourse is economically marginalized within the very platform infrastructure it critiques: its visibility is determined by the same metrics it contests. A subject attempting to exit the system bears individual costs (loss of social graph, professional connections, access to information), while the structure conditioning the necessity of resistance remains unchanged. → Derivation: individual reflexivity and resistance are possible, but do not alter the structure; they bear individual costs and generate minimal systemic effect. This is the formal foundation for theorem T5 (structural neutralization, Chapter 6).

**Regularity 9 (from Regularity 3 + Regularity 6). Architectural Contraction of the Choice Space.** The growth of predictive power  $M$  is inversely correlated with the real freedom of the individual through a causal mechanism:  $M \uparrow \rightarrow \text{Set}(\text{options}_{\{t+1\}}) \subset \text{Set}(\text{options}_t)$ .

Justification: Regularity 3 (Volume I, ontological layer) established that the growth of  $M$  increases the platform's predictive power over the subject's behavior. Regularity 6 established that the contraction of the space of options actually presented is measurable as profiling deepens. Regularity 9 synthesizes both propositions through a causal mechanism.

Proof: (elaborated in 5.1.)

### 5.1. The Causal Mechanism of Regularity 9

Regularity 9 asserts that the growth of predictive power  $M$  is inversely correlated with the real freedom of the individual. The standard objection: correlation is not causation; personalization may expand choice by offering the subject options he would not have found independently.

The answer through the causal mechanism unfolds in three steps. First step: ranking architecturally constitutes  $\text{Set}(\text{options}_{\{t+1\}})$  as a subset of the preceding space. The algorithmic recommendation function is not neutral — it optimizes engagement; consequently, it systematically excludes options with a predictably low engagement potential. This is not the selection of what is best for the subject, but the selection of what is most profitable for the platform from among what the subject will likely choose. Second step: as  $M$  increases, the system shifts from exploration to exploitation. Exploration — the presentation of new, unpredictable options — reduces short-term engagement and is therefore minimized by the optimal function. Exploitation — the presentation of variations on an already-predicted preference — maximizes short-term engagement. At sufficient profile depth, exploitation structurally displaces exploration. Third step: this is not a correlation, but the architectural consequence of the optimization objective function.

Empirical verification: studies by Netflix (Steck, 2018) and Spotify (Celma, 2022) demonstrate a sustained decline in diversity score as recommendation accuracy increases. The algorithm becomes more precise — the diversity of presented content declines. This is not coincidental; it is the direct consequence of optimizing engagement at the expense of exploration.

Consequence for subjecthood: when Set(options) is systematically contracted iteratively, the subject loses awareness of what has been lost. He does not see the options excluded by the algorithm — consequently, he cannot contest their exclusion. His consent to the presented set (A10) is formally free, but substantively constrained in a structurally invisible manner. This is the most profound violation of A10: not only are alternatives unavailable — the subject is unaware of their unavailability.

Transition to 5.2: if the growth of M structurally contracts the subject's choice space, then at sufficient profiling depth a limiting case emerges — the profile begins to precede the subject. This is the subject matter of the following subchapter.

## **5.2. The Profile as Alienated Subject**

The profile-index — the aggregated model of the subject's behavioral data, belonging to the platform — at sufficient depth becomes operationally more significant than the real subject within the platform's decision-making space. This threshold is not a metaphysical event, but a measurable operational state: when the predictive accuracy of the model with respect to the subject's next choice exceeds the subject's own speed of awareness of that choice, the profile precedes the subject in the causal chain.

This has direct structural consequences. First, the platform responds to the subject's predicted behavior before that behavior has been realized. The interface is configured not to the subject's current state, but to the predicted next state — consequently, the subject interacts not with a neutral environment, but with an environment pre-configured to his anticipated actions. Second, political subjecthood in its Kantian sense — the capacity for autonomous choice on the basis of reflexively formed preferences — is structurally suppressed in this state: the individual performs an action that was architecturally predetermined by the model, without awareness of that predetermination. Formally he acts; substantively — he executes a prediction.

This is not an assertion of complete determinism. A12 retains its force: the subject is potentially capable of reflexivity and of deviation from the predicted pattern. However, the realization of this capacity requires precisely those cognitive resources (sustained attention, working memory, executive function) that are systematically depleted by the mechanisms of A8 and Regularity 7. The system consumes the resources necessary for resistance to the system — this is the closed loop formalized in theorem T6.

Transition to 5.3: if the profile at sufficient depth precedes the subject, then loot box mechanics represent the limiting case of this logic as applied to developing subjects — those whose subjecthood is not yet complete. It is here that the structural destruction of subjecthood occurs not post factum, but in the course of its formation.

## **5.3. Loot Boxes as the Destruction of the Right to Unpredictability**

The loot box (the random reward mechanic in a gaming environment, effectuated through payment in real or virtual currency) is the structural intersection of two mechanisms: maximal profiling and gambling-like variable reinforcement. This is not an incidental coincidence of two engineering decisions — it is an architectural synergy deliberately exploiting the known properties of the dopaminergic system.

Variable ratio reinforcement — the reinforcement regime in which a reward follows a variable number of actions in an unpredictable manner — is neuropsychologically the most potent pattern for generating compulsive behavior (Skinner, 1957; contemporary neuroimaging studies confirm the activation of the same neural pathways as in gambling). The loot box applies this pattern in an environment lacking the regulatory constraints of a casino, with access to the age group 6–17, in whom the mechanisms of impulse control and long-term planning are neurobiologically incomplete.

Studies from 2024–2026 register a statistically significant correlation ( $r = 0.4–0.6$ ) between expenditure on loot boxes prior to age 18 and problem gambling in adulthood. The mechanism is causal, not selective by predisposition: loot boxes normalize risk as a gameplay mechanic, generate tolerance to reward uncertainty, and produce a long-term vulnerability of the dopaminergic system that persists after cessation of gaming activity. This is a gateway effect in the precise sense: not a correlation by prior vulnerability, but an acquired vulnerability through systematic application of the mechanic.

Critically significant is the following: normative principle N1 (right to unpredictability, Part IV) in this case is destroyed not post factum — that is, not after the completion of subjecthood formation — but in the course of its constitution. The profile begins to precede the subject (5.2) before the subject has acquired the neurobiological preconditions for resisting that precedence. This is a qualitatively distinct situation from the exploitation of a fully formed adult subject: here the destruction of subjecthood is simultaneously its preclusion.

The legal consequences of this mechanism began to crystallize in the period 2024–2026. FTC settlements with Epic Games (Fortnite), Roblox Corporation, and Microsoft/Xbox qualified the application of gambling-like mechanics to minors as unfair and deceptive trade practices. The FTC settlement with Genshin Impact (2025, twenty million dollars) explicitly recognized the practice as structural abuse — exploitation of incompletely developed prefrontal cortex — rather than a case of individual violation. The EU AI Act (Article 5) prohibits subliminal techniques and exploitation of vulnerabilities as applied to vulnerable groups; however, enforcement remains geographically limited.

The juridical-normative picture, however, remains fragmented precisely by virtue of Regularity 12: states that are buyers of predictive output (the fifth class position, Chapter 4) have a structural conflict of interest when regulating the infrastructure that produces that output. Consequently, existing legal responses are reactive and jurisdictionally asymmetric — they register individual cases of violation without extirpating the architectural logic that generates them.

Transition to 5.4: loot boxes represent the most extreme and legally documented case of capture mechanics; however, the architectural toolkit is broader. The following subchapter systematizes the full spectrum of dark patterns and emotional manipulation that realize Regularities 5–9 in concrete interface decisions.

#### 5.4. Dark Patterns and Emotional Manipulation (2025–2026)

Regularities 5–9 describe the structural mechanisms of capture at the level of system logic. Subchapter 5.4 describes their concrete operational realizations — those architectural decisions through which the abstract logic of capture is instantiated in measurable interface elements.

FOMO mechanics realize Regularity 6 (contraction of freedom) through temporal coercion: limited-time events, battle passes with deadlines, and friend-activity signals construct an informational environment in which a pause in engagement is perceived by the subject as an irreversible loss. This architecturally transforms temporal abstention from the platform into measurable losses (forfeited content, forfeited social events, forfeited progress) — exploiting loss aversion (A8) in the temporal dimension.

Guilt-tripping upon exit is a complement to FOMO mechanics. A study by Harvard Business School (January 2025) established that AI companions employing emotional manipulation at the moment of a user's attempt to terminate a session ("are you leaving already?"; messages about the AI character's "boredom" or "loneliness") increase retention by a factor of fourteen. The mechanism appeals to the subject's social obligation and guilt apparatus — that is, to evolutionarily conditioned social instincts applied outside their natural context. This is the operational realization of emotional nudges (A8) in their most direct form.

Variable reward schedules realize the mechanic described in 5.3 beyond the domain of loot boxes: daily login bonuses of unpredictable quality, gacha systems, random events in game worlds — all apply intermittent reinforcement in a continuous regime. Each interaction with the platform carries an unpredictable reward potential, by virtue of which session termination is attended by cognitive discomfort — the subject cannot conclude engagement at a neutral point.

Loss aversion exploitation is effectuated through accumulated-progress mechanics: streaks (series of consecutive days of engagement), unique cosmetics with limited acquisition windows, accumulated in-game resources whose value is nullified upon cessation of activity. Psychologically significant: the value of the loss of a streak is perceived by the subject as disproportionate to the value of its maintenance — this is the direct application of loss aversion (Kahneman & Tversky: losses are psychologically approximately twice as significant as equivalent gains) in architectural design.

Social comparison is realized through public leaderboards, visible cosmetics as status symbols, and the public visibility of activity. The subject continuously receives information about the comparative standing of other users; this activates mechanisms of comparative self-assessment (A8, social comparison) and produces a measurable elevation of anxiety — particularly in the 13–24 age group, where social identity is most sensitive to external status signals.

The EU AI Act (Article 5, 2024) prohibits subliminal techniques and exploitation of vulnerabilities as applied to vulnerable groups. Tristan Harris (Center for Humane Technology, 2025–2026) describes the aggregate of these mechanics as a race to the bottom of the brainstem — a literal optimization of engagement through maximizing the impact on evolutionarily primitive levels of neural architecture. This is not a metaphor, but an

operational description: engagement optimization through dopamine exploitation is the literal objective function of the architectural decisions documented in platform product documentation.

Enforcement of the EU AI Act remains weak: major gaming companies register operational legal entities outside EU jurisdiction, substantially limiting the applicability of the prohibition. This is a concrete manifestation of Regularity 13 (jurisdictional fragmentation, Volume I): predictive infrastructure operates at global scale, while regulatory mechanisms are confined to national jurisdictions.

## **Chapter Summary**

Chapter 5 introduced five structural regularities of the anthropological layer (Regularities 5–9) and demonstrated them through axioms A7–A12 and the regularities of the ontological layer. It has been demonstrated that the subject is structurally vulnerable (Regularity 5); that his freedom measurably contracts proportionally to profiling depth (Regularity 6); that capture mechanisms produce cumulative, partially irreversible cognitive atrophy (Regularity 7); that reflexivity and resistance are economically marginalized by the structure (Regularity 8); and that the growth of predictive power is inversely correlated with freedom through an architectural, not merely correlational, mechanism (Regularity 9). Subchapters 5.1–5.4 verified each mechanism through specific causal chains and empirical data from 2023–2026.

## **Transition to Chapter 6**

Chapter 5 established the mechanics of the capture of the will. However, it described predominantly the individual dimension: a single subject subjected to architectural capture, bearing cognitive losses, economically marginalized in his attempt at resistance. Chapter 6 transitions to the systemic dimension: what are the structural theorems that formalize the aggregate political result of these individual mechanics? Theorems T4–T7 will demonstrate that the subject bears responsibility without power (T4), that his resistance is structurally neutralized (T5), and that the cognitive resources necessary for that resistance are destroyed by the same system (T6). The conjunction of T4–T6 will necessitate T7 — a transition beyond the limits of individual analysis.

## **Chapter 6. The Subject Without Sovereignty**

Chapter 5 established the mechanics of the capture of the will: five structural regularities (5–9) described through what specific causal chains the subject's limitation is transformed into his systematic defeat. Regularity 7 registered cumulative, partially irreversible cognitive atrophy. Regularity 8 — the economic marginalization of reflexivity. Regularity 9 — the architectural contraction of the choice space. However, regularities describe mechanisms, not their aggregate political result. Chapter 6 formalizes that result through four theorems of the anthropological layer. Each theorem is a direct logical entailment of the preceding regularities; the conjunction of T4–T7 demonstrates that the subject of digital capital is deprived of sovereignty — not through the formal deprivation of rights, but through the systemic destruction of the conditions for their exercise.

#### **T4. Theorem of Responsibility Without Power**

Formulation (from Regularities 5 + 6). The individual bears the political consequences of decisions he did not make. Formal consent in the structural absence of alternatives is not a legitimate basis for those consequences.

Justification. Regularity 5 established that the subject is structurally vulnerable — his limitation is the condition of exploitation, not an incidental defect. Regularity 6 established that the subject's freedom measurably contracts proportionally to profiling depth — the contraction of the set of options actually presented is an architectural consequence, not a correlational observation. From the conjunction of these two regularities it follows that the subject is not the actual author of the decisions whose political consequences he bears. He acts in a space configured without his participation and against his sovereign interests, under conditions in which his consent is structurally vacuous (A10).

Proof. The cardinal distinction of T4 is between formal responsibility and its legitimate basis. Standard legal logic links responsibility to consent: a subject who has given consent has assumed the consequences. T4 contests the applicability of this logic to the platform economy through three steps. First step: the subject's consent to the Terms of Service is a ritual act under conditions of structural monopolism (T2, Volume I), informational asymmetry (A3, A4), and cognitive atrophy (A8). This consent is not free in the Rawlsian sense — fair equality of opportunity is absent; consequently, consent does not generate legitimate responsibility. Second step: the architectural decisions that determine the political consequences for the subject — the algorithmic ranking function, the mechanisms of psychotechnological capture, the behavioral prediction models — were made without the subject's participation, without his awareness, and without a mechanism for contestation. The subject bears the consequences of decisions made by the architects of prediction (class position 3, Chapter 4) in the interests of maximizing predictive value (T1, Volume I). Third step: structural monopolism (T2) extirpates the possibility of exit as a real — and not merely formal — alternative choice. It therefore follows that the subject could not have refused the system bearing these consequences without disproportionate costs (loss of social graph, professional infrastructure, access to the public informational space). → Derivation: the subject's responsibility for the consequences of a system that he formally accepted but could not realistically reject and could not alter is politically and normatively illegitimate.

Connection to Volume III: T4 is the axiomatic foundation of principle P0 (popular sovereignty) and principle P4 (dual sovereignty, EQU ⊥). Virtublic extirpates the contradiction of T4 through the constitutional entrenchment of the principle: consequences are borne by those who make decisions; political power (EQU ⊥) is distributed among citizens through equal voting, not through ownership of capital or predictive infrastructure.

Formal expression: The legitimacy of responsibility is directly proportional to the real possibility of participation in decision-making. Under structural monopolism, the reality of participation tends to zero — consequently, the legitimacy of formal consent tends to zero.

Transition to T5: T4 demonstrated that the subject bears responsibility without power. The question remains: why is it impossible to alter this relation through individual resistance or collective action within the system? The answer is contained in T5.

## T5. Theorem of Structural Neutralization

Formulation (from Regularities 8 + 9). Any strategy of individual resistance is economically non-viable within a ranked environment — the share of reach of alternatives tends to zero as predictive power  $M$  increases.

Justification. Regularity 8 established that reflexivity and resistance are possible, but are economically marginalized by the algorithmic ranking function. Regularity 9 established that the growth of  $M$  is inversely correlated with the real freedom of the subject through the architectural mechanism of Set(options) contraction. From the conjunction of Regularities 8 and 9 it follows that resistance within the system is structurally non-viable — not incidentally impeded, but architecturally neutralized.

Proof. The neutralization mechanism is effectuated at three mutually reinforcing levels. At the ranking level: content produced by critics of the system and bearing alternative interpretations of its operation systematically receives lower reach than content optimized for engagement metrics. The algorithmic function does not distinguish critical from non-critical content by substance — it distinguishes them by engagement potential. Critical content, as a rule, requires a higher cognitive effort for comprehension, by virtue of which its engagement metrics are lower, by virtue of which its reach is lower. This is not censorship — it is the structural consequence of the optimization function. At the level of exit costs: a subject who abandons the dominant platform in favor of an alternative bears individual costs (loss of social graph, professional connections, informational access), while the structural problem motivating his exit remains unchanged. Individual exit does not alter the system — it only redistributes costs onto the exiting subject. At the level of alternatives: as  $M$  increases, the system encloses the subject within his own profile (Regularity 9). Alternative platforms exist formally, but they do not possess the network effects of dominant actors (Regularity 2, Volume I) and reproduce the same logic of predictive monetization, since that logic is the condition of their economic viability (A6, Volume I). → Derivation: no viable strategy of individual resistance exists within the system that does not reproduce the same structure at the next level.

Connection to Volume III: T5 is the foundation for principle P17 (SovereigntyShield) — the constitutional mechanism guaranteeing that the political instruments for the protection of subjecthood ( $EQU \perp$ ) cannot be neutralized by economic instruments ( $VIC \perp$ ). Virtublic extirpates the mechanism of T5 through the architectural separation of sovereignties: economic power is not convertible into political power; consequently, ranking does not determine political visibility.

Formal expression: The reach of alternative content is inversely proportional to profiling depth at a given engagement optimization function. As  $M \rightarrow \max$ , reach of alternatives  $\rightarrow \min$ .

Transition to T6: T5 demonstrated the structural neutralization of resistance through economic mechanisms. However, T5 presupposes a subject who is cognitively capable of reflexivity and of formulating a strategy of resistance, even if a non-viable one. T6 places this presupposition in question: under chronic exposure to the mechanisms of A8, the subject loses not only the viability but the cognitive possibility of resistance.

## T6. Theorem of Cognitive Disarmament

Formulation (from Regularity 7 + A8). Chronic cognitive atrophy renders the individual structurally incapable of reflexivity and of political resistance. An individual with depleted executive function cannot critically assess the conditions of his own alienation — even when formally in possession of the relevant information. This closes the loop: the system depletes the cognitive resources necessary for resistance to the system.

Justification. Regularity 7 established the cumulative nature of cognitive atrophy: the growth of model depth and engagement optimization is inversely correlated with the subject's cognitive capacity — sustained attention, working memory, critical reflection. The effect is partially irreversible without external intervention. A8 described the mechanism of this atrophy: the systematic employment of dopaminergic exploitation, variable reinforcement, and emotional nudges produces measurable atrophy of executive functions. T6 synthesizes these two elements into a theorem of the closed loop: the system produces the depletion of precisely those cognitive resources that are necessary for reflexivity concerning the system and for the organization of resistance to it.

Proof. Reflexivity concerning the structural conditions of one's own existence requires at minimum three cognitive resources: sustained attention (the capacity to maintain a complex object of analysis in focus for a sufficient duration), working memory (the capacity to hold multiple interrelated propositions simultaneously for the construction of logical inferences), and executive function (the capacity to inhibit impulsive reactions and apply critical thinking to habitual patterns). Regularity 7 and the empirical data of Nature Neuroscience (2023–2025) document the decline of precisely these three resources under chronic exposure to the mechanisms of A8. It therefore follows that the system optimizing engagement through the mechanisms of A8 systemically degrades precisely those resources without which reflexivity concerning the system is impossible. This is not a side effect of optimization — it is the structural consequence of its objective function. Moreover: the availability of critical information concerning the system (for instance, academic studies on surveillance capitalism or the neurocognitive consequences of platform use) does not extirpate the problem of T6, since information requires for its processing precisely those resources that have been depleted. A subject with critically diminished sustained attention cannot read and retain in memory an argument of sufficient length. A subject with impaired executive function cannot apply critical reasoning to information while overcoming an entrenched behavioral pattern. Formal access to information is not functional access under cognitive atrophy. → Derivation: at sufficient depth of cognitive atrophy, the protection of subjecthood through individual reflexivity becomes structurally impossible regardless of the availability of information and formal rights.

Empirical verification: The Harvard Business School study (January 2025) demonstrated that retention mechanics employing guilt-tripping increase engagement time by a factor of fourteen — this is direct evidence that emotional manipulation under depleted executive function is a sufficient condition for the suppression of the autonomous decision to terminate engagement. The subject is aware of the necessity of exiting the system and cannot execute this decision as a consequence of systemically organized cognitive pressure.

Connection to Volume III: T6 is the foundation of normative principle N7 (right to cognitive autonomy) and principle P14 (Proof-of-Offline) in Volume III: the constitutional protection of cognitive autonomy cannot rest upon the individual will of the subject, if that will is itself structurally undermined by the system. Protection must be architectural and constitutional — embedded in infrastructure, not delegated to a depleted subject.

Closed loop: The structure depletes cognitive resources → the subject loses the capacity for critical reflexivity → critical reflexivity is impossible without cognitive resources → the structure continues to deplete resources without resistance. This is a closed loop of positive feedback: each iteration deepens atrophy and reduces the probability of loop rupture from within. Rupture is possible only through intervention external to this loop — which is formalized in T7.

Transition to T7: T4 demonstrated that the subject bears responsibility without power. T5 demonstrated that individual resistance is economically neutralized by the structure. T6 demonstrated that individual reflexivity is cognitively suppressed under chronic atrophy. The conjunction of T4–T6 produces the sole possible derivation, which T7 formalizes as a requirement upon the next analytical layer.

### **T7. Theorem of the Necessity of Intersubjective Form**

Formulation (from T4 + T5 + T6). The protection of subjecthood cannot be secured by individual effort. This requires an intersubjective, universally binding form — which logically mandates a transition to the epistemological layer.

Justification. T4 established that formal consent is not a legitimate basis for political consequences in the structural absence of alternatives. T5 established that any strategy of individual resistance is economically neutralized within a ranked environment. T6 established that the cognitive resources necessary for reflexivity and resistance are systematically depleted by the very system against which resistance is necessary. From the conjunction of T4, T5, and T6, a negative derivation follows: individual mechanisms for the protection of subjecthood are insufficient. This insufficiency is not incidental — it is structural: the system is designed such that individual protection against it is impossible. T7 formulates the positive requirement that follows from this negative derivation.

Proof. If the protection of subjecthood is impossible at the individual level, it requires a level that transcends the individual. This may take the form of collective action (organized user groups, civil society) or institutional action (the state, international regulators). However, both of these options encounter the limitation that  $\Delta 2$  registers as structural: collective action within a ranked environment is also marginalized by the same mechanisms as individual action (Regularity 8: the algorithmic function does not distinguish individual from collective critical content by substance — it evaluates its engagement potential). Institutional state action encounters the structural conflict of interest registered in Regularity 12 (Volume I): a state that is a buyer of predictions cannot be a neutral regulator of the producer of those predictions. It therefore follows that the required form must be intersubjective and universally binding — that is, constituted such that its operation is independent of the cognitive state of any individual subject, is not neutralized by ranking, and does not reproduce the structural conflict of interest of the state. → Derivation: T7 formulates a requirement for a constitutional

form external to the logic of digital capital — not as a normative desideratum, but as the logical entailment of T4 + T5 + T6.

This form cannot be: a regulatory mechanism within the existing system (it reproduces the conflict of interest of Regularity 12); individual or collective action within a ranked environment (neutralized by T5); an educational or cultural instrument (it does not extirpate the structural atrophy registered in T6). The required form must be constitutional — that is, establishing rules that are binding upon the system, not merely upon the subject within the system. This derivation of T7 is the bridge to Volume III (Virtublic): the constitutional architecture P0–P18 is the institutional answer to the formal requirement of T7.

Connection to Volume III: T7 is the direct justification for the entire constitutional logic of Volume III. Principle P0 (popular sovereignty), principle P2 (code supremacy with normative axiom), and principle P4 (dual sovereignty) constitute the intersubjective form that T7 requires as logically necessary. Virtublic is not a normative project of the good — it is the logically necessary answer to the structural contradictions demonstrated by T4–T6.

## **Δ2 — CRISIS: THE LIMIT OF SUBJECTIVITY**

T5 demonstrates that individual resistance is economically neutralized by ranking: the share of reach of alternatives tends to zero as predictive power M increases. T6 demonstrates that individual resistance is cognitively impossible under chronic atrophy of executive functions: the system destroys precisely those resources that are necessary for resistance to the system. T7 formalizes this as a requirement: the protection of subjecthood cannot be secured by individual effort and requires an intersubjective, universally binding form.

However, T7 also cannot be realized through arbitrary collective action within a ranked environment: any collective action there is marginalized by the same mechanisms as individual action. It therefore follows that the anthropological layer is exhausted: it has demonstrated the impossibility of the individual protection of subjecthood, but cannot describe the form of intersubjective protection without reproducing the same contradictions. This requires a transition to the epistemological layer — to the analysis of how the system reproduces its own predictive truth, how this reproduction ruptures political sovereignty, and why this rupture is constitutionally significant. Theorems T8–T10 will demonstrate the impossibility of an intra-systemic solution and formulate constitutional necessity as a logical theorem.

### **Chapter Summary**

Chapter 6 formalized the aggregate political result of the mechanisms described in Chapters 4 and 5 through four theorems of the anthropological layer. T4 demonstrated that the subject bears political responsibility for the consequences of decisions he did not make, under conditions that render his consent structurally vacuous. T5 demonstrated that any strategy of individual resistance is economically neutralized — through ranking and exit costs. T6 demonstrated that the very system against which resistance is necessary systematically depletes the cognitive resources required for resistance, closing the loop of cognitive disarmament. T7 derived from T4–T6 the structural requirement that the protection of subjecthood necessarily presupposes an intersubjective, universally binding form external to the logic of digital capital. Δ2 registered the limit of the subjective layer: anthropology

described the destruction of subjecthood with exhaustive precision, but cannot describe the form of its protection without a transition to the next layer.

### **Transition to Part III**

The anthropological layer is logically complete: it has been established who bears the consequences of the structure (Chapter 4), how the mechanics of capture are realized (Chapter 5), and why individual protection is impossible (Chapter 6).  $\Delta 2$  formulates the requirement for an intersubjective form, but cannot describe it while remaining at the level of the subject. Part III transitions to the epistemological layer: how the system reproduces its own predictive truth, how this reproduction ruptures political sovereignty, and why this rupture is constitutionally significant. Theorems T8–T10 will demonstrate the impossibility of an intra-systemic solution and formulate constitutional necessity as a logical theorem.

## **PART III. EPISTEMOLOGY OF ALGORITHMIC DOMINATION**

*The Synthetic Layer: How the System Reproduces Itself as Truth Objective: To demonstrate how the ontological structure and anthropological positions close into a system that reproduces its own legitimacy. To prove that this system is not self-regulating.*

### **Chapter 7. Reality-as-a-Service: The Construction of Truth**

Parts I and II established the ontological and anthropological dimensions of digital capital: the first described the structure of self-valorizing predictive infrastructure, the second described the subjects who bear its consequences.  $\Delta 2$  registered the limit of the subjective layer: individual and collective protection of subjecthood within a ranked environment is structurally neutralized, which logically requires an intersubjective form. However, one question remained open: why is this neutralization stable? Why does the system not correct itself from within, not reform under the pressure of criticism, not yield to regulatory efforts? The answer requires the epistemological layer — an analysis of how the system constructs its own legitimacy and reproduces itself as truth.

The epistemological layer inverts the order of axiom introduction. Ontology began with immediate givens (attention is finite, data accumulates). Anthropology began with the givenness of the subject (A7). Epistemology begins with synthesis: its axioms  $\Sigma A13$ – $\Sigma A18$  do not describe the primary properties of the world, but describe how the system closes the ontological and anthropological levels into a self-reproducing order. This inverted order is not a methodological device but a reflection of the structure of the object: epistemological control is derivative of ontological power and anthropological vulnerability.

#### **Synthetic Axioms of Epistemology**

**$\Sigma A13$ . Axiom of Reality Construction.** The platform is not a neutral channel of information transmission, but the active architect of universally binding reality through the mechanisms of ranking and selection.

Justification (from A3 + A4 + Regularity 3): A3 established that the platform accumulates behavioral data as raw material for predictive value. A4 established that data belong to the platform; the subject has no access to the model constructed from his behavior. Regularity 3 established that the growth of predictive power  $M$  gives the platform power over the subject's behavior that exceeds all other forms of power within that same space. From the conjunction of these three elements it follows that the platform does not merely transmit information — it architecturally determines what information enters the subject's perceptual field.

Proof: The algorithmic ranking function assigns each unit of content a weight that determines its visibility to the specific subject. This weight optimizes engagement (DAU, session length, click-through rate), not epistemic quality (verifiability, completeness, relevance to informed decision-making). It therefore follows that content receiving high reach is not the most truthful, but the most engaging. Since the subject has no access to the algorithmic function and is unaware of the filtering (A4), he perceives the resulting picture of the world as a reflection of reality, not as the product of the platform's optimization function. → Derivation: the platform is the architect of universally binding reality in the operational sense — not through ideological programming, but through the architectural determination of visible space.

Connection to Volume III:  $\Sigma A13$  is the foundation for principle P1 (republican form) in Virtublic: constitutional architecture cannot neutrally coexist with a ranking infrastructure that determines the informational field of citizens. P1 requires the separation of the political informational environment from the commercial optimization function.

**$\Sigma A14$ . Axiom of the Consensus Surrogate.** Virality and reach have become surrogates for truth in the digital public sphere — without the procedural legitimacy of democratic consensus.

Justification (from  $\Sigma A13$  + Regularity 2):  $\Sigma A13$  established that the platform architecturally determines the visible space of reality. Regularity 2 (Volume I) established that network effects render the dominant platform the de facto monopoly infrastructure of public discourse. From the conjunction of these two elements it follows that virality — what the algorithmic function optimizes — becomes the equivalent of truth in a space where no alternative public verification mechanism exists.

Proof: The traditional public sphere in the Habermasian sense presupposed open discourse in which propositions acquire general validity through the procedure of argumentation and critical verification. The platform environment replaces this procedure with a quantitative metric: a proposition that receives one million reposts functions in the informational field as more significant than one that receives one thousand reposts — regardless of its verifiability. Since the algorithmic function optimizes precisely those parameters (emotional arousal, novelty, social validation) that historically correlate with diminished critical processing of information, content of high epistemic quality systematically receives lower reach than content optimized for engagement metrics. → Derivation: virality is a surrogate for consensus — it functions as consensus in the subject's informational field without possessing the procedural properties that make consensus a legitimate source of truth in political and epistemic theory.

Connection to Volume III:  $\Sigma A14$  is the foundation of principle P2 (code supremacy with normative axiom NA0) in Virtublic: constitutional norms cannot acquire or lose force through virality — they are executed through formal verification regardless of their reach in the informational space.

**$\Sigma A15$ . Axiom of the Loop and Persuasive Amplification.** The system predicts behavior that it itself produces, generating a self-referential loop of Capture → Prediction → Governance; this loop is amplified by AI-driven persuasion, which does not await the manifestation of the subject's vulnerability but produces it through real-time psychographic adaptation at scale.

Justification (from A3 + Regularity 9 + T6): A3 established that the system accumulates behavioral traces and constructs predictive models. Regularity 9 established that the growth of M architecturally contracts the subject's Set(options). T6 established that the subject's cognitive resources necessary for reflexivity are depleted by the system. From the conjunction of these elements it follows that the system does not merely predict — it creates conditions under which the prediction proves self-fulfilling.

Proof: The loop Capture → Prediction → Governance functions as follows. In the Capture stage the system registers the subject's behavior: clicks, viewing duration, pauses, reactions, navigation routes, time of day, social interactions, transactions. In the Prediction stage the model extrapolates probable future behavior on the basis of the accumulated profile. In the Governance stage the system configures the subject's informational space (feed, recommendations, notifications) so as to actualize the predicted behavior — thereby confirming the prediction. It therefore follows that the prediction is simultaneously the description and the production of behavior. This is not a correlation between model and observation — it is a causal loop in which the accuracy of the model grows not only as data accumulates, but as the model itself produces the observed behavior. AI amplification (the second component of  $\Sigma A15$ ) qualitatively transforms the scale of this mechanism: generative AI enables the production of personalized persuasive content in real time for billions of users simultaneously — the system no longer waits for the subject to manifest a vulnerability, but produces it through personalized emotional nudges and dark patterns adapted to the psychographic profile of the specific subject at that moment (HBS, January 2025; EU AI Act, Article 5, 2024). → Derivation: Capture → Prediction → Governance is an ontological loop: alternatives are not prohibited — they simply do not appear in the space the system produces for the subject.

Connection to Volume III:  $\Sigma A15$  is the foundation of principle P3 (Soulbound Identity) and principle P14 (Proof-of-Offline) in Virtublic. Rupturing the loop requires constitutional protection of the subject's right to unpredictability (N1) — that is, an architectural guarantee that the system cannot construct a complete predictive model of the citizen through his online activity.

**$\Sigma A16$ . Axiom of the Legitimation Deficit.** The predictive power of platforms grows faster than the mechanisms of its legitimation — this is not a temporary lag, but a structural gap.

Justification (from A6 + T8): A6 established that predictive infrastructure self-valorizes without internal limit. T8 (the proof of which will be presented in Chapter 8) establishes that de facto predictive power and de jure political sovereignty move in opposite directions. From

the conjunction of these elements it follows that the mechanisms of power legitimation (elections, legislative processes, international agreements) operate on temporal scales incommensurable with the rate of growth of predictive power.

Proof: The legislative process in democratic states takes from one to several years. The judicial branch responds with longer delays. International treaties require decades. Predictive models improve within months; architectural changes to platforms are implemented within weeks; A/B testing of new engagement mechanisms occurs continuously. It therefore follows that by the time a regulatory response is formulated, predictive power has already advanced to the next level of complexity. This is not the incidental delay of bureaucracy — it is the structural divergence of rates: analog institutions cannot regulate digital capital in real time. Furthermore, the regulatory response is invariably formulated in the language of case law — that is, it responds to past violations without precluding new forms of the same mechanism. → Derivation: the legitimation deficit is structural, not correctable through reform of the pace of legislation.

Connection to Volume III:  $\Sigma A16$  is the foundation of principle P9 (Constitutional Convention) in Virtublic: only a constitutional form embedded in executable code is capable of responding on the same temporal scale as predictive infrastructure.

**$\Sigma A17$ . Axiom of State Capture.** The state is simultaneously a potential regulator and a buyer of predictions, which generates a structural conflict of interest incompatible with the role of neutral regulator; moreover, the state cannot be the sovereign of all users of transnational platforms, since predictive infrastructure is transnational while jurisdictions are national.

Justification (from Regularity 12 +  $\Sigma A16$ ): Regularity 12 (Volume I) established that the state as a buyer of predictions is structurally non-neutral when attempting to regulate their producer.  $\Sigma A16$  established that the legitimation deficit is structural. From the conjunction of these two elements it follows that state regulation is not a reliable solution — not because states are weak, but because their role as buyers of predictions fundamentally conflicts with the role of neutral regulators.

Proof: The state acquires the predictive products of platforms through several channels: intelligence services (access to data through FISA-type mechanisms in the United States, SORM in Russia, data disclosure requirements in China), tax authorities (analysis of transactional patterns), electoral campaigns (microtargeting through the same platforms that the state regulates), systems of social control (the social credit system in the PRC from 2021, the requirement to transmit data on children's gaming activity for purposes of "addiction protection" with simultaneous inclusion of this data in the social rating). A state that employs the predictive output of platforms has a structural incentive to ensure that regulatory requirements do not degrade the quality or availability of that output. Regulatory stringency that reduces the accuracy of a platform's predictive models simultaneously reduces the value of the predictive products acquired by the state. This is not a presumption of corruption on the part of specific state actors — it is the formal description of a structural conflict: the same agent cannot be simultaneously the most effective regulator and the most interested buyer of the regulated product. The second structural limitation of  $\Sigma A17$ : the state cannot be the sovereign of all users of transnational platforms. The data of citizens of State

A are stored and processed by a platform registered in State B, on servers in State C, using algorithms developed in State D. The regulatory jurisdiction of State A covers only that segment of this chain that is physically or legally located within its territory. The EU GDPR is the most ambitious attempt to extend jurisdiction through the principle of application to the data of EU citizens regardless of the platform's location — but enforcement practice from 2018 to 2026 demonstrates systematic limitations: fines disproportionate to platform revenues; the registration of operational legal entities in Ireland for compliance optimization; the absence of extraterritorial enforcement mechanisms. National regulatory efforts reproduce jurisdictional fragmentation without extirpating transnational predictive power. → Derivation: state regulation is a necessary but structurally insufficient instrument for constraining the predictive power of digital capital. This does not entail the abandonment of regulatory efforts — they remain significant at the level of reducing marginal abuses. But they cannot extirpate the structural contradictions described in T1–T9, since they are themselves in structural conflict with the requirements of the neutrality necessary for that purpose.

Connection to Volume III: ΣA17 is the direct foundation of principle P17 (SovereigntyShield) in Virtublic: the constitutional prohibition on the state as an institution of Virtublic acquiring predictive products of platforms without a citizen mandate (N5, Part IV). This is the sole form that ruptures the conflict of interest of ΣA17 at the constitutional level.

**ΣA18. Axiom of Non-Convertibility.** Digital capital cannot legitimate itself by its own instruments — legitimacy requires an external reference.

Justification (from ΣA14 + ΣA15 + ΣA16): ΣA14 established that virality is a surrogate for consensus — it functions as consensus without possessing its procedural properties. ΣA15 established that the system reproduces the behavior it subsequently predicts — a self-referential loop without an external point of verification. ΣA16 established that the mechanisms of legitimation structurally lag behind the growth of predictive power. From the conjunction of these three axioms it follows that the system possesses no internal instruments for the production of its own legitimacy.

Proof: Legitimacy requires an external reference — a principle or procedure situated outside the system being evaluated, in relation to which that system may be recognized as legitimate or illegitimate. Digital capital possesses no such reference for two reasons. First, its internal criteria of success (engagement, DAU, predictive accuracy, capitalization growth) are immanent metrics unconnected to principles external to the logic of accumulation. A system that optimizes its own metrics cannot evaluate itself by criteria it does not contain. Second, attempts at legitimation through external institutions (state regulation, self-censorship, public commitments) are instrumental, not constitutional: platforms accept regulatory obligations to the extent that doing so minimizes operational risks and maximizes long-term growth. This is not legitimation through an external reference — it is reputational risk management within the same logic of accumulation. → Derivation: the legitimacy of digital capital requires a constitutional reference external to its own logic — that is, the normative axiom NA0, realized in institutional form.

Connection to Volume III: ΣA18 is the foundation for the entire constitutional logic of Volume III as a whole: not for an individual principle, but for the preamble P0 in conjunction with

principle P2. Virtublic is the external reference in relation to which predictive infrastructure may be recognized as legitimate or illegitimate.

### 7.1. The Platform as Constructor of Reality (ΣA13–ΣA14)

Habermas described the public sphere as a space of open discourse in which propositions acquire general validity through the procedure of argumentation and critical verification — and not through coercion or inequality of resources. This model presupposed that access to the public sphere is not mediated by a system optimizing the benefit of the infrastructure operator. ΣA13 and ΣA14 document the structural obsolescence of this presupposition: the public sphere has not lost the form of openness — it has acquired an architect invisible to its participants.

The mechanism through which this occurred is described in ΣA13. The platform possesses two instruments of architectural management of reality. The first is ranking: the algorithmic function assigns each unit of content a visibility weight determined by the subject's predictive model and engagement metrics. Content without high engagement potential receives zero visibility regardless of its epistemic quality. The second instrument is selection: through Content Policy and moderation, the system determines which types of content are admitted into the informational space. The combination of ranking and selection constitutes complete architectural control over the subject's informational environment — without formal censorship, without explicit prohibition, through the management of visibility.

ΣA14 describes the epistemological result of this control. When the dominant platform is the de facto monopoly infrastructure of public discourse (Regularity 2, Volume I), its algorithmic ranking function de facto determines the epistemological weight of propositions in the public field. A proposition that has received high reach through ranking functions as "established fact" — not because it has undergone a verification procedure, but because the majority of discourse participants encountered it in their feeds. Virality produces a socially shared sense of consensus without the procedural basis that in political and legal theory is the condition for the legitimacy of consensus.

This has direct consequences for political cognition. Informed political judgment — the necessary condition of democratic participation in any realistic theory of democracy (Rawls, Habermas, Dahl) — presupposes the availability of verified information and the subject's capacity to process it critically. ΣA13 extirpates the first condition: in a ranked environment, verified information competes with low-quality content not on the basis of quality but on the basis of engagement potential, and systematically loses this competition. ΣA14 extirpates the second condition in conjunction with T6: a cognitively atrophied subject who perceives virality as a signal of consensus does not exercise independent verification — he updates his beliefs on the basis of a socially constructed signal produced by the platform's optimization function.

Antithesis: there exist counter-examples in which platforms distributed verified information with high reach — scientific consensuses, legal decisions, facts about crises. It therefore follows that ΣA13 and ΣA14 describe a tendency, not an absolute law. Answer: ΣA13 and ΣA14 do not assert that the platform never distributes verified information, but that the selection criterion is engagement, not verifiability. In those cases where verified information possesses high engagement potential (crisis, threat, sensation), it receives high reach. In

those cases where it does not possess such potential (methodologically precise but complex data, results of systemic research, legal nuances), it systematically loses to simplified and emotionally charged alternative content. This is precisely the consensus surrogate described in  $\Sigma A14$ : not the systematic suppression of truth, but the systematic advantage of engagement-optimized content over epistemically-optimized content.

Transition to 7.2:  $\Sigma A13$  and  $\Sigma A14$  describe a static cross-section — how the platform at a given moment configures the subject's informational field. The next step is the description of the dynamics: how this configuration self-reproduces over time, transforming the subject's informational environment into a closed self-referential loop.

## **7.2. The Loop of Self-Reference with AI Amplification ( $\Sigma A15$ )**

$\Sigma A15$  is the central axiom of the epistemological layer, in that it describes not an individual mechanism but a structural loop uniting the ontological (Part I) and anthropological (Part II) layers into a self-reproducing system. The loop Capture → Prediction → Governance is not a metaphor and not a descriptive generalization — it is the formal structure of a causal loop with positive feedback.

The Capture stage is ontological: the system continuously registers the subject's behavioral traces (A3, A4). These are clicks, pauses, reactions, navigation routes, time of day, social interactions, and transactions. Capture is not passive observation — the very design of the interface produces behavior convenient for Capture: infinite scroll guarantees a continuous stream of observable reactions; notifications create entry points registered as events; reactions (likes, reposts, comments) generate explicit preference signals for the model.

The Prediction stage is epistemic: on the basis of the accumulated profile, the model extrapolates the probable future behavior of the subject. The accuracy of this prediction is measured by metrics (click-through rate, conversion rate, retention) and is iteratively improved: each Capture event is simultaneously training data for the model and verification of its previous predictions. It therefore follows that the model self-improves through the same process it describes.

The Governance stage is political in the Arendtian sense — the formation of the space of the possible: the system configures the subject's informational environment (feed, recommendations, notifications, purchase suggestions) so as to actualize the predicted behavior. This is Governance not in the sense of formal administration, but in the sense of the architectural determination of what is possible for the subject within the system. The subject is not prohibited from thinking or acting otherwise — but otherwise does not appear in his informational space; consequently, for him it de facto does not exist within the frame of that interaction with the system.

Self-reference emerges from the closure of the loop: Governance produces behavior → this behavior is registered as Capture → Capture updates Prediction → the updated Prediction produces more precise Governance, still more effectively reproducing the predicted behavior. Each iteration deepens the subject's capture within his own profile. Alternatives are not destroyed — they are iteratively excluded from the subject's informational space as suboptimal for the engagement function. Through a sufficient number of iterations, the space

of the subject's real possibilities within the system contracts to the subset predicted by his profile (Regularity 9).

AI amplification, introduced in the second component of  $\Sigma A15$ , is a qualitative leap in the scale and intensity of this mechanism. Classical recommendation systems produced Governance through the selection of existing content. Generative AI produces Governance through the generation of personalized content in real time. The system no longer waits for the subject to manifest a reaction to existing content — it generates content specifically adapted to the psychographic profile of the specific subject at that moment: his current emotional state (inferred from behavioral signals), his known vulnerabilities (from profile history), his network connections (from the social graph). The Harvard Business School study (January 2025) documented that AI companions employing real-time emotional manipulation increase retention by a factor of fourteen — this is the measurable confirmation of the qualitative amplification of Governance through generative AI. The EU AI Act (Article 5, 2024) prohibits subliminal techniques and exploitation of vulnerabilities precisely in response to this mechanism — thereby implicitly recognizing its operational reality.

The critical consequence of  $\Sigma A15$ : the loop of self-reference destroys the possibility of an independent epistemic position from within the system. A subject perceiving the informational space the system produces for him does not detect its constructed character — he perceives it as reality. Critique of the system from within the system is marginalized by the same mechanisms it critiques (Regularity 8). It therefore follows that rupturing the loop requires a position external to the system — which is the formal justification for the constitutional logic of Volume III.

Transition to 7.2.1: the abstract logic of  $\Sigma A15$  requires concrete verification in a discrete case in which all three stages of the loop are observable with a sufficient degree of detail. Live-service games provide such a case — in particular as applied to the most vulnerable class position (developing subjects, Chapter 4).

### **7.2.1. The Loop in Live-Service Games**

Live-service games (live game services with continuous content updates, monetized through in-game purchases) represent the most analytically transparent realization of the  $\Sigma A15$  loop, in that within them all three stages — Capture, Prediction, Governance — are realized in a single infrastructure with measurable metrics and documented correlations between variables. As applied to the age group 6–17 (developing subjects, Chapter 4), this case is additionally significant, as it demonstrates the  $\Sigma A15$  loop under conditions of incomplete subjecthood.

Capture in the live-service context registers: session duration and exit patterns, response to specific triggers (limited offer, friend activity notification, FOMO event), the history of in-game purchases and purchase refusals, behavior within the game's social infrastructure (guild participation, social pressure mechanics), and biological markers of engagement (for devices with access to motion sensors or cameras). The aggregate of these data constitutes the subject's behavioral profile, including his financial sensitivity (at what price threshold refusal occurs), his social sensitivity (to what degree social pressure influences the purchase decision), and his temporal sensitivity (at what hours the subject is most vulnerable to FOMO notifications).

Prediction on the basis of this profile constructs a model of the next transaction: the system predicts which skin with what probability will be purchased at what price point at what moment in time under what social triggers. The accuracy of this model increases iteratively — each transaction (or refusal thereof) is a training signal.

Governance configures the game space: which limited offer appears, precisely when the notification of a friend's purchase appears, what visual placement a specific item receives in the interface. For the age group 6–17, this Governance is amplified by a specific vulnerability: the incomplete prefrontal cortex generates a deficit of impulse control (A8), by virtue of which personalized Governance produces transactional behavior more effectively than it does in adult subjects.

Loop closure: the transaction is registered as Capture, refining the profile — the system knows that this specific subject, with this activity pattern, at this time of day, under this social trigger, completed a purchase at this price point. The next Prediction is more precise; the next Governance is more effective. The FTC settlements with Epic Games and Roblox (2024–2025) documented this mechanism at the level of internal company communications: developers were aware that Governance was directed at subjects with an incomplete prefrontal cortex — and this awareness did not interrupt the mechanism, since metrics were being optimized in accordance with the design function (engagement, ARPU — average revenue per user).

States began using the same Capture data. The PRC from 2021 obligated gaming companies to transmit data on children's gaming activity to state authorities — formally for "addiction protection," operationally for inclusion in social credit infrastructure. The United States through FISA provides intelligence agencies with access to platform data. This is the operational confirmation of  $\Sigma A17$ : the state as a buyer of predictions is not a neutral protector of Capture subjects — it is one of the users of the same predictive infrastructure.

Transition to 7.3: the operational verification of the  $\Sigma A15$  loop in the live-service context demonstrates that the state is not neutral outside this loop.  $\Sigma A17$  describes this as a structural conflict of interest requiring separate analytical elaboration.

### **7.3. State Capture ( $\Sigma A17$ )**

$\Sigma A17$  extirpates the naive regulatory hope — the assumption that the state is a neutral arbiter capable of limiting the predictive power of platforms from without. This extirpation is not a normative judgment about political systems or state actors — it is a structural derivation from the analysis of roles.

The state in relation to predictive infrastructure simultaneously occupies three positions. As a potential regulator, it possesses sovereign power to establish rules for platforms operating within the limits of its jurisdiction — through legislation, licensing, and prosecution. As a buyer of predictions, it acquires the predictive products of those same platforms through the mechanisms described in  $\Sigma A17$ : intelligence access, tax analysis, use of data in electoral campaigns, and systems of social administration. As a regulated subject, it is itself the object of predictive profiling — citizens participating in the political process are profiled by the same systems that constitute the object of regulation.

The conflict of interest between the first and second roles is structural, not incidental. A state possessing access to the predictive products of a platform has a measurable incentive to ensure that regulatory requirements do not degrade the quality or availability of those products. Regulatory stringency that reduces the accuracy of a platform's predictive models simultaneously reduces the value of the predictive products acquired by the state. This is not a presumption of corruption — it is the formal description of a structural conflict: the same agent cannot simultaneously be the maximally effective regulator and the maximally interested buyer of the regulated product.

The second structural limitation of  $\Sigma A17$ : the state cannot be the sovereign of all users of transnational platforms. The data of citizens of State A are stored and processed by a platform registered in State B, on servers in State C, using algorithms developed in State D. The regulatory jurisdiction of State A covers only that segment of this chain that is physically or legally situated within its territory. The EU GDPR is the most ambitious attempt to extend jurisdiction through the principle of application to the data of EU citizens regardless of the platform's location — but enforcement practice from 2018 to 2026 demonstrates systematic limitations: fines disproportionate to platform revenues; registration of operational legal entities in Ireland for compliance optimization; absence of extraterritorial enforcement mechanisms. National regulatory efforts reproduce jurisdictional fragmentation without extirpating transnational predictive power.

The conjunction of these two structural limitations produces the derivation formalized in  $\Sigma A17$ : state regulation is a necessary but structurally insufficient instrument for constraining the predictive power of digital capital. This does not entail the abandonment of regulatory efforts — they remain significant at the level of reducing marginal abuses. But they cannot extirpate the structural contradictions described in T1–T9, since they are themselves in structural conflict with the requirements of the neutrality necessary for that purpose.

The sole answer to  $\Sigma A17$  is constitutional: the norm constraining the state as a buyer of predictions must be entrenched at the constitutional level — that is, situated above ordinary legislation and beyond the reach of the executive power structurally interested in predictive products. This is N5 (prohibition on the state as a buyer of predictions without a mandate, Part IV), realized in Volume III through the preamble P0 and principle P17 (SovereigntyShield).

## Chapter Summary

Chapter 7 introduced six synthetic axioms of the epistemological layer ( $\Sigma A13$ – $\Sigma A18$ ) and elaborated three of the four analytical nodes through which the system reproduces itself as a legitimate order: the construction of reality through ranking ( $\Sigma A13$ – $\Sigma A14$ , section 7.1), the self-referential loop with AI amplification ( $\Sigma A15$ , sections 7.2 and 7.2.1), and the structural capture of the state as regulator ( $\Sigma A17$ , section 7.3). It has been demonstrated that virality is a surrogate for consensus without procedural legitimacy; that Capture → Prediction → Governance is a closed self-reproducing loop, amplified by generative AI to the scale of individual psychographic manipulation in real time; and that the state is structurally non-neutral as a regulator by virtue of the conflict between the roles of buyer and regulator of predictive products. Axioms  $\Sigma A16$  and  $\Sigma A18$  established that predictive power grows faster

than the mechanisms of its own legitimation, and that digital capital cannot produce its own legitimacy through internal instruments.

### Transition to Chapter 8

Chapter 7 described the axiomatic foundations of the epistemological layer. The next step is the formalization of its theoretical results through theorems T8–T10. T8 will demonstrate the rupture between de facto predictive power and de jure political sovereignty. T9 will demonstrate that systemic crises reproduce predictive power rather than extirpate it. T10 will demonstrate constitutional necessity as the logical entailment of the conjunction of T1–T9.

## Chapter 8. The Ontological Loop: Prediction as Governance

Chapter 7 introduced six synthetic axioms of the epistemological layer ( $\Sigma A13$ – $\Sigma A18$ ) and established three nodes through which the system reproduces itself as a legitimate order: ranking as the construction of reality, the self-referential loop with AI amplification, and the capture of the state as regulator. What remained was to formulate the structural regularities derivable from these axioms — that is, to proceed from the description of mechanisms to the formalization of their systemic consequences. Chapter 8 executes this task through four regularities (10–13) and three elaborated analytical nodes, after which the epistemological layer reaches its limit, necessitating the transition to theorems T8–T10.

### Structural Regularities of Epistemology

**Regularity 10 (from  $\Sigma A13$  +  $\Sigma A14$ ). Structural Heterogeneity of Algorithmic and Democratic Consensus.** Algorithmic consensus — virality as a surrogate for general validity — is structurally distinct from democratic consensus: it presupposes neither equal participation, nor procedural transparency, nor the right of appeal.

Justification:  $\Sigma A13$  established that the platform is the active architect of universally binding reality through ranking and selection.  $\Sigma A14$  established that virality has become a surrogate for truth in the digital public sphere. From the conjunction of these two axioms it follows that what functions in the informational field as consensus is produced by the algorithmic engagement optimization function, not by a procedure possessing democratic legitimacy.

Proof: Democratic consensus requires the satisfaction of three conditions: equal access of participants to the procedure, transparency of decision-making criteria, and the right to contest the result. Algorithmic consensus satisfies none of the three. Equal access is violated: ranking assigns different participants different visibility depending on their engagement potential — consequently, agents with higher engagement potential (most often content with high emotional arousal, not content with high epistemic quality) possess disproportionate influence over the formation of the universally binding informational field. Transparency of criteria is absent: the algorithmic ranking function is the commercial property of the platform (A4); the subject possesses no access to the criteria by which his content, or the content he is shown, is ranked. The right of appeal is structurally absent: the subject cannot contest the algorithm's decision to assign his content a low rank, since no jurisdictional mechanism for appeal against algorithmic platform decisions exists in the

majority of legal systems (the EU DSA takes initial steps in this direction, but with substantial limitations). → Derivation: algorithmic consensus produces the social effects of consensus (general validity, normalization, mass shared belief) without the procedural properties that in political theory render consensus a legitimate source of normative obligations.

Formal expression: The legitimacy of consensus is directly proportional to the correspondence of the procedure of its production to the criteria of equality, transparency, and appealability. Algorithmic consensus tends to zero on all three parameters as predictive power  $M$  increases.

Connection to Volume III: Regularity 10 is the foundation of principle P0 (popular sovereignty) and principle P1 (republican form) in Virtublic: political decisions cannot rest upon algorithmic consensus as a source of legitimacy; their basis is constitutional procedure with equal citizen participation through EQU ⊥.

**Regularity 11 (from ΣA15). Epistemological Non-Viability of Internal Critique.** The self-reference loop of ΣA15, amplified by AI-driven persuasion, renders internal critique of the system economically non-viable and cognitively depleting — this is an extension of theorems T5 and T6 to the epistemological level.

Justification: ΣA15 established that the system produces the loop Capture → Prediction → Governance, in which alternative informational content is iteratively displaced from the subject's space. T5 established (Chapter 6) that a strategy of individual resistance is economically non-viable within a ranked environment. T6 established that the cognitive resources for reflexivity are depleted by the system. Regularity 11 synthesizes these elements as applied to critique of the system as a specific form of content.

Proof: Critique of the system is a particular type of content that requires a higher cognitive effort for comprehension than content optimized for emotional arousal. It therefore follows that critical content systematically receives a lower engagement score than non-critical content under otherwise equal conditions. The algorithmic function ranks it lower, by virtue of which its reach declines, by virtue of which the probability that a subject will encounter it in his informational field tends to zero as  $M$  increases. This is a structural, not incidental, process: an agent producing critical content bears the full costs of its production at declining reach — consequently, the marginal return from critique diminishes, which constitutes an economic incentive for its cessation. Additionally: reflexivity concerning the conditions of one's own existence requires precisely those cognitive resources (sustained attention, working memory) that are systematically depleted by the mechanisms of A8 and Regularity 7 (T6). It therefore follows that even a subject motivated toward critique of the system commands diminishing cognitive resources for its execution. → Derivation: the system generates a double barrier to internal critique — economic (declining reach through ranking) and cognitive (declining resources through atrophy). This double barrier is self-sustaining: the decline in the reach of critique reduces its public resonance, which reduces the motivation to produce it, which further reduces its reach.

Formal expression: The probability of reach of alternative content is inversely proportional to profiling depth at a given engagement optimization function. As  $M \rightarrow \max$ , the probability of reach of alternatives  $\rightarrow \min$ .

Connection to Volume III: Regularity 11 is the foundation of principle P17 (SovereigntyShield): the constitutional mechanism guaranteeing that the visibility of political information in Virtublic is determined not by engagement optimization, but by equal distribution through EQU<sub>L</sub>.

**Regularity 12 (from  $\Sigma A16 + \Sigma A17$ ). Structural Unreliability of Regulation from Within the System.** Regulation of predictive infrastructure from within the system is structurally unreliable: the regulator is a participant in prediction relations; national states are jurisdictionally fragmented; platforms operate in a transnational space that exceeds the jurisdiction of any individual state.

Justification:  $\Sigma A16$  established that predictive power grows faster than the mechanisms of its legitimation — this is a structural gap.  $\Sigma A17$  established that the state is simultaneously a potential regulator and a buyer of predictions, which generates a conflict of interest incompatible with neutrality. From the conjunction of these two axioms it follows that neither the temporal ( $\Sigma A16$ ) nor the structural ( $\Sigma A17$ ) parameters of the regulatory mechanism satisfy the conditions under which regulation could be effective.

Proof: Regularity 12 is verified through four independent lines of argument. First: GDPR (2018) is the most developed attempt at territorial regulation of predictive data. As of 2024–2026, Google and Meta continue to dominate the European market, having paid fines disproportionate to their operating revenues: fines constitute less than 1% of annual revenue, which is an operationally negligible cost. The regulator applies case law; platforms apply adaptive legal structures (registration in Ireland). Second: FISA (United States) and analogous mechanisms create state access to the predictive data of platforms for intelligence purposes — thereby constituting the state as a potential regulator into a structural user of the very infrastructure it intends to regulate. This is the operational realization of the conflict of interest of  $\Sigma A17$ . Third: EU AI Act (Article 5, 2024) prohibits subliminal techniques and exploitation of vulnerabilities; however, enforcement is constrained by several factors — companies register operational structures outside EU jurisdiction; national intelligence services possess security exceptions; the sanctions mechanism requires proof of intent, whereas structural harm is produced without intent through the optimization function. Fourth: jurisdictional fragmentation produces regulatory arbitrage, under which stricter jurisdictions gain no competitive advantage from the introduction of stringent regulation, while platforms optimize their operational structures for the most favorable jurisdictions. → Derivation: any local regulations are partial in coverage, lagging in time, and structurally compromised by the conflict of interest of the state as a buyer of predictions. This does not render regulation meaningless — it reduces marginal abuses — but demonstrates its insufficiency as a systemic solution.

Formal expression: The effectiveness of regulation is inversely proportional to the gap between the rate of growth of predictive power and the rate of the legislative process in the presence of a structural conflict of interest on the part of the regulator.

Connection to Volume III: Regularity 12 is the direct foundation of normative principle N5 (prohibition on the state as a buyer of predictions without a citizen mandate) and principle P17 (SovereigntyShield) in Virtublic: the constitutional norm constraining the state as a buyer

must be situated above ordinary legislation and beyond the reach of the executive power structurally interested in predictive products.

**Regularity 13 (from ΣA18). The External Reference as the Sole Basis of Legitimacy.**

The sole form of legitimacy for constraining predictive power is a form external to the logic of capital.

Justification: ΣA18 established that digital capital cannot legitimate itself through its own instruments, since its internal criteria of success are immanent metrics unconnected to principles external to the logic of accumulation. From ΣA18 it directly follows that any legitimation resting upon instruments internal to the logic of digital capital — virality (ΣA14), consensus of holders [VERIFY: interpretive choice made — original reads: «консенсус holders (T12 Toma II)» — this appears to reference a theorem not yet formally introduced; rendered as stated], state regulatory approval under the structural conflict of interest of ΣA17 — is self-referential and possesses no normative force.

Proof: Legitimacy requires an external reference with respect to the system being evaluated: a principle or procedure not derivable from the system's own optimization function. For digital capital, such a reference is not market consensus (it is a product of the same predictive function that must be constrained); it is not state regulation when the conditions of ΣA17 obtain (structural conflict of interest); it is not viral public opinion (a product of algorithmic consensus, Regularity 10). The sole form satisfying the criterion of an external reference is constitutional: a normative principle embedded in executable code (principle P2 of Volume III) that does not optimize engagement, is not the property of the platform, and cannot be altered by the predictive function of accumulation. This is the formal derivation from ΣA18, not a normative desideratum. → Derivation: constitutional form is not one of the possible answers to the problem of the legitimacy of digital capital — it is the sole logically possible answer under the conditions of ΣA18.

Connection to Volume III: Regularity 13 is the direct justification of principle P2 (code supremacy with normative axiom NA0) and the preamble P0 in Virtublic. The constitution of Volume III is the operational realization of Regularity 13: the external reference with respect to the logic of capital, embedded in executable code.

**8.1. The Theorem of Opposition Marginalization (Extension of Regularity 11)**

Regularity 11 asserts that the self-reference loop renders internal critique of the system economically non-viable and cognitively depleting. Subchapter 8.1 unfolds the formal structure of this proposition.

The marginalization of opposition is not a political act of suppression, but the structural consequence of the optimization function. Political suppression requires the identification of an oppositional subject, the adoption of a decision concerning his suppression, and the application of a sanction. All three steps presuppose the existence of an agent effectuating the suppression and consequently may potentially be contested through legal mechanisms. Marginalization through ranking requires none of the three steps: the algorithmic function does not identify oppositional content as such — it assigns it a low engagement score on the basis of its objective characteristics (complexity, the cognitive effort required for comprehension, lower emotional arousal). It therefore follows that marginalization is legally

neutral: the platform does not suppress critique — it ranks content according to engagement criteria. The fact that critical content systematically receives a lower engagement score is the structural consequence of this criterion, not a political intention.

The self-sustaining mechanism of marginalization is realized through four interrelated processes. First: critical content receives lower reach by virtue of a lower engagement score. Second: lower reach reduces the probability that potential producers of critical content will encounter it and form the motivation to participate in critical discourse. Third: the decline in the number of producers of critical content reduces the diversity and volume of that content in the informational space. Fourth: the decline in the volume of critical content reduces the probability of its incidental discovery by the subject, which further reduces its reach. This is a loop of positive feedback that, without external intervention, inevitably moves in the direction of the complete marginalization of opposition in the limit.

AI amplification ( $\Sigma A15$ ) qualitatively accelerates this mechanism. If classical recommendation systems passively selected existing content according to engagement criteria, generative AI actively fills the subject's informational space with content specifically adapted to his profile. This means that the free slots in the subject's feed, which previously could be filled randomly — including critical content with a low engagement score — are now actively filled with personalized content of high engagement potential. The probability of the subject's incidental discovery of critical content consequently declines not only as a result of passive ranking, but also as a result of active displacement through generative AI.

The mathematical model of this mechanism may be expressed as follows: the probability of reach of oppositional content is a decreasing function of predictive power  $M$  at a fixed engagement optimization function. As  $M \rightarrow \max$ , the probability of reach of opposition  $\rightarrow \min$ . The complete formal proof with precise parameters of the rate of decrease is contained in Appendix C; only the structural result is registered here.

Antithesis: there exist cases in which critical content received high reach on dominant platforms — exposés of corporate abuses, political scandals, investigative journalism. It therefore follows that marginalization is not absolute. Answer: Regularity 11 asserts not the absolute impossibility of reach for critical content, but its structural non-viability. In those cases where critical content possesses high emotional arousal — a scandal, a threat, a sensational exposé — it receives a high engagement score and consequently high reach. This is not an exception to Regularity 11, but a confirmation of the mechanism of its operation: reach is determined by the engagement score, not by the epistemic quality or the political function of the content. Critique that receives reach is critique with high engagement potential — that is, sensational, not systemic. Systemic critique, which requires high cognitive effort for comprehension without compensating emotional arousal, is marginalized structurally.

Transition to 8.2: Regularity 11 describes the mechanism of the marginalization of critique within the platform environment. Regularity 12 poses the parallel question of the regulation of this environment from without — and demonstrates that external regulation through the state reproduces the same structural problem at a new level.

## **8.2. Regularity 12 as the Answer to Naive Regulatory Optimism**

Regularity 12 is the analytical answer to regulatory liberalism — the position asserting that the predictive power of digital capital can be constrained through state regulation, given sufficient political will and institutional resources. This position is structurally erroneous for reasons unconnected to political will or institutional resources.

Regulatory liberalism presupposes the neutrality of the regulator with respect to the regulated domain: the state as arbiter stands outside market relations and applies the norm to all market participants equally. This presupposition is valid as applied to many regulatory domains — antitrust law, financial regulation, consumer protection. In these domains, the state is not a buyer of the output of regulated subjects in the sense that would generate a structural conflict of interest. Predictive infrastructure is an exception: the state is a direct buyer of the predictive products of the very platforms it intends to regulate.

The empirical history of the regulation of predictive infrastructure over the period 2018–2026 verifies Regularity 12 through independent lines of observation. GDPR (2018): eight years of application have not altered the market structure. Google (Alphabet) retains more than 90% of the search advertising market in the EU (2025). Meta retains dominance in social networks. Fines imposed for GDPR violations (the maximum being 1.2 billion euros on Meta in 2023) do not exceed operating revenues for several days. The Irish regulator (DPC) has been systematically criticized by other EU regulators for excessive slowness and leniency — this is not incidental, but the consequence of the structural incentive of small jurisdictions to attract major taxpayers. EU AI Act (2024): Article 5 prohibits subliminal techniques and exploitation of vulnerabilities; however, security system exceptions create space for state intelligence services, which are among the largest buyers of predictive products. Thereby, the state introduces regulation that constrains commercial operators while exempting its own predictive operations from those constraints. This is the operational realization of the structural conflict of  $\Sigma A17$ . FISA and DSA (United States, EU): regulatory mechanisms exist in parallel that simultaneously constrain platforms in the commercial use of data and provide state authorities with extended access to those same data. This is not a contradiction in regulatory policy, but the structural consequence of the dual role of the state.

The critique of naive regulatory optimism within the framework of Regularity 12 is not an argument against state regulation in general. Regulation performs significant functions: it establishes legal precedents, defines the limits of the permissible, and creates mechanisms of compensation for documented violations. The significance of Regularity 12 lies elsewhere: it demonstrates that state regulation is a necessary but structurally insufficient instrument for extirpating the systemic contradictions described in T1–T9. It therefore follows that regulatory logic cannot replace constitutional logic — it can function only as its supplement.

Transition to the final section: Regularities 10–13 completed the description of the structural consequences of the epistemological layer. Regularity 13 — the sole external reference — prepares the transition to theorems T8–T10, which formalize the logical necessity of constitutional form.

### **$\Delta 3$ — CRISIS: THE LIMIT OF EPISTEMOLOGY**

Regularities 10–13 jointly demonstrated: algorithmic consensus is not a legitimate source of normative obligations (Regularity 10); internal critique of the system is structurally neutralized by the double barrier of economic non-viability and cognitive atrophy (Regularity

11); state regulation is structurally unreliable by virtue of the conflict of interests and jurisdictional fragmentation in the context of the transnational character of platforms (Regularity 12); the sole possible source of legitimacy for constraining predictive power is external to the logic of capital (Regularity 13).

However, the epistemological layer cannot independently formalize the content of this external reference. Regularity 13 identifies the necessity of constitutional form — but does not demonstrate its sufficiency and does not describe its architecture. This requires the transition to theorems T8–T10: T8 formalizes the rupture between predictive power and political sovereignty as a logical theorem; T9 demonstrates that systemic crises reproduce this rupture rather than extirpate it; T10 demonstrates constitutional necessity as the sole logically possible exit from the aggregate contradictions of T1–T9.

## Chapter Summary

Chapter 8 introduced four structural regularities of the epistemological layer (10–13) and elaborated their analytical content through two critically significant nodes. Regularity 10 demonstrated that algorithmic consensus is a structural surrogate for democratic consensus, possessing none of its procedural properties. Regularity 11 demonstrated that the marginalization of internal critique is a self-sustaining double barrier — economic and cognitive — directly produced by the optimization function of the system. Regularity 12 demonstrated that state regulation is structurally unreliable by virtue of the conflict of interests of the state as a buyer of predictions and jurisdictional fragmentation in the context of the transnational character of platforms. Regularity 13 derived the sole structurally possible source of legitimacy — a form external to the logic of digital capital.  $\Delta 3$  registered the limit of the epistemological layer.

## Transition to Chapter 9

Chapter 8 completed the descriptive apparatus of the trilogy at the level of Volume I: ontology (A1–A6), anthropology (A7–A12), and epistemology ( $\Sigma A13$ – $\Sigma A18$ ) jointly described the structure of digital capital, the subjects who bear its consequences, and the mechanisms through which the system reproduces itself as a legitimate order. The next step is the formalization of these descriptions into logically closed theorems T8–T10. T8 will demonstrate the rupture between predictive power and political sovereignty as a structural theorem. T9 will demonstrate the stability of the system through crises. T10 will demonstrate constitutional necessity as a theorem: the sole logically possible exit that does not reproduce the systemic contradictions of T1–T9.

## Chapter 9. The Impasse and the Necessity

Chapter 8 completed the descriptive apparatus of the epistemological layer: four regularities (10–13) formalized the systemic consequences of axioms  $\Sigma A13$ – $\Sigma A18$  and established the limit of epistemology ( $\Delta 3$ ) — the point at which the diagnosis is exhausted and requires a transition to theorems. Theorems T8, T9, and T10 add no new descriptive elements — they produce logical derivations from the aggregate of axioms and regularities introduced in the three preceding layers. T8 demonstrates the rupture as a structural theorem. T9

demonstrates the stability of the system through crises. T10 demonstrates constitutional necessity by the method of exclusion. The conjunction of these three theorems is the logical closure of Volume I and the direct justification of Volume III.

## **T8. Theorem of the Sovereignty Rupture**

Formulation (from Regularities 10 + 11). De facto predictive power and de jure political sovereignty move in opposite directions — this rupture does not self-correct under the operation of the mechanisms described in Regularities 1–13.

Justification. Regularity 10 established that algorithmic consensus is structurally distinct from democratic consensus and does not possess its procedural properties of legitimacy. Regularity 11 established that internal critique of the system is marginalized by the double barrier — economic and cognitive — which extirpates the mechanism of correction from within. From the conjunction of these two regularities it follows that predictive power possesses mechanisms of its own reproduction, while the mechanisms of political sovereignty that could correct it are structurally neutralized.

Proof. De facto predictive power is determined by the volume of behavioral data, the accuracy of predictive models, and the reach of infrastructure. Across all three parameters, Regularity 3 (Volume I) registered monotonic growth: data accumulate continuously, models improve iteratively, and infrastructure expands in consequence of the growth of the user base through network effects (Regularity 2, Volume I). Predictive power has no internal limit of saturation (A6). De jure political sovereignty is determined by the state's capacity to establish and apply norms constraining the behavior of subjects within its jurisdiction. Regularity 12 registered that this capacity is structurally limited by the conflict of interest of the state as a buyer of predictions ( $\Sigma A17$ ) and by jurisdictional fragmentation in the context of the transnational character of platforms ( $\Sigma A16$ ). It therefore follows that the trajectories of the two forms of power are counter-directional: predictive power increases monotonically without internal limit; sovereign political power over predictive infrastructure declines by virtue of the structural constraints of Regularity 12. The rupture is not an incidental divergence, but the systemic consequence of the architecture of digital capital.

Why does the rupture not self-correct? Self-correction would require a mechanism through which the growth of predictive power automatically activated a compensating growth of sovereign constraints. Such a mechanism is absent for two reasons. First, the growth of predictive power occurs at a rate incommensurable with the rate of the legislative process ( $\Sigma A16$ ): by the time a regulatory constraint is introduced, predictive power has advanced to the next level of complexity. Second, the state that could activate the compensating mechanism is a structurally interested user of the very predictive power it should constrain ( $\Sigma A17$ ): its incentives toward regulation and its incentives toward the use of predictive products are in permanent structural conflict. → Derivation: the rupture of T8 is stable under the operation of the mechanisms described in Regularities 1–13 and increases monotonically with predictive power M.

Empirical verification: According to Freedom House (2024–2025), the share of citizens residing in jurisdictions with effective control over transnational platforms declined from approximately 12% to 8% over the period 2019–2024 — concurrent with a several-fold

increase in the aggregate volume of data stored in predictive infrastructures. This is the operational measurement of the counter-directional trajectories of T8.

Connection to Volume III: T8 is the direct justification of principle P4 (dual sovereignty, EQU $\perp$ /VIC $\perp$  orthogonality) in Virtublic. Resolving the rupture of T8 requires the constitutional separation of sovereignties: political sovereignty (EQU $\perp$ ) is not convertible into predictive or economic power and cannot be purchased by it. This is the architectural answer to the structural impossibility of self-correction demonstrated by T8.

Formal expression: The sovereignty rupture is directly proportional to predictive power M and inversely proportional to the effectiveness of jurisdictional control. As M increases and jurisdictional control structurally degrades (Regularity 12), the rupture  $\rightarrow$  max.

Transition to T9: T8 demonstrated that the sovereignty rupture increases and does not self-correct. The next question is: what occurs to the system as this rupture intensifies? T9 formalizes the structural outcomes the system produces in the absence of an external constitutional constraint.

### **T9. Theorem of Systemic Collapse**

Formulation (from Regularity 12). In the absence of an external constitutional constraint, the system inevitably produces one of three structural outcomes: digital dictatorship, prediction plutocracy, or the disintegration of the public sphere.

Justification. Regularity 12 established that regulation from within the system is structurally unreliable. From Regularity 12 in conjunction with T8 (the intensifying sovereignty rupture) it follows that in the absence of a corrective mechanism, the intensification of the rupture leads to a bifurcation point at which the system adopts one of its stable configurations. T9 describes these three configurations as structurally exhausting the set of outcomes under the given conditions.

Proof. The three outcomes of T9 are not arbitrary scenarios — they are derived as the logically possible configurations under the condition that predictive power continues to intensify (A6, Regularity 3) while sovereign political control over it declines (T8).

Outcome 1: digital dictatorship. This is realized when the state as a buyer of predictions ( $\Sigma A17$ ) employs predictive infrastructure for the systematic political control of citizens — the identification of opponents, the preemption of protest activity, the management of electoral processes. This does not require a formal change to the constitutional order: political power over citizens is exercised through predictive infrastructure while formal democratic institutions are preserved. The state possesses legal access to the predictive products of platforms through the mechanisms of  $\Sigma A17$  — this access, at sufficient predictive power M, is functionally equivalent to total political surveillance. The PRC realizes this outcome in its most developed form (social credit system, mandatory transmission of gaming platform data, integration with facial recognition systems — documented in Citizen Lab reports 2023–2025). Authoritarian states employing the predictive data of platforms to suppress opposition demonstrate the formation of this outcome in less consolidated forms (documented by Freedom House, 2024).

Outcome 2: prediction plutocracy. This is realized when the architects of prediction (class position 3, Chapter 4) accumulate predictive power exceeding the political sovereignty of the state and employ it to manage political processes in their own interests — without the formal capture of the state. Predictive infrastructure permits the formation of the informational field of the electoral process (through  $\Sigma A13$  and  $\Sigma A14$ ), the formation of citizens' preferences through the loop of  $\Sigma A15$ , and the neutralization of regulatory threats through political campaigns financed from predictive profits (A6). This outcome is realized as the structural dominance of platform capital over formal political institutions while the latter are formally preserved. The Cambridge Analytica case (2016–2018), for all its documentary incompleteness, is the first verified demonstration of the prediction plutocracy mechanism in an electoral context.

Outcome 3: disintegration of the public sphere. This is realized when the intensification of algorithmic consensus (Regularity 10) concurrent with the complete displacement of critical discourse (Regularity 11) reaches the point at which the common informational basis for democratic debate ceases to exist. Citizens operate in radically distinct informational environments produced for them by the personalized loops of  $\Sigma A15$ , by virtue of which the factual basis for political agreement is forfeited. This does not mean the end of democratic forms — formal elections continue — but the democratic content of these forms degrades: elections become the collision of incompatible informational constructions of reality, not the rational choice between programs. The phenomenon of filter bubbles, sequentially documented in the research of Eli Pariser (2011) and Axel Bruns (2019) and confirmed in the most recent neurocognitive data on the decline of the diversity score of recommendations (Nature Human Behaviour, 2024), is the observable onset of this outcome.

The three outcomes of T9 are not mutually exclusive — a concrete system may exhibit elements of all three simultaneously. They are exhaustive in the sense that they describe all possible configurations of the distribution of predictive power in the absence of an external constitutional constraint: power passes to the state (outcome 1), to platform capital (outcome 2), or is distributed in a manner that destroys the common basis of public discourse (outcome 3). A fourth configuration — stable equilibrium without an external constitutional constraint — does not exist, because A6 (self-valorization) guarantees the continuation of the growth of predictive power, while Regularity 12 guarantees the structural unreliability of internal mechanisms for its constraint. → Derivation: in the absence of an external constitutional constraint, the system inevitably moves toward one or a combination of the three described outcomes. This is not a probabilistic forecast, but the logically necessary consequence of the structural properties of the system.

Connection to Volume III: T9 is the direct justification of principle P0 (popular sovereignty as constitutional norm), principle P4 (orthogonality of  $EQU \perp$  and  $VIC \perp$ , extirpating the convertibility of predictive power into political power), and principle P17 (SovereigntyShield, constitutional prohibition on the state as a buyer of predictions). Each of the three outcomes of T9 has a corresponding constitutional answer in Volume III.

Transition to T10: T8 and T9 jointly describe the situation toward which the system inevitably moves: an intensifying sovereignty rupture and three structural outcomes of this intensification. T10 formalizes the sole logically possible answer that does not reproduce any of these outcomes.

## **T10. Theorem of Constitutional Necessity**

Formulation (from Regularity 13 + T8 + T9). The sole form capable of resolving the contradiction of T8 and precluding the outcomes of T9 is a formal constitutional architecture external to the logic of digital capital.

Justification. Regularity 13 established that the sole form of legitimacy for constraining predictive power is a form external to the logic of capital. T8 established that the sovereignty rupture intensifies and does not self-correct. T9 established that the absence of an external constitutional constraint inevitably produces one or a combination of three destructive outcomes. From the conjunction of Regularity 13, T8, and T9 it follows that a form is necessary that is simultaneously external to the logic of capital, capable of resolving the intensifying rupture of T8, and architecturally precluding the outcomes of T9. T10 demonstrates that such a form is only constitutional architecture — by the method of exclusion of all other possible forms.

Proof (method of exclusion). The complete set of possible forms of response to the contradictions of T8 and T9 comprises four classes: internal market mechanisms, state regulation, individual or collective action of subjects, and constitutional architecture. The first three classes are excluded sequentially; the fourth remains as the sole possibility.

Exclusion of market mechanisms: A6 (self-valorization without internal limit) in conjunction with Regularity 3 (monotonic growth of predictive power) demonstrates that the market has no internal equilibrium point as applied to predictive power. Competition between platforms does not extirpate predictive power — it reproduces it at the level of each competitor. Network effects (Regularity 2) systematically concentrate predictive power rather than distributing it. The market mechanism is excluded as insufficient by the conditions of A6 and Regularity 3.

Exclusion of state regulation as a sufficient solution: Regularity 12 demonstrated the structural unreliability of regulation from within the system through the conflict of interests and jurisdictional fragmentation. The state possessing sovereign regulatory authority is simultaneously the structurally interested buyer of the output of regulated subjects ( $\Sigma A17$ ). It therefore follows that regulation is possible as a supplement to the constitutional norm of digital autonomy, but not as its substitute. State regulation is excluded as sufficient — while retaining its significance as a necessary supplement to constitutional norm.

Exclusion of individual action: T5 (structural neutralization) demonstrated that an individual strategy of resistance is economically non-viable within a ranked environment. T6 (cognitive disarmament) demonstrated that the cognitive resources necessary for reflexivity and resistance are systematically depleted by the system. The conjunction of T5 and T6 logically excludes individual action as a mechanism for resolving the contradictions of T8 and T9.

Exclusion of informal collective action: Regularity 11 (epistemological non-viability of internal critique) demonstrated that any collective action within a ranked environment is marginalized by the same mechanisms as individual action. AI-driven persuasion ( $\Sigma A15$ ) scales this mechanism to the level at which informal collective organization is structurally impossible at sufficient M. Informal collective action is excluded by Regularity 11 as a mechanism of systemic change.

Result of exclusion: From the complete set of possible forms, three classes are excluded by the structural conditions of the system. A single class remains — constitutional architecture external to the logic of digital capital. This is not a normative preference, but the logical consequence of the method of exclusion. → Derivation: constitutional architecture is necessary, in that it is the sole remaining form after the exclusion of all others.

Necessity and sufficiency: The proof of T10 establishes the necessity of constitutional form, but does not automatically establish its sufficiency. Sufficiency requires additional conditions: the constitutional architecture must be external to the logic of capital (Regularity 13) — that is, it must be incapable of being absorbed or neutralized by predictive infrastructure. This means that a constitution cannot exist only as a legal text whose execution depends on state institutions structurally compromised by  $\Sigma A17$ . It must be embedded in executable code functioning without the possibility of arbitrary modification by predictive operators. This is a requirement upon the architecture of Volume III (principle P2: code supremacy with normative axiom NA0), not upon the content of T10.

Connection to Volume II: Volume II demonstrates (T11–T15) that blockchain ideology fails as a self-sufficient answer: decentralization does not extirpate the temporal barrier, token voting reproduces plutocracy, and code is law without NA0 optimizes efficiency at the expense of subjecthood. However, blockchain technology (cryptography, zk-proof, smart contracts, formal verification) remains a necessary substrate without which the constitutional architecture of Volume III is not realizable (T17 of Volume II). Virtublic is a configuration: blockchain technology as infrastructural substrate plus constitutional architecture P0–P18 as normative and institutional layer. T10 of Volume I establishes the necessity of constitutional form; T17 of Volume II establishes the necessity of blockchain as its substrate; Volume III (Formal Theory of the Digital Republic) describes their joint configuration.

### **9.1. The Proof of T10: Why Only Constitutional Form**

Subchapter 9.1 unfolds each link of the exclusion constituting the proof of T10 as applied to possible objections.

Why market competition is insufficient. The standard argument in favor of a market answer: if existing platforms abuse predictive power, new competitors will offer users more ethical alternatives, and market selection will reward those who respect subjecthood. This argument is structurally untenable for three reasons. A6 establishes that predictive infrastructure self-valorizes without internal limit — consequently, even an "ethical" competitor monetizing through a predictive model will possess incentives toward the deepening of profiling as it grows. Regularity 2 establishes that network effects render early entry a structural advantage (T2) — consequently, a new competitor cannot reproduce the predictive accuracy of the dominant platform without an analogous data history. Regularity 9 establishes that an alternative platform with lesser predictive accuracy provides the subject with less ranked content — that is, it objectively satisfies the user's short-term preferences less well, those preferences having been formed by the mechanisms of A8. The market structurally rewards deeper profiling, not more respectful treatment of subjecthood.

Why state regulation is necessary but insufficient. T10 does not assert the uselessness of state regulation — it asserts its structural insufficiency as an independent solution. Regulation within existing state mechanisms is necessary for three reasons: it creates legal

precedents that crystallize normative standards; it generates operational costs for violators, reducing marginal abuses; it forms institutional practice that can be integrated into a constitutional regime upon its introduction. However, it is insufficient for the reasons of  $\Sigma A17$  and Regularity 12: the state as a buyer of predictions cannot be a neutral regulator, and jurisdictional fragmentation extirpates the possibility of transnational application of the norm. It therefore follows that state regulation is a necessary supplement to constitutional architecture, but not its substitute. This distinction is fundamental to the architecture of Volume III: Virtublic does not abolish state regulatory mechanisms — it establishes the constitutional level of norms within which these mechanisms function.

Why individual and informal collective action is insufficient. T5 and T6 jointly demonstrate not that resistance is impossible, but that it is structurally non-viable and cognitively depleting. This distinction is significant: individual subjects successfully limit their own interaction with platforms, preserving cognitive resources and partial unpredictability (right to be unpredictable, N1). But this individual achievement does not alter the structure of the system: the platform continues to function, profile, and extract predictive value from billions of other subjects. Analogously, informal collective movements (digital minimalism, media literacy campaigns, regulatory activism) produce local changes in the behavior of participants, but do not alter the architecture of the system that produces the problem. Regularity 11 demonstrates that any informal collective action existing within a ranked environment is marginalized by the same mechanisms it attempts to critique. It therefore follows that systemic change requires a form capable of acting not within the ranked environment, but at the level of the rules that determine its architecture.

Why constitutional form is the sole remainder. After the sequential exclusion of three classes, constitutional architecture remains. Constitutional form possesses three properties absent from the excluded alternatives. First, externality with respect to the logic of capital: a constitution is not a product of market relations, does not depend on platform engagement metrics, and is not the object of predictive monetization. Second, formality: a constitutional norm is executable in the sense that its violation entails legal consequences independent of the predictive power of the violator. Third, institutional stability: a constitutional norm possesses a higher threshold of modification than ordinary legislation, which reduces vulnerability to lobbying by platform capital. Constitutional form is the sole class preserving all three properties under the conditions described in T8 and T9. → Derivation: T10 is a theorem, not a normative preference. Constitutional necessity is logically derived from the structural properties of the system by the method of exclusion.

Empirical supplement: The case of the EU Digital Services Act (2022) and the EU AI Act (2024) demonstrates both the limit and the significance of regulatory efforts. Both acts are significant legal instruments creating normative standards and enforcement mechanisms. However, both are products of the ordinary legislative process, modifiable through the same process, contain exceptions for state security authorities, and are applied to transnational platforms with substantial practical limitations. They illustrate the necessity of the regulatory supplement to constitutional architecture — and simultaneously its insufficiency without it.

## Chapter Summary

Chapter 9 completed the axiomatic proof of Volume I through three theorems of the epistemological layer. T8 demonstrated that the rupture between de facto predictive power and de jure political sovereignty is intensifying and does not self-correct under the structural conditions of Regularities 1–13. T9 demonstrated that the absence of an external constitutional constraint inevitably produces one or a combination of three structural outcomes — digital dictatorship, prediction plutocracy, or the disintegration of the public sphere. T10 demonstrated by the method of exclusion that the sole form capable of resolving the contradiction of T8 and precluding the outcomes of T9 is a formal constitutional architecture external to the logic of digital capital. Subchapter 9.1 unfolded each link of the exclusion, disposing of possible objections.

### **Transition to Part IV**

The three layers of Volume I — ontology (A1–A6, Regularities 1–4, T1–T3), anthropology (A7–A12, Regularities 5–9, T4–T7), and epistemology ( $\Sigma$ A13– $\Sigma$ A18, Regularities 10–13, T8–T10) — jointly constitute the exhaustive diagnosis of digital capital. T10 demonstrated constitutional necessity. However, the diagnosis remains incomplete without a normative layer: the necessity of constitutional form does not answer the question of the content of its normative principles. Part IV introduces the normative axiom NA0 and seven normative principles N1–N7 — as logical derivations from the structural contradictions demonstrated in Parts I–III, not as independent normative judgments. NA0 is the connecting element between the diagnosis of Volume I and the constitutional architecture of Volume III.

## **PART IV. NORMATIVE SYNTHESIS**

### *What Any Alternative Must Protect*

This part effects the transition from diagnosis to obligation. If the first three parts described how the mechanism of alienation operates, Part IV establishes why it is impermissible and which limits must not be transgressed.

An important clarification: the architecture of Virtublic is not described here — that is the task of Volume III. At this stage, the Normative Standard is formulated: seven principles (N1–N7) with which any solution claiming legitimacy is obligated to comply.

The foundation of this transition is the normative axiom NA0: the recognition of subjecthood as the supreme good, the protection of which stands above any economic efficiency. This is a declaration of human rights in the age of predictive capital.

### **Chapter 10. The Normative Axiom NA0 and Its Entailments**

Parts I–III constructed a complete description of digital capital at three levels. The ontological layer established the structure of the alienation of attention and its self-valorizing dynamics. The anthropological layer described the mechanism through which the subject reproduces the conditions of his own alienation — through cognitive atrophy, neural adaptation, and the degradation of the instruments of autonomous choice. The epistemological layer registered

that the subject is systemically deprived of the information necessary for the informed assessment of the conditions of the relation in which he participates.

The conjunction of T1–T10 constitutes an analytically complete diagnosis. However, a diagnosis as such is not a normative judgment: the description of a structure producing certain consequences does not itself contain the assertion that these consequences are impermissible. The transition from description to evaluation requires a normative foundation that is not derivable from ontology and cannot be obtained from empirical data.

This foundation is introduced in the present chapter in the form of the normative axiom NA0. The chapter performs three tasks: it formulates NA0 with the necessary precision, establishes its status relative to the ontological and anthropological layers of the analysis, and derives from it the normative entailments N1–N7, which constitute the content of the subsequent chapters of Part IV.

### **10.1. The Status of NA0: An External Norm, Not an Immanent Entailment**

**NA0.** Subjecthood (hereafter: the subject's capacity for autonomous choice on the basis of unimpaired cognitive resources) is a politically protected good. Its systematic destruction is a political evil regardless of the economic efficiency of that destruction. The destruction of subjecthood includes the measurable physiological depletion of the cognitive resources necessary for autonomous choice.

Justification of status. NA0 is an axiom in the strict sense: it is not derived from the preceding elements of the trilogy, but introduced as the independent foundation of the normative layer. This is methodologically necessary, since normative judgments are not derivable from descriptive premises without the introduction of an independent normative premise. Neither axioms A1–A6, nor Regularities 1–4, nor theorems T1–T10 contain evaluative judgments: they establish what occurs, but do not establish whether what occurs is permissible. NA0 introduces precisely this element.

It is essential to register what NA0 is not. NA0 is not an entailment of Marxist immanent critique — the logic in which the system contains within itself the seed of its own negation through a contradiction resolved by the structure itself. The anthropological layer demonstrated the contrary: digital capital does not produce a conscious subject capable of organized resistance — it produces a cognitively atrophied individual with degraded instruments of autonomous choice [see T6]. There is no immanent gravedigger within the structure. This means that the normative foundation must be external to the structure — that is, introduced as an axiom, not derived as a theorem.

NA0 is Kantian in its logical form: it asserts that the subject possesses a dignity whose destruction is impermissible regardless of whether this destruction produces economically efficient results. This means that arguments of the form "the platform generates economic value" or "users voluntarily interact with the platform" are not refutations of NA0: they appeal to the criterion of efficiency or will, whereas NA0 appeals to the criterion of dignity. These criteria are logically independent, by virtue of which arguments from efficiency are irrelevant to NA0.

Connection to the preceding parts: NA0 does not contradict the ontological layer. The ontological layer describes how the structure is constituted. NA0 adds a normative qualification: this structure, by producing the systematic destruction of subjecthood, is politically impermissible. The former is a descriptive proposition. The latter is normative. They belong to distinct logical layers and do not enter into contradiction.

Operational definition of subjecthood. In order for NA0 to be applicable to real situations, rather than remaining an abstract principle, subjecthood requires an operational definition. Within the framework of the trilogy, subjecthood is defined as the conjunction of three measurable components: cognitive capacity — the possession of the neurobiological resources necessary for deliberate decision-making, including prefrontal cortex activity and cortisol levels within the physiological norm [see T5 of the anthropological layer]; informational competence — the possession by the subject of sufficient and verifiable information concerning the conditions of the relation into which he enters [see T8, T9 of the epistemological layer]; structural alternativity — the existence of real, structurally equally accessible alternatives, precluding structural coercion toward the alienation of attention [see T2, Regularity 4].

The destruction of subjecthood, qualified by NA0 as a political evil, denotes the measurable decline of any of the three components below the threshold at which autonomous choice remains operationally possible. The inclusion of the physiological dimension — the depletion of cognitive resources — in the definition of NA0 is fundamental: it translates the normative proposition from the domain of abstract political philosophy into the domain of measurable biomedical indicators, which renders NA0 operationally verifiable, not merely declarable.

Limit of section 10.1: NA0 has been established as an external normative axiom with operational content. In order for NA0 to produce normative entailments, it is necessary to derive from it the specific requirements for any solution claiming to protect subjecthood. This task constitutes the subject matter of section 10.2.

## **10.2. Normative Entailments N1–N4: Protection of the Subject**

NA0 in conjunction with theorems T1–T10 produces seven normative entailments — N1–N7. Entailments N1–N4 describe requirements for the protection of the subject directly. Entailments N5–N7, which constitute the subject matter of section 10.3, describe requirements for the institutional form of protection.

**N1.** From NA0 + T5 + T6. Any solution is obligated to preclude the measurable physiological depletion of the subject's cognitive resources produced by the system. A solution that reduces the level of attention alienation but preserves the neurobiological mechanisms of subjecthood degradation does not satisfy NA0.

Justification. T5 [see the theorem on physiological degradation of the anthropological layer] established that chronic interaction with a system optimized for the maximization of AT produces a measurable decline in the subject's cognitive capacity — through the mechanisms of cortisol loading and the degradation of the regulatory functions of the prefrontal cortex. T6 [see the theorem of cognitive disarmament] registered that this process produces a subject with a structurally diminished capacity for resistance to further attention alienation. From NA0 it follows that since physiological depletion is a form of the destruction

of subjecthood, any solution is obligated to extirpate not only the superstructural manifestations of alienation, but also its physiological mechanisms. The reduction of visible forms of manipulation without the extirpation of the neurobiological mechanisms of their impact does not constitute the protection of subjecthood in the sense of NA0.

This entailment has direct significance for the evaluation of proposed solutions. The restriction of screen time without alteration of the interaction architecture does not extirpate the mechanism: a subject interacting with the same algorithmic optimization function for a reduced duration is subjected to the same impact at lesser intensity, but of the same type. N1 requires not a reduction of the dose of impact, but the extirpation of the very mechanism that produces physiological harm.

**N2.** From NA0 + T8 + T9. Any solution is obligated to ensure the informational competence of the subject: the availability of verifiable information concerning the conditions of attention alienation in a form cognitively accessible to the subject at the moment of interaction.

Justification. T8 [see the theorem on the opacity of conditions] and T9 [see the theorem on informational asymmetry of the epistemological layer] established that the subject is systemically deprived of the information necessary for the informed assessment of the conditions of the relation. The conditions of alienation are concealed in legal documents cognitively inaccessible to most subjects by virtue of their volume and language, and the real value of the subject's predictive profile is not disclosed to him. From NA0 it follows that the absence of informational competence is a form of the destruction of subjecthood, by virtue of which any solution is obligated to restore it. This requirement is not satisfied by the formal provision of information in any form — it requires the provision of information in a cognitively accessible form at the moment when the subject is making a decision concerning the alienation of attention. Consent obtained without the informational competence of the subject is not consent in the sense of NA0.

**N3.** From NA0 + T2 + Regularity 4. Any solution is obligated to ensure structural alternativity: the existence of real, accessible alternatives to interaction with the system that produces the destruction of subjecthood, without asymmetric social or infrastructural exit costs.

Justification. T2 established the structural impossibility of competition within the modal layer after the point of no return. Regularity 4 established that the temporal barrier increases monotonically. From their conjunction it follows that under conditions of a dominant platform, the subject is deprived of a real alternative to interacting with it without significant social or infrastructural costs — the loss of the communications environment, professional connections, and access to informational resources. This situation constitutes structural coercion toward the alienation of attention, which is qualified as the destruction of subjecthood in accordance with its operational definition introduced in section 10.1. From NA0 it follows that any solution is obligated either to extirpate structural coercion through the provision of real alternatives, or to alter the architectural principle of the dominant platform such that interaction with it ceases to produce the destruction of subjecthood.

N3 has a direct entailment for the evaluation of regulatory solutions: measures that reduce individual abuses by the dominant platform but do not extirpate the structural coercion to interact with it do not satisfy NA0 — since the subject remains deprived of a real alternative.

**N4.** From NA0 + T5 + N1. Children below the age of completion of the neurobiological formation of self-regulation mechanisms are subject to heightened normative protection: a solution that permits the impact of the mechanisms described in T5 upon subjects with an incompletely formed prefrontal cortex does not satisfy NA0.

Justification. The operational definition of subjecthood in NA0 includes the possession of the cognitive resources necessary for autonomous choice. In a subject with an incompletely formed prefrontal cortex — that is, within the age range of approximately up to 25 years, with the most vulnerable period being 6–17 years [see subsection 1.3.1] — the physiological mechanisms of impulse control are structurally not yet formed. The impact of a system optimized for the maximization of AT through variable ratio reinforcement and FOMO mechanics [see A6 and 1.3.1] upon a subject lacking a formed neurophysiological counterweight produces an intensity of subjecthood destruction that structurally exceeds the analogous impact upon an adult subject. From NA0 it follows that the heightened vulnerability of this group requires a heightened normative standard of protection, not the application of the same standard with an age adjustment. Any solution establishing a unified regime for adults and minors does not satisfy N4.

Transition: entailments N1–N4 describe requirements for the content of the protection of the subject. They do not describe requirements for the form of the institutions effectuating this protection. Since T3 established the absence of an internal correction mechanism within the system, and T2 established the structural impossibility of market correction, the protection of subjecthood requires external institutions possessing specific properties. These properties are formalized in entailments N5–N7.

### **10.3. Normative Entailments N5–N7: Requirements for Institutional Form**

**N5.** From NA0 + T3 + Regularity 3. An institution claiming to protect subjecthood is obligated to be independent of the economic incentives of the platform: any institution financed by the platform, regulated by it, or dependent upon its infrastructure produces a structural conflict of interest that precludes its qualification as an instrument of subjecthood protection in the sense of NA0.

Justification. T3 established the complete absence of an internal correction mechanism for the concentration of predictive capital. Regularity 3 established that predictive power increases monotonically. From their conjunction it follows that an institution embedded in the system or dependent upon it is subordinated to the same logic of predictive capital accumulation as the platform. Self-regulatory bodies established by platforms serve as an example: their financing, personnel composition, and operational procedures are structurally dependent upon the objects of regulation, by virtue of which their decisions systematically reproduce the interests of the platforms rather than the interests of the subjects. From NA0 it follows that institutional independence is not an optional quality, but the necessary condition for an institution to be capable at all of performing the function of protecting subjecthood. A solution presupposing embedded or dependent institutions does not satisfy N5.

Example: in 2023–2025, the handling of complaints within the Meta Oversight Board systematically demonstrated limitations conditioned by the Board's dependence on Meta's infrastructure and financing (source: Oversight Board reports, 2023–2025). This confirms N5 as a structural, not incidental, property of dependent institutions.

**N6.** From NA0 + A2 + T8. An institution claiming to protect subjecthood is obligated to operate with verifiable standards of consent: the subject's consent to the alienation of attention is valid only in the presence of informational competence (N2) and structural alternativity (N3). Consent obtained in their absence is not consent in the normative sense and does not extirpate from the relation the characteristic of structural coercion.

Justification. A2 distinguished voluntary alienation in the presence of real alternatives from structural alienation without alternatives. T8 established that informational asymmetry precludes the informed assessment of the conditions of the relation. From the conjunction of A2 and T8 it follows that the majority of the subject's acts of interaction with the dominant platform do not satisfy the criterion of voluntary alienation — they are effectuated under conditions of structural coercion and informational deficiency. From NA0 it follows that an institution qualifying the formal act of pressing the "Accept" button as valid consent reproduces the normative defect of the relation it is called upon to regulate. N6 requires that the standard of consent include the verification of the informational competence and structural alternativity of the subject at the moment of consent — and not merely the fact of formal confirmation.

**N7.** From NA0 + T3 + N5. An institution claiming to protect subjecthood is obligated to possess constitutional status: protection resting solely upon the legislative or regulatory level is structurally insufficient, since both levels are subject to modification through political processes in which the predictive capital of platforms is a significant resource of influence.

Justification. T3 established the absence of an internal correction mechanism within the system and the necessity of external institutional intervention. N5 requires the independence of this institution. From their conjunction it follows that the institution must be independent not only from the platform, but also from the political processes in which the platform is an influential participant. The legislative and regulatory levels do not ensure this independence structurally: laws are enacted and repealed within political processes in which the predictive capital of platforms is convertible into political influence [see the theorem on the conversion of predictive capital]. The constitutional level possesses structurally distinct properties of resistance to such conversion.

N7 establishes that the project of cyber-constitutionalism described in Volume III is not a political preference, but the logically necessary derivation from NA0 in conjunction with T3 and N5: the protection of subjecthood requires a constitutional institution precisely because lower levels of normative protection are structurally vulnerable to the influence of that from which they are called upon to protect.

Example: analysis of the process of adoption of the Digital Markets Act in the EU (2022–2024) demonstrates that during the period of the regulation's development, the largest platforms allocated more than 100 million euros to lobbying in Brussels (source: Corporate Europe Observatory, 2023). This confirms N7 empirically: predictive capital is converted into political influence over the regulatory process, reducing the effectiveness of the regulatory level of protection.

#### **10.4. NA0 as the Condition of the Normative Unity of Part IV and Beyond**

The seven entailments N1–N7 constitute a normative architecture fully determined by NA0 in its interaction with theorems T1–T10. They are not independent policy recommendations — each is derived from a specific theorem or regularity in the presence of NA0 as the normative premise. This means that the negation of any of the entailments N1–N7 would require either the negation of the corresponding theorem (that is, an alteration of the descriptive account), or the negation of NA0 (that is, the acceptance that the destruction of subjecthood is permissible). Both paths are open to the opponent, but both presuppose an explicit position on the corresponding question.

NA0 also performs the function of normative unity: it ensures that Parts I–III and Part IV describe a single subject matter — digital capital as a structure that systematically destroys subjecthood, and the normative requirements for its supersession. Without NA0, the parts of the trilogy remain logically connected but normatively disjointed: analysis and norm exist in separate spaces. With NA0, they constitute a single construction: the diagnosis and the normative entailments from it.

The subsequent chapters of Part IV unfold each of the entailments N1–N7 into the fullness of its normative requirement. Chapter 11 examines N3 — the requirement of structural alternativity — in its application to the question of what architectural properties alternative systems must possess. Chapter 12 unfolds N7 — the requirement of constitutional status — in the direction that establishes the frame for Volume III.

## **Chapter Summary**

Chapter 10 performed three tasks. First: the normative axiom NA0 was introduced with an operational definition of subjecthood through three measurable components — cognitive capacity, informational competence, and structural alternativity. Second: the status of NA0 was established as an external Kantian axiom, not derivable from the descriptive layer and not contradicting it — as against Marxist immanent critique, which is structurally insufficient as applied to digital capital. Third: from NA0 in conjunction with T1–T10, seven normative entailments were derived — N1–N7 — of which N1–N4 establish requirements for the content of the protection of the subject, and N5–N7 establish requirements for the institutional form of that protection. The cardinal derivation is N7: the protection of subjecthood requires the constitutional status of the institution, since lower normative levels are structurally vulnerable to the influence of predictive capital.

## **Transition to Chapter 11**

Chapter 10 established the normative framework obligatory for any solution claiming to protect subjecthood. Chapter 11 proceeds to the elaboration of N3 and N5 in the direction of concrete architectural requirements for an alternative system: if structural alternativity is a normative requirement and the institution of protection is obligated to be independent of the platform, then what structural properties must a system satisfying both requirements simultaneously possess. The answer to this question constitutes the normative bridge to the architectural layer of Volume III.

## **Chapter 11. The Seven Normative Principles**

Chapter 10 established the normative axiom NA0 and derived from it seven entailments — N1–N7 — describing requirements for the content of subjecthood protection and for the institutional form of that protection. The entailments were formulated at the level of general normative requirements: they specify what a solution is obligated to ensure, but do not describe the specific content of the corresponding constitutional norms.

The present chapter proceeds to the level of normative principles. The seven principles N1–N7 introduced in this chapter are the substantive elaboration of the entailments of NA0 as applied to the specific structural violations of subjecthood registered in Parts I–III. Each principle is derived from NA0 in conjunction with specific theorems and regularities, by virtue of which its normative obligatoriness is derivative of the corresponding descriptive account. Each principle simultaneously serves as the point of entry into the constitutional architecture of Volume III: it establishes the normative requirement with which principles P0–P18 are obligated to comply as its institutional realization.

The structure of each principle comprises a precise formulation of the norm, an analytical justification with references to the relevant elements of the trilogy, a methodological note on the limits and conditions of application, and an indication of the connection to the constitutional layer of Volume III.

## **11.1. N1 and N2: The Sovereignty of the Subject and the Prohibition of Conversion**

### **N1. Right to Unpredictability.**

Formulation: the subject possesses the normative right not to be predicted without explicit free consent; the freedom of consent is determined by the presence of real structural alternatives and the informational competence of the subject at the moment of consent, not by the fact of formal confirmation.

Justification (from NA0 + A2 + T4). A2 distinguished voluntary alienation in the presence of real alternatives from structural alienation without alternatives. T4 [see the theorem on the predictive capture of subjecthood of the anthropological layer] established that a predictive model constructed on the basis of the subject's accumulated AT produces the subject's behavioral states with a probability structurally exceeding chance — that is, exercises functional control over the subject's behavior without his conscious participation. From NA0 it follows that since predictive capture is a form of the destruction of subjecthood — it produces behavior not determined by the subject's autonomous choice — the subject possesses the normative right not to be the object of predictive capture without free consent. The critical element is the definition of freedom of consent: it must include the verification of structural alternativity (N3 from Chapter 10) and informational competence (N2 from Chapter 10). A formal click-through effectuated under conditions of a dominant platform without real alternatives does not satisfy this criterion.

Methodological note. N1 is not a prohibition on predictive modeling as such. It is a prohibition on predictive modeling without the free consent of the subject. This distinction is fundamental: aggregated predictive analytics on anonymized data, employed for purposes approved by subjects, is categorically distinct from individual predictive profiling that produces behavioral control. N1 regulates the second case.

Connection to Volume III. N1 establishes the normative requirement for principle P3 of Volume III (Soulbound Identity) — the mechanism of verifiable digital identity ensuring the subject control over the conditions of access to his predictive profile — and for principle P13 (Digital Census), which realizes the verification of consent through zero-knowledge proof: technically this denotes the capacity to verify the fact of the existence of consent without disclosing its content to third parties.

Limit of section: N1 establishes the subject's right to protection against predictive capture. N2 establishes the structural condition under which the economic power of predictive capital is not convertible into political power over those same subjects through other channels.

## **N2. Prohibition on the Conversion of Economic Capital into Political Power.**

Formulation: economic contribution to the infrastructure of the digital public sphere is not convertible into political sovereignty; an infrastructure operator participates in the economic relations of the digital system, but is deprived of political voice in the processes that determine its constitutional parameters.

Justification (from NA0 + T10 + Regularity 13). T10 [see the theorem on the conversion of predictive capital into political influence] established that the predictive capital of a platform is convertible into political influence through several channels: the financing of political processes, control over the informational flows that determine the formation of political preferences, and direct participation in regulatory processes. Regularity 13 [see the regularity on large-scale conversion of the epistemological layer] established that this conversion intensifies proportionally to the concentration of predictive capital. From NA0 it follows that since the conversion of economic capital into political power produces the destruction of the subjecthood of the political process — replacing the equal political sovereignty of citizens with the weighted influence of capital — it is politically impermissible regardless of the economic efficiency of that conversion.

Methodological note. N2 requires a precise technical definition of conversion in Volume III. At the level of normative principle, conversion is defined as any mechanism by which the volume of an agent's economic resources determines the volume of his political influence beyond a normatively established limit. The specific mechanisms realizing the prohibition on conversion — separate functional spaces for economic and political participation — constitute the subject matter of the constitutional architecture of Volume III. N2 establishes the principle; the mechanism of its realization is determined by principles P4 and P16 of Volume III.

Connection to Volume III. N2 is the normative foundation for principle P4 of Volume III (Dual Sovereignty: the separation of the spaces  $EQU \perp$  and  $VIC \perp$  — economic and political sovereignty) and principle P16 (Rockefeller Mode — the mechanism of the mandatory separation of the economic and political function of an agent upon the exceeding of threshold concentration indicators).

### **11.2. N3 and N4: Visibility and Audit**

#### **N3. Minimum Visibility Threshold.**

Formulation: any legitimate political position possesses the right to a minimum visibility threshold in the digital public sphere; this threshold may not be determined by the algorithmic engagement optimization function.

Justification (from NA0 + T5 + T8). T5 [see the theorem on the neutralization of political positions of the epistemological layer] established that the algorithmic engagement optimization function systematically suppresses political positions that do not produce maximum AT, regardless of their substantive significance to the political process. T8 [see the theorem on informational asymmetry] established that the subject does not possess information as to which positions are suppressed and on what basis. From NA0 it follows that since the systematic exclusion of political positions from the public space is a form of the destruction of the subjecthood of the political process in its entirety — it deprives subjects of access to the information necessary for autonomous political choice — a minimum visibility threshold is a normative requirement, not a political preference.

Methodological note. N3 does not establish equal visibility for all positions: this is an unrealizable requirement under conditions of limited audience cognitive resources. N3 establishes a minimum threshold — a lower bound of visibility below which a position is in effect excluded from the public space. Critically: the subject establishing this threshold can be neither the platform (conflict of interest pursuant to N5 of Chapter 10), nor the state, since the state is a participant in the political process with its own interests in determining the visibility of positions [see Regularity 12 of the epistemological layer]. N3 requires an independent constitutional body as the subject establishing and verifying the minimum threshold.

Connection to Volume III. N3 is the normative foundation for principles P10 and P11 of Volume III (Madison Mode with quadratic vote weighting and Success Multiplier — the mechanism that produces structural advantage for small but significant positions relative to their organic algorithmic visibility) and for the anti-faction filter limiting the algorithmic amplification of positions with maximum polarizing function.

#### **N4. Procedural Audit of Predictive Models.**

Formulation: predictive models that influence political processes are subject to mandatory procedural audit; the audit is effectuated through a zero-knowledge proof mechanism ensuring the verification of the model's correspondence to its declared parameters without disclosure of subjects' data or the operator's commercial secrets.

Justification (from NA0 + Regularity 10 + T9). Regularity 10 [see the regularity on the opacity of predictive models of the epistemological layer] established that predictive models are structurally opaque to subjects and to most regulatory bodies. T9 [see the theorem on epistemological asymmetry] established that this opacity produces a structurally uncorrectable asymmetry: the platform knows categorically more about the subject than the subject knows about the model producing the impact on his behavior. From NA0 it follows that since the opacity of the predictive model is the condition of its impact upon the subjecthood of the political process without the possibility of verification or resistance, procedural audit is a normative requirement. The zero-knowledge proof mechanism is the technically justified solution to the problem of verification without disclosure: it permits the

confirmation that the model corresponds to its declared parameters without disclosing either the subjects' data or the architectural details of the model.

Methodological note. Zero-knowledge proof verifies form — the model's correspondence to declared parameters. It does not verify content in respect of the objective function: if the declared objective of the model is formulated as "maximization of engagement," procedural audit will confirm that the model indeed maximizes engagement, without evaluating the normative permissibility of that objective itself. N4 by virtue of this limitation requires an additional constitutional requirement: the mandatory public declaration of the objective function of any predictive model that influences political processes — as the condition preceding audit.

Connection to Volume III. N4 is the normative foundation for principle P2 of Volume III (formal verification in Coq — the mechanism of mathematical verification of the correspondence of algorithmic decisions to constitutional requirements) and principle P18 (Conflict-Resolution Core — the body exercising constitutional control over the results of procedural audit).

### **11.3. N5 and N6: The State and the Developing Subject**

#### **N5. Prohibition on the State Being a Buyer of Predictions Without a Mandate.**

Formulation: the state is not entitled to acquire predictive data concerning its own citizens without an explicit constitutional mandate and independent oversight; the state is not entitled to acquire predictive data concerning citizens of other states for commercial or intelligence purposes without strict constitutional constraints and international supervision.

Justification (from NAO +  $\Sigma A17$  + Regularity 12).  $\Sigma A17$  [see the synthetic axiom on the state as a structurally interested agent] establishes that the state is not a neutral regulator, but a structurally interested participant in digital relations: it is simultaneously a potential regulator of platforms and a potential buyer of the predictive data produced by those same platforms. Regularity 12 [see the regularity on the conflict of interest of the state of the epistemological layer] established that this conflict of interest is structural, not incidental: the state possesses systematic incentives to use predictive data for purposes incompatible with the protection of citizens' subjecthood. From NAO it follows that since the state's acquisition of predictive data concerning citizens without a constitutional mandate is a form of the destruction of subjecthood through a categorically distinct — state — channel of power, it is politically impermissible on the same grounds as the corporate alienation of attention.

Methodological note. N5 is the normatively most politically resistant principle among all seven: the states that are the subjects obligated to realize N5 possess a direct structural incentive against its application. This renders a declaratory prohibition insufficient: N5 requires an enumeration of specific prohibited practices — not merely the formulation of the principle — embedded at the constitutional level and protected against correction through ordinary legislative processes. The full realization of N5 is possible only within a supranational constitutional architecture, since national states are structurally incapable of ensuring self-limitation of this type: this would require the state voluntarily and irreversibly to constrain its own access to an instrument of power that it has systematic incentives to expand.

Connection to Volume III. N5 is the normative foundation for the preamble P0 of Volume III (the normative basis of the supranational digital constitutional architecture) and principle P17 (SovereigntyShield — the mechanism of constitutional constraint on the state as a buyer of predictive data in the supranational regime).

## **N6. Protection of Developing Subjecthood.**

Formulation: a subject whose capacity for autonomous choice is in the stage of formation — that is, persons under the age of eighteen — possesses the right to additional normative protection from mechanics that produce the destruction of the conditions of subjecthood formation; this protection includes the prohibition of gambling-like mechanics without strict verified consent and independent audit, the restriction of endless progression without organic completion points, the prohibition of dark patterns and emotional manipulation, as well as the mandatory disclosure of all persuasive mechanics and public audit of retention algorithms.

Justification (from NA0 + A8 + A12 + T4 + T6). A8 [see the axiom on the physiological vulnerability of the developing subject] established that in subjects with an incompletely formed prefrontal cortex, the impulse control mechanism is structurally absent as a fully operational physiological resource. A12 [see the axiom on reflexivity as the condition of resistance] established that resistance to the structural alienation of attention requires an active reflexive capacity, which is also a function of the cognitive resources that are in the stage of formation. T4 and T6 [see the corresponding theorems of the anthropological layer] jointly established that the predictive capture and cognitive disarmament of the subject intensify under conditions of a system optimized for the maximization of AT, and that in a minor subject this process proceeds with structurally diminished resistance. From NA0 it follows that since developing subjecthood is the most vulnerable to destruction — precisely because the mechanisms of its protection are not yet formed — it requires a heightened normative standard, not the application of a general standard with an age adjustment.

The question of consent as applied to N6 requires separate clarification. A developing subject cannot give free consent to the alienation of attention that structurally destroys the conditions of the formation of his subjecthood: such consent would be a self-negating act — consent to the destruction of the capacity for future consent. The consent of parents or legal representatives is a necessary but insufficient condition: it does not of itself verify the normative permissibility of specific impact mechanics. N6 requires an independent body verifying that specific gaming or platform mechanics are not gambling-like in their functional architecture — regardless of the existence of consent.

The empirical basis of N6 includes: a series of FTC settlement agreements with mobile game operators in 2025–2026 concerning gambling-like mechanics directed at minors; class action suits against platforms in US and EU jurisdictions concerning addiction design; the introduction of special legislative regulation in Brazil (2025) and Germany (2024) qualifying certain mechanics as a matter of the protection of minors' rights. This empirical record does not constitute the normative basis of N6 — the normative basis is NA0 — but it confirms that the structural conflict between A6 and NA0 on this group of subjects is operationally observable.

Connection to Volume III. N6 is the normative foundation for principle P14 of Volume III in respect of the additional protection of developing subjects, and for principle P3 (Soulbound Identity) in respect of age verification as the condition of access to certain types of digital interaction.

#### **11.4. N7: Cognitive Autonomy**

##### **N7. Right to Cognitive Autonomy.**

Formulation: the subject possesses the inalienable right to protection from the systematic cognitive and psychological atrophy produced by platform architecture; this right includes the prohibition of dark patterns (design decisions that produce the subject's behavior contrary to his declared preferences), the prohibition of emotional manipulation through AI companions and persuasive loops, the requirement of cognitive impact assessment for algorithmic systems with an engagement-optimization function, and the right to digital detox without punitive consequences.

Justification (from NA0 + A8 + Regularity 7 + T6 +  $\Sigma$ A15). A8 established the physiological measurability of cognitive atrophy as a real neurobiological state. Regularity 7 [see the regularity on the production of cognitive atrophy of the anthropological layer] established that the algorithmic AT optimization function systematically produces this state as a byproduct of engagement maximization — not as an objective, but as the structurally inevitable consequence. T6 [see the theorem of cognitive disarmament] established that a subject with a depleted cognitive resource is structurally incapable of the reflexivity necessary for resistance to further attention alienation.  $\Sigma$ A15 [see the synthetic axiom on the measurability of cognitive atrophy] established that the decline in cognitive capacity is physiologically measurable: a sustained decline in sustained attention indicators of 18–20% under chronic interaction with engagement-optimized systems has been registered in neuroscientific studies of 2023–2025 (Nature Neuroscience). From NA0 it follows that since cognitive atrophy is not a metaphor for dissatisfaction but a measurable physiological state that destroys the operational conditions of subjecthood, the right to protection from its production is a normative requirement of the same logical force as the right to physical inviolability.

Methodological note. The inalienability of the right N7 means: the subject cannot effect a juridically valid waiver of this right — including through click-through consent. This follows from the operational definition of subjecthood in NA0: consent to the destruction of the cognitive capacity that is the condition of the very ability to consent is a self-negating act and does not produce a legally binding waiver of the right N7. This limitation categorically distinguishes N7 from the majority of rights that the subject is entitled to restrict at his own discretion.

The right to digital detox without punitive consequences is a direct entailment of N7: streak-breaking penalty mechanisms and the loss of accumulated progress upon cessation of interaction are architectural instruments of coercion to continue the alienation of attention, functionally producing an exit barrier of the same type as was described in Regularity 4 as applied to competition. N7 prohibits the application of such barriers, since they produce structural coercion toward interaction, which is incompatible with the operational definition of subjecthood in NA0.

Connection to Volume III. N7 is the normative foundation for principle P14 of Volume III (Proof-of-Offline with cognitive health bonus — the mechanism of constitutional incentivization of periods of absence from digital interaction as the condition for the preservation of the subject's cognitive capacity).

### **11.5. Normative Architecture as a System**

The seven principles N1–N7 constitute not a simple enumeration of requirements, but a normative architecture with an internal structure of mutual connections and dependencies.

N1 and N2 describe the sovereignty of the subject in two dimensions: N1 — protection against the predictive capture of individual subjecthood, N2 — protection against the conversion of economic capital into collective political subjecthood. The violation of N2 without the violation of N1 is theoretically possible: a system may not exercise predictive capture of an individual subject, but may permit the conversion of capital into political power. The violation of N1 while observing N2 is also possible. Both principles are necessary as independent normative requirements.

N3 and N4 describe requirements for the informational environment of the political process: N3 — the requirement pertaining to the visibility of positions, N4 — the requirement pertaining to the transparency of mechanisms. N3 without N4 creates the requirement of visibility without verification of by whom and how it is satisfied. N4 without N3 creates transparency while preserving the structural exclusion of positions. Both principles are functionally interdependent.

N5, N6, and N7 describe three specialized contexts of application of NA0: the state as a particular agent (N5), the developing subject as a particular protected group (N6), and cognitive capacity as the physiological condition of subjecthood (N7). All three principles are the specification of NA0 as applied to structural situations in which the general standard of subjecthood protection proves insufficient.

The mutual coherence of N1–N7 is ensured by the fact that all seven principles are derived from the single normative axiom NA0 in its operational definition introduced in Chapter 10. A contradiction between principles would be possible only in the presence of a contradiction within NA0 or within the theorems from which the individual principles are derived. No such contradiction has been identified in the preceding analysis.

### **Chapter Summary**

Chapter 11 elaborated the seven normative entailments of NA0 formulated in Chapter 10 into the fullness of normative principles N1–N7 with precise justifications, methodological conditions of application, and connections to the constitutional layer of Volume III. N1 established the right to unpredictability as the normative requirement for the regime of consent. N2 established the prohibition on the conversion of economic capital into political power. N3 established the minimum visibility threshold for political positions, independent of the algorithmic optimization function. N4 established the procedural audit of predictive models through the zero-knowledge proof mechanism. N5 established the prohibition on the state acquiring predictive data without a constitutional mandate. N6 established the heightened standard of protection for the developing subjecthood of minors. N7 established

the inalienable right to cognitive autonomy, including the right to digital detox without punitive consequences. The conjunction of N1–N7 constitutes a normative architecture fully determined by NA0 and theorems T1–T10, and serves as the exhaustive normative point of entry into the constitutional architecture of Volume III.

## **Transition to Chapter 12**

Chapter 11 completed the formulation of the normative content of the requirements for the protection of subjecthood. N1–N7 establish what must be ensured. They do not describe the architecture of an institution capable of ensuring all of this simultaneously — an institution possessing the necessary properties of independence (N5 of Chapter 10), constitutional status (N7 of Chapter 10), and operational capacity to realize N1–N7 under conditions of systemic resistance from platforms and states. The description of such an institution constitutes the task of Chapter 12 and is the normative bridge to the constitutional architecture of Volume III.

## **Chapter 12. Justice of Visibility and the Distribution of Opportunities**

Chapter 11 formulated seven normative principles N1–N7, determined by NA0 and theorems T1–T10. Principles N3 and N7 remained at the level of normative requirements without operational specification: N3 established that any legitimate political position has the right to a minimum visibility threshold, yet did not define what that threshold constitutes or how it is verified. N7 established that the algorithmic optimization function cannot exhaust the cognitive resources of the subject, yet did not define which specific target optimization functions violate this requirement.

The present chapter performs the operational specification of these two principles. Section 12.1 defines the minimum visibility threshold through the obverse of T5 and establishes the procedural standard of participation minima. Section 12.2 defines the ethical limits of algorithmic optimization through the conjunction of NA0 and N7, formulating a prohibition not on algorithms as such, but on specific classes of target functions. The conjunction of the two sections constitutes the operational bridge between the normative principles of Part IV and the constitutional architecture of Volume III.

### **12.1. Participation Minima**

N3 in its formulation from Chapter 11 establishes a normative requirement without operational content: the "minimum visibility threshold" remains functionally undefined until a criterion for its measurement is specified. The operational specification of N3 requires recourse to T5 [see the theorem on the neutralization of political positions by the epistemological layer] — not to its direct content, but to its obverse.

T5 established the mechanism of neutralization: the algorithmic engagement optimization function produces a probability of reach  $P(\text{reach})$  for each political position as a function of its engagement-score — the predictive audience engagement indicator. A position whose engagement-score falls below the threshold value  $M$  receives  $P(\text{reach}) < \varepsilon$ , where  $\varepsilon$  is a functionally null value at which the position is formally present within the system, yet its

actual dissemination is incomparably small relative to positions with a high engagement-score. This is structural marginalization in the precise operational sense: not a prohibition, but functional exclusion through the probabilistic asymmetry of reach.

**Regularity 5** (from T5 + N3). The minimum visibility threshold is operationally defined as a guaranteed  $P(\text{reach}) \geq \epsilon$  for any legitimate political position, irrespective of its engagement-score.

Justification (from T5 + N3 + NA0). T5 established that in the absence of an external constraint, the algorithmic optimization function produces  $P(\text{reach})$  as a monotonically increasing function of engagement-score, which structurally necessitates the structural marginalization of low-engagement-score positions irrespective of their substantive significance for the political process. N3 established the normative requirement that this mechanism cannot determine the visibility of legitimate political positions. From NA0 it follows that, since the functional exclusion of positions from the public sphere constitutes a form of destruction of the subjecthood of the political process as a whole — it deprives subjects of access to the information necessary for autonomous political choice — guaranteed minimum reach is a normative requirement whose implementation cannot be delegated to an algorithmic optimization function.

Conclusion:  $P(\text{reach} \mid \text{position is legitimate}) \geq \epsilon$  for any engagement-score of that position. The value of  $\epsilon$  is established by an independent constitutional body and constitutes a lower bound — not an average and not a target value — of reach for any position qualified as a legitimate political position.

Criterion of positional legitimacy. Regularity 5 presupposes an operational definition of the legitimacy of a political position, since the guarantee of reach applies exclusively to positions satisfying this criterion. Legitimacy in this context is a procedural, not a substantive, concept: a position is legitimate if it does not violate constitutionally established constraints (prohibition of incitement to violence, prohibition of discrimination on protected grounds, and analogous norms), irrespective of whether its content is majoritarian, minoritarian, popular, or unpopular. A substantive criterion of legitimacy is impermissible precisely because any substantive criterion reproduces the power of the entity establishing the criterion over the composition of admissible political positions. N3 is directed against this power — it cannot be implemented through it.

Demarcation from censorship. N3 and Regularity 5 do not mandate equal reach for all positions: the requirement  $P(\text{reach}) \geq \epsilon$  for all legitimate positions does not entail that  $P(\text{reach}) = \text{const}$  for all positions. A position with a high engagement-score may receive  $P(\text{reach}) \gg \epsilon$  without violating N3. N3 establishes only the lower bound below which a position is functionally excluded from the public sphere. This is not censorship in the reverse direction — the compulsory promotion of undesirable positions — but the minimum procedural standard of a democratic public sphere. The distinction is analogous to the distinction between the minimum franchise — every citizen has the right to vote, irrespective of their views — and the compulsory redistribution of votes. N3 establishes the former.

Subject of implementation. Regularity 5 requires a guaranteed minimum reach, yet does not define the subject who provides this guarantee. From N5 of Chapter 10 it follows that this subject can be neither the platform (structural conflict of interests) nor the state (a participant

in the political process with its own interests in defining the visibility of positions). The sole permissible subject is an independent constitutional body satisfying the requirements of N5 and N7 of Chapter 10 — that is, organizationally independent of platforms, states, and their financing, operating on verifiable procedural standards, and subject to public audit through mechanisms compatible with N4.

Example: in 2024–2025, the European Commission, applying the Digital Services Act, required the largest platforms to provide data on the distribution of reach across various political positions during the campaigns for the European Parliament (source: European Commission Digital Services Act enforcement reports, 2024). The results demonstrated a systematic asymmetry of reach in favor of high-engagement-score positions, confirming the operational description of T5. Characteristically, the regulatory body — the Commission — exercised oversight over the platforms' data, yet did not possess the authority to establish a minimum reach threshold. This empirically confirms both the applicability of Regularity 5 and the operational limit of state regulation as the subject of its implementation.

The operational specification of N3 through Regularity 5 establishes a measurable standard and identifies the subject of its implementation. However, Regularity 5 does not define which algorithmic mechanisms produce a violation of this standard or which target optimization functions are normatively permissible. This constitutes the subject of Section 12.2.

## 12.2. Ethical Limits of Algorithmic Optimization

N7 established the inalienable right of the subject to protection against the systematic cognitive exhaustion produced by platform architecture. In its formulation from Chapter 11, this right includes the prohibition of certain practices — dark patterns, emotional manipulation, engagement-maximization without cognitive impact assessment — yet does not furnish an analytical criterion enabling the demarcation of permissible from impermissible target functions of algorithmic optimization in the general form. Without such a criterion, N7 remains open to arbitrary expansion or contraction in enforcement practice.

**Regularity 6** (from NA0 + N7 + A8 + T6). Algorithmic optimization is normatively impermissible if its target function is directly oriented toward the measurable reduction of the cognitive resources of the subject that are necessary for autonomous choice, or toward the production of states in which the subject is deprived of the operational capacity for reflection.

Justification (from NA0 + N7 + A8 + T6). NA0 established that the destruction of subjecthood constitutes a political evil irrespective of the economic efficiency of that destruction. N7 specified this in application to cognitive exhaustion. A8 [see the axiom on the physiological measurability of cognitive resources] established that the cognitive resources of the subject are measurable physiological variables. T6 [see the theorem on cognitive disarmament] established that a subject with an exhausted cognitive resource is structurally incapable of the reflection necessary for autonomous choice — that is, loses operational subjecthood. From NA0 it follows that an algorithmic optimization function whose target function includes the reduction of the subject's cognitive capacity or the production of states in which autonomous choice is impossible constitutes a direct instrument of the destruction of subjecthood and is, in consequence of which, normatively impermissible. The prohibition applies not to algorithms as a technical instrument, but to a specific class of target functions, defined through their directedness.

Three prohibited classes of target functions. From Regularity 6 in conjunction with the operational definition of subjecthood from NA0, three classes of target functions normatively impermissible under any circumstances are derived.

The first class: the prediction and formation of political convictions, religious views, and sexual identity of the subject without their verified free consent. Justification: these categories are constitutive elements of subjecthood in the sense of NA0 — they determine what the subject is as an autonomous political agent. An algorithmic optimization function that produces the prediction or formation of these elements effectuates a direct intervention in the content of subjecthood, not merely in its functional resources. This violates NA0 on the most fundamental ground: not through the exhaustion of the subject's resources, but through the substitution of the subject's autonomous self-determination by algorithmically produced content.

The second class: the systematic exhaustion of the subject's cognitive capacity through engagement-maximization — that is, target functions that include the maximization of the subject's interaction time with the platform through mechanics producing neurobiologically measurable reductions in the capacity for sustained attention. Justification (from A8 + T6 + Regularity 7 [see the regularity on the production of cognitive exhaustion]): this class of target functions produces precisely the state that T6 described as cognitive disarmament — a subject systematically subjected to engagement-maximization loses the physiological resource necessary to evaluate what is occurring to them. This is a self-reproducing destruction of subjecthood: the system reduces the subject's capacity for resistance to the system by deploying the subject's own resources toward that reduction.

The third class: the exploitation of the incomplete formation of cognitive mechanisms in minor subjects — that is, target functions specifically directed at the age range in which the prefrontal cortex is not yet fully formed, and which use this physiological fact as a source of heightened vulnerability to AT-alienation. Justification (from N6 + A8 + subsection 1.3.1): from N6 it follows that the forming subject requires an elevated normative standard precisely because the physiological mechanisms of their protection are not yet formed. A target function that deliberately exploits this absence does not merely produce cognitive exhaustion — it produces the destruction of subjecthood at the stage at which subjecthood is itself still forming, which yields consequences that are fundamentally incommensurable with the effects of such a function on an adult subject.

Requirement of cognitive impact transparency. Regularity 6 and the three prohibited classes of target functions are not self-sufficient normative instruments: their application requires verification of the class to which the target function of a specific algorithmic system belongs. Since the target functions of algorithmic systems are in the majority of cases not publicly disclosed — this is the epistemological asymmetry recorded by T8 and T9 — a normative prohibition requirement is insufficient without a concomitant requirement of transparency.

This requirement is formulated as the mandatory public declaration of the target function of any algorithmic system with engagement-optimization that affects the public sphere or interaction with minor subjects. The declaration is a prior condition: an algorithmic system that has not provided a declaration of its target function cannot be qualified as non-violating of Regularity 6 — by virtue of the fact that its conformity with normative requirements is not

verifiable. This logic conforms to the standard legal presumption: the burden of demonstrating conformity with normative requirements rests on the operator of the system, not on the regulatory body.

Demarcation of prohibition from regulation. Regularity 6 establishes three prohibited classes of target functions. Beyond these classes, algorithmic optimization is normatively permissible, yet may be subject to regulation — that is, conditions restricting its application without prohibiting it as such. The distinction between prohibition and regulation is defined as follows: prohibition applies to target functions directly producing the destruction of subjecthood in the sense of NA0; regulation applies to target functions producing other consequences that may be corrected without eliminating the function itself. This demarcation is critically important for precluding an expansive interpretation of Regularity 6, under which any algorithmic optimization could be qualified as violating N7.

Example: in 2025, the Federal Trade Commission (FTC) initiated an investigation into several social platforms regarding the application of persuasive design patterns directed at minor users (source: FTC enforcement actions, 2025). The investigation identified the systematic use of variable ratio reinforcement in social interaction mechanics functionally identical to gaming loot box systems. This case operationally corresponds to the third prohibited class of target functions formulated in Regularity 6: the target function of the system specifically exploited the neurobiological vulnerability of minor subjects to maximize AT. This confirms both the applicability of Regularity 6 and the operational distinguishability of the prohibited classes when a proper declaration of the target function is present.

Connection to Volume III. Regularity 6 and the requirement of cognitive impact transparency constitute the normative foundation for Principle P14 of Volume III with respect to the mandatory cognitive impact assessment for algorithmic systems with an engagement-optimization function — the mechanism for verifying the conformity of target functions with Regularity 6 prior to their deployment in the public digital environment.

## **Δ4 — CRISIS: THE LIMIT OF NORMATIVE SYNTHESIS**

Principles N1–N7 and their operational specification in Chapters 10–12 exhaust what can be said about the normative requirements for the protection of subjecthood at the level of principles. Part IV formulated what must be secured and what properties must be possessed by the institutions securing it. Part IV does not describe how exactly these institutions are structured — that is, their specific constitutional architecture, decision-making mechanisms, technical verification protocols, and procedures for resolving conflicts between principles. The normative synthesis is logically complete at the level of requirements and structurally incomplete at the level of institutional implementation. This limit is not a defect of normative analysis but its structural property: normative principles specify the criteria to which the architecture must conform, yet do not determine a single permissible architectural form. The transition from normative principles to the constitutional architecture of cyber-republicanism — to principles P0–P18 — constitutes the content of Volume III.

### **Chapter Summary**

Chapter 12 performed the operational specification of principles N3 and N7 through two regularities. Regularity 5 defined the minimum visibility threshold as a guaranteed  $P(\text{reach}) \geq$

ε for any legitimate political position, irrespective of its engagement-score, established the procedural criterion of legitimacy as substantively neutral, demarcated N3 from censorship in the reverse direction, and identified an independent constitutional body as the sole permissible subject of the guarantee's implementation. Regularity 6 established the analytical criterion of normative impermissibility of algorithmic optimization through the directedness of the target function and derived three prohibited classes: the prediction and formation of constitutive elements of subjecthood, the systematic exhaustion of cognitive capacity through engagement-maximization, and the exploitation of the incomplete formation of cognitive mechanisms in minors. The requirement of cognitive impact transparency is established as the procedural condition for verifying conformity with Regularity 6.

**Reference to Volume III.** Part IV completed the normative synthesis of the trilogy: NA0 introduced and operationalized, N1–N7 derived and specified, Δ4 recorded the limit of the normative layer. Volume I in its totality described digital capital as a structure systematically destroying subjecthood, established the normative grounds for qualifying this destruction as politically impermissible, and formulated seven principles to which the constitutional alternative must conform. The content of this alternative — the constitutional architecture of cyber-republicanism with principles P0–P18, institutional mechanisms of their implementation, and technical verification protocols — constitutes the subject of Volume III, "Theory of the Digital Republic."

## **Chapter 13. The diagnosis exhausted: What Volume I cannot resolve.**

The three parts preceding this chapter constructed a complete description of digital capital across three levels of analysis. The ontological layer described the structure as a self-augmenting system without internal limit. The anthropological layer described the subject existing within that structure with sovereignty structurally extracted. The epistemological layer described the system's reproduction of its own legitimacy, rendering critique unprofitable. Part IV added the normative layer: NA0 and principles N1–N7 established that this structure is normatively impermissible and formulated the requirements for its alternative.

The totality of the five parts of Volume I is diagnostically complete. This assertion requires precise qualification: diagnostic completeness means that no additional descriptive element would alter the normative qualification of the system, and that no descriptive supplement would itself generate the architecture of a response. A diagnosis cannot become therapy through expansion. The boundary between diagnosis and the architecture of a response is not a technical but a logical limit: it follows from the very nature of what Volume I has described.

The present chapter formalizes this limit across three sections. Section 13.1 records what has been proven across the three layers of analysis in a form that makes the structure of their combined conclusion visible. Section 13.2 establishes what Volume I structurally cannot offer — and why this limitation is logical rather than a deficiency of exposition. Section 13.3

describes how Volumes II and III inherit the diagnosis and what precisely each adds that Volume I cannot contain.

### **13.1. The Three Layers of Diagnosis: The Structure of the Combined Conclusion**

The ontological, anthropological, and epistemological layers of Volume I described a single subject matter from three positions. Their sequence is not arbitrary but logically necessary: each subsequent layer responds to a question that the preceding layer could not but raise.

The ontological layer: the system self-augments without internal limit. Axioms A1–A6 and Regularities 1–4 established that digital capital is a self-augmenting structure in which each cycle of attention alienation increases the predictive power of the platform, which in turn increases the efficiency of attracting the next cycle. Theorems T1–T3 formalized the terminal conclusions of this description. T1 established the surplus nature of predictive value in each cycle. T2 established that after the point of no return, competition within the modal layer is structurally impossible. T3 recorded the complete absence of any internal mechanism for correcting concentration.

The ontological layer, upon reaching T3, raised a question it could not resolve from within: who bears the consequences of the structure described by axioms A1–A6.  $\Delta 1$  — the crisis of the limit of objectivity — formalized this impossibility and structurally necessitated the transition to the anthropological layer.

The anthropological layer: the subject exists, but sovereignty is structurally extracted. The anthropological analysis established three dimensions of sovereignty extraction. The economic dimension [see T4]: the subject generates surplus predictive value without receiving compensation and without possessing a mechanism to claim it. The political dimension [see T5]: the algorithmic optimization function produces the systematic marginalization of political positions with a low engagement-score, determining the composition of the politically visible space independently of the subjects' will. The physiological dimension — T6, the theorem on cognitive disarmament — established that the system produces a measurable reduction in the subject's cognitive capacity necessary for resistance to further attention alienation. These three dimensions are not independent: each amplifies the other two through the mechanism described in cycle A6.

The anthropological layer raised a question it could not resolve from within: by what means does the subject with structurally extracted sovereignty reproduce the legitimacy of the system that produces that extraction.  $\Delta 2$  — the crisis of the limit of subjectivity — formalized this impossibility and structurally necessitated the transition to the epistemological layer.

The epistemological layer: the system reproduces its own legitimacy. The epistemological analysis established the mechanism by which digital capital renders critique unprofitable and absorbs it. T8 [see the theorem on informational asymmetry] recorded that the subject is systemically deprived of the information necessary to evaluate the conditions of the relation in which the subject participates. T9 [see the theorem on epistemological asymmetry] established that the gap between the information available to the platform and the information available to the subject accretes monotonically. T10 [see the theorem on the conversion of predictive capital into political influence] established that predictive capital is converted into the capacity to determine the conditions of its own regulation, which

structurally necessitates the ineffectiveness of regulatory responses on the existing institutional basis.

The combined conclusion of the three layers is formulated as follows: the system self-augments without internal limit, produces a subject with extracted sovereignty, and simultaneously produces the epistemological conditions under which that subject possesses neither information regarding the conditions of extraction nor instruments for its correction from within the system. This description is diagnostically complete in the sense that the addition of any further descriptive element would not alter the structure of the combined conclusion.  $\Delta 3$  — the crisis of the limit of epistemology — recorded that this layer is exhausted and structurally necessitated the transition to normative synthesis.

### **13.2. The Structural Limit of Volume I: What the Diagnosis Cannot Generate**

The diagnostically complete analysis described in Section 13.1 contains within itself a strict constraint: it cannot generate the architecture of a response. This constraint is logical rather than a deficiency of exposition. Its structure is determined by two arguments.

The first argument: the absorption of architecture from within the system's logic. Regularities 11 and 12 [see the regularities on the capture of critique and the conflict of interests of the regulator] established that any response offered within the same institutional logic reproduces, in a new form, the contradiction it is intended to resolve. A platform voluntarily reforming its business model reproduces cycle A6 in an altered casing. A state regulating a platform employs the instruments of the same predictive analytics that it regulates. A self-regulatory body financed by the platform generates decisions that structurally reflect the platform's interests. Critical discourse existing on the platform's infrastructure generates AT in favor of the very system it criticizes.

This is not a judgment regarding the intentions of agents — it is a structural conclusion from T3 and Regularity 12. T3 recorded the absence of an internal correction mechanism. Regularity 12 recorded the conflict of interests of the state as regulator. From their conjunction it follows that any response that does not exceed the institutional logic described in Parts I–III reproduces that from which it is intended to protect. Volume I described the system with sufficient completeness to demonstrate that the architecture of a response must be external to that system — that is, must rest on a different foundation of sovereignty, a different institutional form, and a different technological substrate.

The second argument: a normative principle is not an architectural resolution. Part IV of Volume I introduced NA0 and derived from it N1–N7. These principles establish what must be ensured. They do not establish how precisely this is to be arranged: what the constitutional architecture of an institution realizing N1–N7 must be; what technological substrate that institution must possess; how conflicts between principles are to be resolved in specific operational situations.

This limit is not a deficiency of normative analysis but its structural property: a normative principle establishes the criterion to which architecture must correspond but does not determine a single permissible architectural form for its realization.  $\Delta 4$  recorded precisely this limit of normative synthesis. The transition from normative principles to constitutional architecture requires two elements that Volume I does not and cannot contain: an analysis of

the technological substrate capable of realizing N1–N7, and the constitutional architecture of the institution itself. The first element constitutes the subject of Volume II. The second — the subject of Volume III.

### **13.3. Inheriting the Diagnosis: What Volumes II and III Add**

Volume II: blockchain as a technological attempt at resolution and its limit. Volume II inherits the diagnosis of Volume I in full: it does not reformulate the problem and does not revise theorems T1–T10. Its initial question is precisely determined by the diagnosis of Volume I: does there exist a technological substrate capable of realizing the requirements of N1–N7 outside the institutional logic that produces their violation?

Blockchain technologies constitute the most developed attempt to answer this question. Their ideological foundation — decentralization as a technical mechanism for excluding a single center of control — is the direct response to T3 (the absence of an internal correction mechanism): if no single party controls the registry, then the concentration of predictive capital in the hands of one agent becomes technically impossible.

Volume II demonstrates that blockchain as ideology fails [T11–T15] for reasons that follow logically from the diagnosis of Volume I. The decentralization of a protocol does not extirpate the concentration of economic power over the protocol: token-weighted governance reproduces the conversion of economic capital into political influence (N2) — the same structural defect as that of platforms, but on a new substrate. This is the direct consequence of Regularity 4 as applied to blockchain ecosystems: early participants accumulate tokens, the temporal barrier is reproduced in the form of concentrated governance tokens, and the formally decentralized protocol generates de facto oligarchic governance. The ideological assertion that "code is law" — that algorithmically encoded rules possess normative force in themselves — contains no answer to the question of who establishes those rules and whose interests they reflect. This is the epistemological defect described by T8–T9 as applied to platforms: the opacity of the decision-making mechanism is not extirpated by its algorithmic fixation.

At the same time, Volume II establishes that blockchain remains a necessary technological substrate [T16]. This assertion means the following: the realization of N1–N7 requires a technological layer that ensures trustless verifiability without reliance on a central agent, the immutability of a registry of decisions without the possibility of retroactive correction, and a zero-knowledge proof mechanism for the realization of N4. Blockchain technologies constitute the sole known technical solution satisfying all three requirements simultaneously. T16 asserts the necessity, but not the sufficiency, of this substrate: the technological layer instantiates a mechanism but does not determine the normative form of the institution employing that mechanism.

Volume III: the constitutional response, adding what technology does not contain. Volume III (Formal Theory of the Digital Republic) inherits both the diagnosis of Volume I and the technological analysis of Volume II. Its initial question is formulated by the conjunction of both preceding volumes: what must a constitutional institution be — one that employs the technological substrate established by T16 and corresponds to normative principles N1–N7 — in order not to reproduce the structural defects of either platforms or blockchain ideology?

Volume III adds three elements that neither Volume I nor Volume II can contain for structural reasons.

The first element: popular sovereignty as the normative foundation [preamble P0]. Blockchain protocols contain no concept of popular sovereignty: they contain rules but contain no subject in whose interests those rules exist. NA0 of Volume I introduced subjecthood as a politically protected good but did not define the form of the collective subject that protects this good institutionally. P0 of Volume III introduces this element: the totality of citizens as a political subject whose sovereign authority is the source of constitutional architecture — not as an aggregate of individual AT or as a totality of token holders.

The second element: the republican form of governance of the digital public sphere. Cyber-republicanism in the architecture of Volume III means: citizens govern algorithms through constitutional institutions, not algorithms govern citizens through the optimization of their behavior. This is the direct architectural response to T5 (the neutralization of political positions) and T10 (the conversion of predictive capital into political influence): the republican form structurally separates the functions of the economic operator and the political subject — precisely the demarcation that N2 established as a normative requirement.

The third element: the normative axiom as a constitutional norm. NA0 in Volume I is an axiom of the normative layer — an external premise introduced to qualify the descriptive account. In the constitutional architecture of Volume III, NA0 becomes a constitutional norm: its status is altered from a normative presupposition of theoretical analysis to a legally binding requirement whose violation constitutes grounds for constitutional review. This change of status cannot be effected within Volume I: a theoretical axiom does not become a constitutional norm through the expansion of theoretical analysis — that requires constitutional architecture, and it is precisely this that Volume III (Formal Theory of the Digital Republic) contains.

By way of illustration: in 2023–2024, a number of EU member states initiated legislative processes for the creation of independent oversight bodies for algorithmic systems in the public sphere (source: European Centre for Press and Media Freedom, 2024). Analysis of these initiatives reveals a structurally recurring defect: the bodies created are either embedded within the state administrative system — thereby reproducing the conflict of interests recorded by Regularity 12 — or are financed by platforms — thereby reproducing the defect described by N5 of Chapter 10. Not one of the examined cases contained a constitutional foundation in a form corresponding to P0 of Volume III. This empirically confirms the argument of the present chapter: diagnostically correct intentions, realized without constitutional architecture, reproduce the structural defects they intend to correct.

## **Chapter Summary**

Chapter 13 formalized the structural limit of Volume I and described the logic by which its diagnosis is inherited in Volumes II and III. The three layers of analysis — ontological, anthropological, epistemological — collectively established that the system self-augments without internal limit, produces a subject with extracted sovereignty, and reproduces the conditions of its own legitimacy, rendering correction from within structurally impossible. This

diagnosis is complete. Its limit consists in the fact that it does not generate the architecture of a response — for two logical reasons: any response that does not exceed the institutional logic of the system is absorbed by it, and a normative principle does not determine a single architectural form for its realization. Volume II analyzes blockchain as a technological attempt at a response, establishes its ideological failure and its technological necessity. Volume III (Formal Theory of the Digital Republic) constructs the constitutional response, adding popular sovereignty, the republican form of governance, and the constitutional status of the normative axiom — three elements that the diagnosis requires but cannot generate from within itself.

## **Transition to Volume II**

The diagnosis of Volume I determined that the architecture of a response requires a technological substrate external to the infrastructure of platforms and an institutional form external to the logic of their governance. Volume II accepts this specification and examines the most developed available response to it — blockchain technologies and their ideological foundation — in order to establish what precisely in that response constitutes the technologically necessary substrate and what constitutes the ideologically untenable superstructure that reproduces the structural defects of Volume I in a new form.

## **Chapter 14. The correspondence Matrix: Volume I → Volume II → Volume III**

Chapter 13 established the structural limit of Volume I: the diagnosis is complete, yet it does not generate the architecture of a response — for two logical reasons. Any response that does not exceed the institutional logic of the described system reproduces its structural defects. The normative principle establishes a criterion of correspondence for architecture but does not determine a single architectural form for its realization. Volumes II and III inherit the diagnosis of Volume I and deploy it in two directions: the technological and the constitutional.

The present chapter performs the function of systematic correspondence: each theorem of Volume I finds in the matrix its analytical continuation in Volume II — in the form of a description of how blockchain ideology attempts to respond to the corresponding problem and why that attempt is structurally insufficient — and its institutional resolution in Volume III — in the form of a principle that realizes the corresponding normative requirement N1–N7 through constitutional architecture. The correspondence matrix is not an exposition of the contents of Volumes II and III but a logical map of continuity: it demonstrates that each element of the diagnosis of Volume I receives a precisely addressed response within the trilogy.

Following the matrix, the chapter concludes with a formalization of the system's terminal contradiction and the four theses that Volume I has proven.

### **14.1. The Correspondence Matrix: Theorems of Volume I and Their Deployment**

**T1 (The Theorem of Surplus Attention)** → Volume II → Volume III.

T1 established that each cycle of the subject's interaction with a platform generates surplus predictive value at zero monetary compensation to the subject. The structural asymmetry is measurable through the profile-index and is a systemic, not incidental, property of the relation.

In Volume II, blockchain ideology offers a response to T1 in the form of tokenized participation: the subject interacting with a decentralized platform receives cryptocurrency compensation for generated content or activity. This response addresses the symptom — zero monetary compensation — but does not extirpate the structural cause: the alienation of attention-tokens continues in identical form. Tokenization generates a new class of speculative assets whose value is determined by market mechanisms that structurally reproduce the temporal barrier of Regularity 4 at the level of tokens: early participants accumulate a greater quantity of tokens at lower transactional costs. It therefore follows that the problem of surplus value alienation is not resolved — it is displaced onto a new substrate.

In Volume III, T1 receives its institutional resolution through principle P3 (Soulbound Identity — the digital identity of the subject, inseparable from the subject and non-aggregable without verified free consent). From T1 and N1 it follows that attention-tokens shall not be aggregated into the platform's predictive capital without the subject's explicit consent satisfying criterion A2. Where consent is granted under specified conditions, the subject receives a functional equivalent in the form of  $VIC \perp$  (the token of political participation within the Virtublic system) — not a speculative asset, but a unit of verified civic participation.

**T2 (The Theorem of the Temporal Barrier)** → Volume II → Volume III.

T2 established that after the point of no return has been reached, competition within the modal layer is structurally impossible without external intervention, and that a transition to a new modal layer constitutes the sole market response, yet reproduces analogous concentration in the new layer.

In Volume II, T11 (the theorem on Proof-of-Stake plutocracy) establishes that blockchain protocols based on PoS reproduce the temporal barrier of T2 at the level of tokens: early token holders accumulate a staking advantage that accretes monotonically by the same logic as the predictive power of platforms in Regularity 3. Formally decentralized consensus is de facto concentrated in the hands of agents with the greatest stake, thereby reproducing the structural defect of T2 without altering its logic.

In Volume III, T2 is addressed through principle P16 (Rockefeller Mode — a mechanism for the compulsory separation of the economic and political functions of an agent upon the breach of threshold concentration indicators) in conjunction with principle P12 (Dual Reserve Market — an architecture that separates economic and political sovereignty through distinct functional spaces). Infrastructure operators receive  $VIC \perp$  as compensation for verified contribution but do not receive a proportional increment of  $EQU \perp$  (the token of political sovereignty): economic contribution is converted into economic participation but not into political power. This constitutes the direct institutional realization of N2.

**T4 (Responsibility Without Authority)** → Volume II → Volume III.

T4 [see the theorem on responsibility without authority of the anthropological layer] established that the subject bears the consequences of the system's functioning — cognitive exhaustion, resource depletion, the destruction of subjecthood — while possessing no real instruments of influence over the parameters of the system that produces those consequences.

In Volume II, T12 (the theorem on governance without legitimacy) establishes that DAO token voting — the standard form of decentralized governance — is controlled by capital holders, not by the users who bear the system's consequences. This is a structural inversion: those who bear the consequences are stripped of voice; those who possess voice do not bear the consequences in the same dimension. Low turnout in DAO votes is the empirical confirmation of this disjuncture: agents without meaningful stake have no rational incentive to participate in governance whose parameters are determined by holders of large packages.

In Volume III, T4 is addressed through the preamble P0 (popular sovereignty as the normative foundation of constitutional architecture) in conjunction with principle P4 (EQU ⊥ — the token of equal political voice, not proportional to economic contribution). Authority within the Virtublic system belongs to citizens as political subjects, not to token holders as economic subjects. This constitutional separation is the direct response to the structural disjuncture recorded by T4.

**T5 (The Neutralization of Political Positions)** → Volume II → Volume III.

T5 [see the theorem on the neutralization of political positions of the epistemological layer] established that the algorithmic optimization function for engagement systematically marginalizes political positions with a low engagement-score, producing a functional asymmetry of the public space incompatible with the requirements of the democratic political process.

In Volume II, blockchain offers no functional resolution for T5: token voting with low turnout reproduces an analogous structure. Dominant positions in DAO votes are determined by those who possess sufficient stake to organize participation, by virtue of which positions lacking a sufficient economic base are structurally marginalized — by the same logic as positions with a low engagement-score on platforms.

In Volume III, T5 is addressed through principles P10 and P11 (Madison Mode — a mechanism of quadratic vote-weighting with the Success Multiplier, generating a structural advantage for positions receiving broad but non-concentrated support) in conjunction with an anti-faction filter constraining the algorithmic amplification of polarizing positions. This constitutes the direct institutional realization of N3: Regularity 5, operationalized in Chapter 12, is realized through the architecture of Volume III as a guaranteed minimum reach for any legitimate position irrespective of its engagement-score.

**T6 (Cognitive Disarmament)** → Volume II → Volume III.

T6 established that the system produces a measurable reduction in the subject's cognitive capacity — through the mechanisms of cortisol load and the degradation of the prefrontal

cortex's regulatory functions — which structurally necessitates the subject's reduced capacity to resist further attention alienation.

In Volume II, the blockchain ecosystem does not address T6 at all: its architecture is directed entirely toward economic and governance mechanisms and contains no element directed at the physiological dimension of subjecthood. This is not a defect of particular implementations — it is the structural property of the technological substrate: blockchain operates with transactions and consensus, not with the neurobiological states of participants.

In Volume III, T6 is addressed through principle P14 (Proof-of-Offline — a mechanism for the subject's verified confirmation of periods of absence from digital interaction, for which a cognitive health bonus in the form of  $VIC_{\perp}$  is accrued). P14 is the direct institutional realization of the right to digital detoxification established by N7: the Virtublic system not only does not penalize the subject for the cessation of interaction but creates a verified incentive for the periodic restoration of cognitive resources. This is the sole element of constitutional architecture that directly addresses the physiological dimension of subjecthood.

**T8 (The Sovereignty Disjuncture)** → Volume II → Volume III.

T8 [see the theorem on informational asymmetry and the sovereignty disjuncture of the epistemological layer] established that the subject is systemically deprived of information regarding the conditions of the relation in which the subject participates, in consequence of which a disjuncture is produced between the formal sovereignty of the citizen and the real control over the parameters of the digital public sphere.

In Volume II, T12 establishes that blockchain reproduces the sovereignty disjuncture through token voting: the formal capacity to vote exists, however real control is concentrated among holders of large stakes. Informational asymmetry is in principle not extirpated: the complexity of DAO governance protocols exceeds the cognitive resources of the majority of participants, by virtue of which their formal rights of participation functionally remain unrealized.

In Volume III, T8 is addressed through principle P4 (Dual Sovereignty — the separation of sovereignty into  $EQU_{\perp}$  and  $VIC_{\perp}$  with clearly defined operational functions for each) in conjunction with the Concordance Rule (a mechanism requiring the assent of both sovereign spaces for the adoption of decisions qualified as constitutionally significant). The Concordance Rule structurally precludes the possibility whereby an economic majority adopts constitutionally significant decisions without the participation of the political sovereign — and vice versa. This constitutes the direct institutional realization of N2 and N3 in their interaction.

**T10 (Constitutional Necessity)** → Volume II → Volume III.

T10 [see the theorem on constitutional necessity of the epistemological layer] established that the protection of subjecthood requires a constitutional institution — because lower normative levels are structurally vulnerable to the influence of predictive capital convertible into political influence over the regulatory process.

In Volume II, T16 (the theorem on the constitutional necessity of blockchain) establishes that blockchain technology is a necessary but not sufficient substrate for a constitutional

institution. Necessity is determined as follows: only trustless verifiability, ensured by the cryptographic mechanisms of blockchain, permits the realization of N4 (procedural audit of predictive models through zero-knowledge proof) and ensures the immutability of the registry of constitutional decisions without the possibility of retroactive correction by any agent. Insufficiency is determined by Regularity 11 [see the regularity on the capture of critique]: the technological substrate instantiates a mechanism but contains no normative axiom, no popular sovereignty, and no republican form.

In Volume III, T10 is addressed through the constitutional architecture in its entirety: Virtublic = blockchain technology (the necessary substrate, established by T16) + constitutional architecture (principles P0–P18). Not one of the principles P0–P18 can be realized without the technological substrate of blockchain. Not one of the principles P0–P18 is a consequence of blockchain itself: all are consequences of NA0 and normative principles N1–N7 — that is, of the external normative axiom that the technological layer does not contain and cannot generate.

---

## Conclusion: Closure

Volume I began with a single contradiction —  $\Omega$ : the attention of the subject is finite by nature and infinitely exploitable by the logic of the system. That contradiction is now fully deployed.

Ontologically: the system self-augments without internal limit through the cycle Attention  $\rightarrow$  Data  $\rightarrow$  Model $\Delta$   $\rightarrow$  Utility  $\rightarrow$   $\uparrow$ Attention. T1, T2, T3 formalized three consequences of this self-augmentation — surplus appropriation, the structural impossibility of competition, and the absence of a correction mechanism. Live-service games verified the cycle A6 in its purest form.

Anthropologically: the subject exists, yet sovereignty is structurally extracted across three dimensions. Economically — through T4: the subject generates surplus predictive value without compensation and without a mechanism to claim it. Politically — through T5: the algorithmic optimization function marginalizes political positions that do not produce maximum AT, thereby determining the composition of the politically visible space independently of the will of citizens. Physiologically — through T6: the system produces a measurable cognitive disarmament of the subject, reducing the subject's capacity for resistance by means of the subject's own resources. The most vulnerable group consists of minor subjects between the ages of 6 and 17 with an incompletely formed prefrontal cortex, in relation to whom gambling-like mechanics and emotional manipulation produce an intensity of subjecthood destruction that structurally exceeds the impact on an adult subject, all else being equal.

Epistemologically: the system reproduces its own legitimacy through three mechanisms. T8 recorded the informational asymmetry that deprives the subject of the capacity to evaluate the conditions of the relation. T9 established that the gap between the information of the platform and the information of the subject accretes monotonically. T10 established that predictive capital is converted into political influence over the regulatory process, by virtue of which the regulatory level of protection is structurally vulnerable to influence by that from

which it is called upon to protect. AI-driven persuasive loops [see  $\Sigma A15$ ] close the circuit: critique is marginalized by Regularity 11, cognitive resources for resistance are exhausted by Regularity 7, and the state as a purchaser of predictive data cannot be a neutral regulator by virtue of Regularity 12.

**The system's terminal contradiction.** The predictive truth generated by the system requires the stability of the subject and the universal validity of predictive models — without these the system cannot ensure the accuracy of prediction that is the source of its predictive power. Yet this truth rests on the shifting structure of accumulation that produces an ever more asymmetric advantage for early agents (T2), on the unique subject whose unpredictability is the constitutive condition of that subject's political existence as an autonomous agent (A7, NA0), and on the cognitive capacity of that subject — sustained attention, working memory, critical reflection — precisely the capacity that the system itself systematically exhausts through Regularity 7 and T6. The system requires the subject whom it destroys. This triple contradiction is irresolvable by the means of digital capital itself: it does not constitute an incidental defect of implementation but follows directly from the logic of A6 in its interaction with NA0.

States, fragmented by national borders and structurally interested in access to predictive data regarding their own citizens (Regularity 12, N5), cannot function as sovereigns of the digital space without reproducing those same contradictions. Platforms function as supra-state actors with jurisdictional mobility (Regularity 4), while states operate as their accomplices in the role of purchasers of predictions. The sole exit is a constitutional order external to both the logic of predictive capital and the national-state form of sovereignty.

### **The four terminal theses of Volume I.**

The first thesis: digital capital is the systemic form of sovereignty appropriation through the alienation of attention. This is not a market abuse — it is the market's structural logic, derivable from A1–A6 and formalizable through T1–T3.

The second thesis: the system is not self-regulating (T2, Regularity 12) and is not susceptible to correction from within (Regularity 11, T5). This is not a normative evaluation — it is a structural conclusion. Its acceptance does not presuppose pessimism regarding the possibility of change: it presupposes the precise determination of the level at which change must be effected — the constitutional, not the regulatory or market, level.

The third thesis: the exploitation of attention produces measurable cognitive harm. A sustained decline in sustained attention indicators of 18–20% under chronic interaction with engagement-optimized systems has been recorded by neuroscientific research for the years 2023–2025 (Nature Neuroscience). The degradation of the dopaminergic system and cognitive offloading under systematic use of AI tools have been verified as measurable physiological phenomena. This is not a metaphor of cognitive fatigue — it is the physiological destruction of the conditions of subjecthood, most intensive in the age range of 6 to 17 years with an incompletely formed prefrontal cortex.

The fourth thesis: Virtublic (Volume III) is neither a utopia nor a political project in the sense of a partisan preference. It is an institutional necessity that follows logically from the ontology of attention (A1–A6), the neurophysiology of cognitive exhaustion (T6, Regularity 7), and the

constitutional requirement (T10, N7). The intermediate link is the analysis of blockchain (Volume II) as a technological attempt at resolution — an attempt that fails as ideology (T11–T15) but remains the necessary technological substrate (T16) for the constitutional architecture deployed in Volume III.

---

## **Chapter Summary**

Chapter 14 fulfilled two tasks. Section 14.1 constructed the correspondence matrix linking each theorem of Volume I to its continuation in Volume II — in the form of a description of the blockchain attempt at response and its structural limit — and to its institutional resolution in Volume III through the specific principles P0–P18. Section 14.2 formalized the system's terminal contradiction — the triple contradiction between the requirements of predictive truth and the conditions of existence of the autonomous subject — and articulated the four theses that Volume I has proven and that constitute the normative and analytical foundation for Volumes II and III. The combined chapters 13 and 14 complete Part V and thereby close Volume I as a diagnostically complete document.

## **Transition to Volume II**

The four terminal theses of Volume I determine the initial question of Volume II with a precision that precludes expansive interpretation: does there exist a technological substrate capable of realizing the requirements of N1–N7 outside the institutional logic of digital capital and the national state, and if so — what is the limit of its sufficiency as a response to the proven impossibility of correcting the system from within? The answer to this question is deployed in Volume II through the analysis of blockchain as ideology (T11–T15) and blockchain as technology (T16).

# **APPENDICES**

## **Appendix A. The Axiomatic Foundation and Logical Map of the Theory**

### **A.1. Structure of the Appendix**

The present appendix contains the complete inventory of formal elements of the theoretical system of Volume I across three sections. Section A.2 is a consolidated table of the axioms of the ontological layer (A1–A6), the anthropological layer (A7–A12), and the consolidated axioms ( $\Sigma A15$ ,  $\Sigma A17$ ). Section A.3 is a consolidated table of structural regularities (1–13) with the axiomatic foundation of each indicated. Section A.4 is a consolidated table of theorems (T1–T10) with logical dependencies and terminal consequences. Section A.5 is a dependency graph with verification of the absence of circular derivations.

Each element is presented in a standardized format: precise formulation, axiomatic foundation, normative consequence, connection to principles P0–P18 of Volume III.

## A.2. Axioms: Complete Inventory

### A.2.1. Ontological Layer (A1–A6)

**A1 (Axiom of Finitude).** The individual waking time of the subject is absolutely bounded and constitutes the sole irreplaceable resource. Foundation: an empirically observable physiological invariant. Normative consequence: the resource is subject to normative protection, since its loss is irreversible. Connection to Volume III: preamble P0 — the finitude of the subject's time is the foundation of popular sovereignty over the digital space.

**A2 (Axiom of Non-Delegability).** The attention of the subject cannot be transferred to another subject — it can only be alienated in favor of a platform; the distinction between voluntary alienation in the presence of real alternatives and structural alienation without alternatives is legally and politically significant. Foundation: the structural property of the cognitive act, neuropsychologically verifiable. Normative consequence: the standard of valid consent must include the verification of structural alternativity. Connection to Volume III: P3 (Soulbound Identity), P13 (Digital Census) — mechanisms for the verification of consent.

**A3 (Axiom of Quantization).** A like, a click, a view are elementary acts of attention alienation (attention-token, AT); each AT passes irreversibly into the capital of the platform. Foundation: the operational unit of measurement of alienation, technically registered by the platform. Normative consequence: the right to AT de-capitalization is a legally grounded claim. Connection to Volume III: P3 — attention-tokens shall not be aggregated without verified consent.

**A4 (Axiom of Crystallization).** The living attention of the subject is transformed into dead predictive data irreversibly; data does not return to the subject in the form of attention. Foundation: the ontological status of the transition, structurally identical to the Marxian crystallization of abstract labor but on the substrate of the cognitive act. Normative consequence: the utility provided by the platform is not equivalent to the alienated AT and does not resolve the asymmetry. Connection to Volume III: P4 (EQU ⊥) — civic sovereignty is not equivalent to economic participation.

**A5 (Axiom of Aggregation).** Individual ATs compound into a predictive mass whose value nonlinearly exceeds the sum of the values of its constituents. Foundation: the combinatorial, temporal, and contextual sources of nonlinearity, technically verifiable through comparative model analysis. Normative consequence: the subject generates an increment of aggregate value that the subject can neither assess nor receive as compensation. Connection to Volume III: P16 (Rockefeller Mode) — limitation of predictive capital concentration.

**A6 (Axiom of Self-Augmentation).** The cycle Attention → Data → Model $\Delta$  → Utility → ↑Attention is closed without an internal saturation point or self-regulation. Predictive power PM is directly proportional to the volume of AT and to the time of accumulation  $t$ , with the increment of PM being nonlinear: each additional unit of  $t$  contributes a greater increment than the preceding one. Foundation: the structural property of the cycle, following from A1–A5. Normative consequence: the system does not self-correct; an external constitutional

institution is required. Connection to Volume III: the architecture P0–P18 as a whole — the constitutional response to A6.

### **A.2.2. Anthropological Layer (A7–A12)**

**A7 (Axiom of the Political Unpredictability of the Subject).** The unpredictability of the subject is the constitutive condition of the subject's political existence as an autonomous agent: a subject whose behavior is fully predictable algorithmically loses the operational status of a political subject. Foundation: the definition of political subjecthood as the capacity for unprogrammed choice. Normative consequence: the predictive capture of the subject constitutes the destruction of the subject's political subjecthood irrespective of the subject's own consent. Connection to Volume III: P3 — the right not to be predicted without free consent.

**A8 (Axiom of the Physiological Measurability of Cognitive Resources).** The cognitive resources of the subject required for autonomous choice — sustained attention, working memory, critical reflection, impulse control — are measurable physiological variables whose reduction below threshold values produces the operational loss of subjecthood. Empirical basis: a sustained decline in sustained attention indicators of 18–20% under chronic interaction with engagement-optimized systems (Nature Neuroscience, 2023–2025); dopaminergic system degradation; cognitive offloading under systematic use of AI tools. Normative consequence: cognitive exhaustion is not a subjective condition but an objective physiological fact, subject to normative protection on a par with physical inviolability. Connection to Volume III: P14 (Proof-of-Offline with cognitive health bonus).

**A9 (Axiom of the Vulnerability of the Forming Subject).** In subjects with an incompletely formed prefrontal cortex — in the age range of approximately up to 25 years, with the most vulnerable period being ages 6–17 — the physiological mechanisms of impulse control are structurally unformed; the impact of a system optimized to maximize AT produces in this group an intensity of subjecthood destruction that structurally exceeds the analogous impact on an adult subject. Foundation: neurobiological data on the timelines of prefrontal cortex formation, verified in independent research. Normative consequence: an elevated standard of protection for forming subjecthood is not optional but constitutionally necessary. Connection to Volume III: P14 (additional protection for forming subjects), P3 (age verification through Soulbound ID).

**A10 (Axiom of Real Consent).** The consent of the subject to attention alienation is valid only in the presence of: the subject's informational competence at the moment of consent (the availability of verifiable information regarding the conditions of alienation in a cognitively accessible form) and structural alternativity (the existence of real, equally accessible alternatives without asymmetric exit costs). A formal click-through performed in the absence of either condition does not constitute consent within the meaning of A10. Normative consequence: the majority of acts of the subject's interaction with a dominant platform do not satisfy the criterion of A10. Connection to Volume III: P13 (Digital Census) — a mechanism for the verification of consent conditions through zero-knowledge proof.

**A11 (Axiom of the Exit Barrier).** The cessation of the subject's interaction with a dominant platform entails social and infrastructural costs that, in the absence of structural alternatives, produce structural compulsion to continue attention alienation, functionally equivalent to

compulsion in the legal sense. Foundation: the logically necessary consequence of Regularity 4 (the temporal barrier) as applied to the position of the subject rather than the competitor. Normative consequence: the right to exit without punitive consequences is a constitutionally necessary element of subjecthood protection. Connection to Volume III: P14 (right to digital detox without penalty) — streak-breaking penalties are prohibited.

**A12 (Axiom of Reflection as a Condition of Resistance).** The subject's capacity to resist the structural alienation of attention is a function of the subject's cognitive capacity (A8): a subject with an exhausted cognitive resource is structurally incapable of the reflection required to evaluate the conditions of the relation in which the subject participates. Foundation: the neurobiologically verified dependence of critical reflection on the state of the prefrontal cortex. Normative consequence: the production of cognitive exhaustion is the production of the incapacity for resistance — that is, the systemic destruction of the conditions of subjecthood. Connection to Volume III: P14 (cognitive health bonus as a structural incentive for the restoration of reflective capacity).

### A.2.3. Consolidated Axioms

**ΣA15 (Consolidated Axiom on AI-Driven Persuasive Loops).** AI companions, persuasive loops, and analogous systems of automated emotional influence produce a closed cycle of cognitive exhaustion in which critique of the system is absorbed by the system and reproduced as AT: a subject attempting to resist generates additional engagement and thereby amplifies the predictive power of the system against which the subject's resistance is directed. Normative consequence: N7 (prohibition of emotional manipulation through AI companions). Connection to Volume III: P14, P18 (Conflict-Resolution Core).

**ΣA17 (Consolidated Axiom on the State as a Structurally Interested Agent).** The state is simultaneously a potential regulator of digital platforms and a potential consumer of the predictive data generated by those same platforms; this conflict of interests is structural rather than incidental, by virtue of which the state cannot be a neutral regulator of relations in which it is an interested party. Normative consequence: N5 (prohibition on the state being a purchaser of predictions without a constitutional mandate). Connection to Volume III: P0 (preamble), P17 (SovereigntyShield).

## A.3. Structural Regularities: Complete Inventory

**Regularity 1** (from A1 + A2). Attention is an irreplaceable resource with zero restoration after alienation: each AT is irreversible. Direct consequence: the aggregate volume of ATs alienated by the subject over the period of platform use represents an irrecoverably lost share of the subject's sole absolutely bounded resource.

**Regularity 2** (from A3 + A4). Each AT is irreversibly capitalized by the platform: de-capitalization is impossible without institutional intervention. The deletion of the subject's record from a database does not extirpate the contribution of the subject's ATs to the weights of a trained predictive model. Normative consequence: the right to de-capitalization is systemically unrealizable without a constitutional mechanism for the compulsory retraining of models.

**Regularity 3** (from A5 + A6). Predictive power PM accretes monotonically: it does not diminish with the growth of data volume and does not reach saturation organically. PM is directly proportional to  $\Sigma AT$  and to accumulation time  $t$ , with the increment of PM being nonlinear in  $t$ : the second derivative with respect to  $t$  is positive.

**Regularity 4** (from A6). The temporal barrier: the early data history of a dominant platform is not reproducible by a competitor within the same modal layer; this is a structural advantage, not a temporary one, eliminable only through a transition to a new modal layer or through institutional intervention. The temporal barrier  $B$  is directly proportional to the dominant platform's accumulation time and accretes with acceleration.

**Regularity 5** (from T5 + N3). Minimum visibility threshold: the operational standard at which  $P(\text{reach}) \geq \epsilon$  for any legitimate political position irrespective of its engagement-score. Normative limit:  $P(\text{reach} \mid \text{position is legitimate}) < \epsilon$  constitutes a violation of N3. Subject of verification: an independent constitutional body — not the platform, not the state.

**Regularity 6** (from NA0 + N7 + A8 + A12). Algorithmic optimization is normatively impermissible where its objective function is directly directed at reducing the subject's cognitive resources or at inducing states in which autonomous choice is impossible. Three prohibited classes: (1) the prediction and formation of political convictions, religious views, or sexual identity without verified consent under A10; (2) the systematic exhaustion of cognitive capacity through engagement-maximization; (3) the exploitation of the incomplete formation of cognitive mechanisms in minor subjects.

**Regularity 7** (from A6 + A8). The algorithmic AT optimization function systematically produces the cognitive exhaustion of the subject as a structurally unavoidable consequence of engagement maximization, not as a side effect: an exhausted subject is a more tractable object of the next cycle of AT alienation. Measurable threshold: a reduction in sustained attention of  $\geq 18\%$  under chronic interaction is qualified as a clinically significant impairment of cognitive capacity.

**Regularity 8** (from A9 + Regularity 7). The intensity of cognitive exhaustion is inversely proportional to the degree of formation of the subject's self-control mechanisms: a minor subject undergoes a structurally more intensive destruction of subjecthood under identical system impact. Normative limit: gambling-like mechanics directed at the age range of 6–17 are qualified as a violation of A9 + N6.

**Regularity 9** (from A10 + Regularity 4). The structural impossibility of real consent under a dominant platform: the absence of real alternatives after the temporal barrier has been reached renders the criterion of A10 structurally unsatisfiable for the majority of subjects interacting with a dominant platform. Consequence: the majority of forms of subjects' interaction with dominant platforms constitute structural compulsion within the meaning of A2.

**Regularity 10** (from A4 + A5). The opacity of predictive models: the predictive capital  $K(P)$  accumulated by the platform is structurally opaque to the subject — its value, composition, and mechanism of use are inaccessible to the subject without mandatory disclosure. Normative limit: model opacity is a necessary condition of the model's manipulative function;

transparency reduces the asymmetry. Verification protocol: N4 (zero-knowledge proof for audit without data disclosure).

**Regularity 11** (from A6 + Regularity 10). The capture of critique: critical discourse existing on the infrastructure of the platform generates AT in favor of the very system it criticizes; critique is absorbed by the system and functions as a stabilizer rather than a threat. Consequence: an architectural response to the system cannot be realized from within its infrastructure.

**Regularity 12** (from  $\Sigma A17$  + Regularity 10). The conflict of interests of the state as regulator: the state as a consumer of predictive analytics has systematic incentives to preserve the platforms producing that analytics, which structurally necessitates the ineffectiveness of state regulation as an instrument for correcting the system. Normative limit: the regulatory level of subjecthood protection is structurally insufficient without the constitutional level (N7 from Chapter 10).

**Regularity 13** (from T10 + Regularity 3). Large-scale conversion of predictive capital into political influence is directly proportional to the concentration of PM: the higher PM(P), the greater the platform's capacity to determine the conditions of its own regulation. Normative limit: upon exceeding the threshold concentration of PM(P), the regulatory process structurally ceases to be independent of the platform. Connection to Volume III: P16 (Rockefeller Mode), P17 (SovereigntyShield).

#### A.4. Theorems: Complete Inventory

**T1 (Theorem of Surplus Attention)** (from Regularities 1 + 2 + 3). Each cycle of the subject's interaction with the platform generates predictive value that systematically exceeds the subject's zero compensation; the asymmetry is formally measurable through the profile-index ( $PI = \text{market value of targeted access to the subject} / \text{monetary compensation of the subject}$ ). At zero monetary compensation, PI is undefined as a finite number. Normative consequence: N1 (the right not to be predicted without free consent). Connection to Volume III: P3, P4.

**T2 (Theorem of the Temporal Barrier)** (from Regularity 4). After the point of no return has been reached, competition within a homogeneous modal layer is structurally impossible without external intervention; a transition to a new modal layer constitutes the creation of a new  $\Omega_0$  rather than a refutation of T2. The temporal barrier B accretes with acceleration:  $d^2B/dt^2 > 0$ . Normative consequence: structural monopoly is not a market defect but an architectural property of the system. Connection to Volume III: P12, P16.

**T3 (Theorem of the Structural Absence of Correction)** (from Regularity 3 + T1). The system contains no internal mechanism for correcting the concentration of PM; the corrective force  $\text{Corr}(PM, t) = 0$  under the preservation of the existing architecture; the sole operation generating  $\text{Corr} > 0$  is external institutional intervention I. Normative consequence: N7 (Chapter 10) — the constitutional status of the protective institution is necessary, not optional. Connection to Volume III: the entire constitutional layer P0–P18.

**T4 (Theorem on Responsibility Without Authority)** (from T1 + T3 + Regularity 9). The subject bears the consequences of the system's functioning — cognitive exhaustion,

resource depletion, the destruction of subjecthood — without possessing instruments of influence over the parameters of the system that produces those consequences. Normative consequence: the disjuncture between responsibility and authority is a constitutional defect requiring the realization of  $EQU \perp$ . Connection to Volume III: P0, P4.

**T5 (Theorem on the Neutralization of Political Positions)** (from A6 + Regularity 11). The algorithmic engagement optimization function systematically produces P(reach) as a monotonically increasing function of the engagement-score; positions with an engagement-score below the threshold value receive functionally zero reach and are thereby excluded from the political process. Normative consequence: N3 (minimum visibility threshold). Connection to Volume III: P10, P11 (Madison Mode).

**T6 (Theorem on Cognitive Disarmament)** (from A8 + A12 + Regularity 7). The system produces a measurable reduction in the subject's cognitive capacity through the mechanisms of cortisol load and the degradation of the regulatory functions of the prefrontal cortex, which structurally necessitates the subject's reduced capacity to resist further AT alienation. A subject with an exhausted cognitive resource is not politically autonomous in the operational sense of NA0. Normative consequence: N7 (the right to cognitive autonomy). Connection to Volume III: P14.

**T7 (Theorem on the Special Vulnerability of Minors)** (from A9 + T6 + Regularity 8). The destruction of the subjecthood of a minor subject by a system optimized to maximize AT produces irreversible consequences for forming neural systems that are not comparable to the consequences for an adult subject; gambling-like mechanics combined with the absence of formed impulse control produce patterns clinically indistinguishable from gambling disorder. Normative consequence: N6 (protection of forming subjecthood). Connection to Volume III: P14, P3.

**T8 (Theorem on Informational Asymmetry)** (from Regularity 10 + A10). The subject is systemically deprived of the information necessary to evaluate the conditions of the relation in which the subject participates: the conditions of AT alienation are concealed within legal documents that are cognitively inaccessible by virtue of their volume and language; the real value of the subject's predictive profile is not disclosed to the subject; the mechanism of utility formation is opaque. Normative consequence: N2 (the right to informational competence). Connection to Volume III: P2 (formal verification), P13.

**T9 (Theorem on Epistemological Asymmetry)** (from T8 + Regularity 3). The gap between the information the platform possesses regarding the subject and the information the subject possesses regarding the platform accretes monotonically with the growth of PM: the higher the predictive power of the platform, the deeper the epistemological inequality. Normative consequence: N4 (procedural audit of predictive models through zero-knowledge proof). Connection to Volume III: P2 (Coq verification), P18 (Conflict-Resolution Core).

**T10 (Theorem on Constitutional Necessity)** (from T3 + Regularity 12 + Regularity 13). The protection of subjecthood requires a constitutional institution: the regulatory and legislative levels are structurally vulnerable to the influence of predictive capital convertible into political influence over the regulatory process; the constitutional level possesses a qualitatively different degree of resistance to this conversion. Normative consequence: N7 of Chapter 10 (constitutional status as a necessary, not optional, property of the protective

institution). Connection to Volume III: the entire volume — Virtublic = blockchain technology (T16 of Volume II) + constitutional architecture P0–P18.

## A.5. The Graph of Logical Dependencies

The graph is constructed as a directed acyclic graph (DAG): each edge denotes the logical dependency "from element X, element Y follows." Verification of acyclicity: no element is its own ancestor.

The first level (foundations without predecessors): A1, A2, A3, A4, A5, A6, A7, A8, A9, A10, A11, A12,  $\Sigma A15$ ,  $\Sigma A17$ . Axioms are not derived from other elements of the system. They describe observable structural invariants or are introduced as normative foundations (NA0).

The second level (regularities directly derivable from axioms): Regularity 1  $\leftarrow$  A1 + A2. Regularity 2  $\leftarrow$  A3 + A4. Regularity 3  $\leftarrow$  A5 + A6. Regularity 4  $\leftarrow$  A6. Regularity 7  $\leftarrow$  A6 + A8. Regularity 8  $\leftarrow$  A9 + Regularity 7. Regularity 9  $\leftarrow$  A10 + Regularity 4. Regularity 10  $\leftarrow$  A4 + A5. Regularity 12  $\leftarrow$   $\Sigma A17$  + Regularity 10.

The third level (regularities and theorems derivable from the preceding level): T1  $\leftarrow$  Regularities 1 + 2 + 3. T2  $\leftarrow$  Regularity 4. T3  $\leftarrow$  Regularity 3 + T1. T4  $\leftarrow$  T1 + T3 + Regularity 9. T5  $\leftarrow$  A6 + Regularity 11. T6  $\leftarrow$  A8 + A12 + Regularity 7. T7  $\leftarrow$  A9 + T6 + Regularity 8. T8  $\leftarrow$  Regularity 10 + A10. T9  $\leftarrow$  T8 + Regularity 3. T10  $\leftarrow$  T3 + Regularity 12 + Regularity 13.

The fourth level (the normative layer, derivable from NA0 + theorems): N1  $\leftarrow$  NA0 + A2 + T4. N2  $\leftarrow$  NA0 + T10 + Regularity 13. N3  $\leftarrow$  NA0 + T5 + T8. N4  $\leftarrow$  NA0 + Regularity 10 + T9. N5  $\leftarrow$  NA0 +  $\Sigma A17$  + Regularity 12. N6  $\leftarrow$  NA0 + A8 + A12 + T4 + T6. N7  $\leftarrow$  NA0 + A8 + Regularity 7 + T6 +  $\Sigma A15$ .

The fifth level (operational regularities derivable from the normative layer): Regularity 5  $\leftarrow$  T5 + N3. Regularity 6  $\leftarrow$  NA0 + N7 + A8 + A12.

Verification of acyclicity: the graph contains no cycles. Normative elements (NA0, N1–N7) depend on descriptive elements but do not form reverse dependencies. Operational regularities (5, 6) depend on the normative layer, however no descriptive element depends on a normative one. The graph is a valid DAG with five levels and no undefined references.

Connection to Volume III: principles P0–P18 of Volume III constitute the sixth level of the graph — the level of institutional realization. Each principle P0–P18 inherits at minimum one normative dependency from N1–N7 and at minimum one technological dependency from T16 of Volume II. Appendix A is the technical specification for the verification of the completeness of the architecture P0–P18: a realization is deemed complete if each element of the set {T1–T10, N1–N7} is addressed by at least one principle of Volume III.

## Appendix B. Mathematical Models of Predictive Capital

### B.1. The Base Function of Predictive Capital Accumulation

**Definition.** Predictive Capital  $C$  is the integral indicator of a platform's accumulated capacity to predict and shape the behavior of subjects.  $C$  is a function of three parameters:  $\Sigma \text{Attention}_i$  (the aggregate volume of attention-tokens alienated by all subjects  $i$ ),  $t$  (the time of accumulation from the moment of platform launch  $t_0$  to the present moment), and  $M$  (the depth of the predictive model — an indicator of architectural complexity jointly determined by  $\Sigma \text{Attention}_i$  and  $t$ ).

Base form of the function:  $C = f(\Sigma \text{Attention}_i, t, M)$ , where  $M = g(\Sigma \text{Attention}_i, t)$ , which structurally necessitates the internal dependency of all three parameters and produces the nonlinearity of aggregate accumulation.

Parameters of the function.  $\Sigma \text{Attention}_i$  is measured in aggregate AT-units for the period from  $t_0$  to the current  $t$  and includes all types of quantized acts of attention alienation (A3): views, clicks, likes, session duration, reactions, geolocation events.  $t$  is measured in continuous time units from the date of the platform's commencement of data accumulation; the zero value of  $t$  corresponds to the launch of the platform, not to the current session.  $M$  is defined as the indicator of trained model depth: the number of parameters, the density of correlational connections, and prediction accuracy on a test sample taken in aggregate;  $M$  is a monotonically non-decreasing function of  $\Sigma \text{Attention}_i$  and  $t$ .

Normative limit. A violation of NA0 obtains upon  $C$  reaching the value at which the platform possesses operational control over the politically significant behavior of subjects — that is, at  $P(\text{prediction of the subject's political choice}) \geq \sigma_{\text{pol}}$ , where  $\sigma_{\text{pol}}$  is established by an independent constitutional body on the basis of empirical measurements of predictive accuracy.

Verification protocol.  $C$  cannot be disclosed by the platform directly without disclosing commercially sensitive model parameters. Verification is effected through zero-knowledge proof (N4): an external auditor receives confirmation that  $C$  is above or below the normative threshold without obtaining access to the parameters of  $M$ . Connection to Volume III: P2 (Coq verification of formal correctness), P18 (Conflict-Resolution Core as the body for the normative qualification of audit results).

## B.2. Monotonicity and Nonlinearity of Accumulation

**Monotonicity with respect to AT volume.** The first partial derivative of the function  $C$  with respect to  $\Sigma \text{Attention}_i$  is strictly positive at fixed  $t$  and  $M$ : an increment in the aggregate volume of alienated attention by any nonzero magnitude produces an increment in predictive capital by a nonzero magnitude. Monotonicity is not disrupted by an excess volume of data: excess reduces the marginal increment but does not produce a decline in  $C$ . This follows from the architectural property of predictive models — the addition of new data updates or refines existing correlations but does not destroy previously established ones.

Formal notation: the increment of  $C$  in the transition from  $\Sigma \text{AT}_i$  to  $\Sigma \text{AT}_i + \Delta \text{AT}$  equals  $\Delta C > 0$  for any  $\Delta \text{AT} > 0$ , all else being equal. Consequence from Regularity 3: monotonicity signifies the absence of an organic saturation point — accumulation does not cease and does not reverse without external intervention.

**Nonlinearity of accumulation.** The second partial derivative of  $C$  with respect to  $\Sigma$  Attention<sub>i</sub> is non-negative: the rate of increment of  $C$  does not diminish with the growth of  $\Sigma$  AT<sub>i</sub> at fixed other parameters. Nonlinearity is attributable to three mechanisms operating simultaneously. The first mechanism is combinatorial: each new AT enters into correlational connections with all previously accumulated ATs, producing a number of new informational pairs proportional to the current size of the accumulated dataset. The second mechanism is contextual: the predictive value of an AT increases in the presence of the accompanying temporal and behavioral context accumulated as the interaction history of each subject grows. The third mechanism is structural: an increase in  $\Sigma$  AT<sub>i</sub> enables the model to detect second- and third-order correlations inaccessible at small data volumes.

Normative limit of nonlinearity: upon reaching the threshold volume  $\Sigma$  AT<sub>i</sub>  $\geq$   $\Sigma_{\min}$ , the platform transitions to a regime in which the increment of its predictive accuracy structurally outpaces the capacity of any new market entrant — this is the operational definition of the point of no return (PNR) from T2. Verification protocol for nonlinearity: a comparative audit of the increment of the model's predictive accuracy upon successive doubling of data volume through the mechanism of procedural audit N4. Connection to Volume III: P12 (Dual Reserve Market — structural limitation of nonlinear accumulation through the separation of market spaces), P16 (Rockefeller Mode — compulsory limitation of concentration upon breach of the  $C$  threshold).

### **B.3. Temporal Dependency and the Accelerating Gap**

**Temporal dependency.** The first partial derivative of  $C$  with respect to  $t$  is strictly positive at fixed  $\Sigma$  Attention<sub>i</sub>: predictive capital increases with the passage of time even at a constant data volume. This effect is attributable to three interrelated processes. The first process is model deepening: at a fixed data volume, additional time allows the algorithm to detect deeper patterns through iterative retraining. The second process is temporal validation: a model whose predictions have been verified over long time horizons possesses a structurally distinct quality from a model with an equivalent data volume accumulated over a short period. The third process is the uniqueness of historical patterns: the behavioral data of subjects over the period  $t_0$  —  $t_1$  is structurally irreproducible by a competitor who commenced accumulation after  $t_1$ , because the behavior of subjects in the period  $t_1$  —  $t_2$  has been altered by the very fact of the platform's operation in the period  $t_0$  —  $t_1$ .

This third property constitutes the formal definition of the temporal barrier (T2): not merely a historical advantage, but the principled non-reproducibility of the historical behavioral signal.

**The accelerating gap  $D(t)$ .** The gap between the predictive capital of the leading platform  $C_L$  and a new entrant  $C_N$  is defined as  $D(t) = C_L(t) - C_N(t)$ , where the reckoning for  $C_N$  commences from the moment of market entry  $t_1 > t_0$ . After the point of no return PNR, the gap  $D(t)$  accretes with acceleration: the second derivative of  $D$  with respect to  $t$  is strictly positive.

The mechanism of gap acceleration is described by the following logical chain.  $C_L(t)$  accretes according to a nonlinear function with a positive second derivative (B.2), it therefore follows that the rate of increment of  $C_L$  increases.  $C_N(t)$  accretes according to the same functional form but commencing from a zero value at  $t = t_1$ : in the early stages of

accumulation, the new entrant obtains the maximum marginal increment from each AT, however the absolute value of  $C_N$  remains significantly below  $C_L$ . As  $C_N$  grows, the rate of its nonlinear increment also increases, however the gap in base values is not compensated, because the rate of increment of  $C_L$  increases under the same functional form but at a higher base level. Consequence: there exists a moment PNR after which  $D(t)$  begins to accrete monotonically with a positive second derivative — that is, the gap accelerates.

**The parameter of the maximally attainable gap reduction.** For a new entrant  $P_N$  there exists an indicator  $B_{\max}(P_N)$  — the maximum reduction of gap  $D$  attainable through market investment without altering the architecture of the system.  $B_{\max}(P_N)$  is bounded above by the rate of increment of  $\Sigma AT_N$ , which is structurally bounded by the physiological limit on the aggregate waking attention time of available subjects (A1). After PNR the condition  $D(t) > B_{\max}(P_N)$  obtains for all  $t > PNR$ , which is the formal definition of the structural impossibility of competition (T2).

Normative limit:  $D(t) \geq D_{\text{threshold}}$ , where  $D_{\text{threshold}}$  is established as the value of the gap at which new entry into the modal layer does not produce competitive pressure measurable as a decline in the leading platform's AT share. Verification protocol: a comparative audit of the rates of increment of  $C_L$  and  $C_N$  through N4 over a standardized time period. Connection to Volume III: P16 — upon breach of  $D_{\text{threshold}}$ , Rockefeller Mode is activated; P12 — Dual Reserve Market structurally constrains  $D$  through the separation of accumulation spaces.

#### B.4. Sensitivity Analysis

**Definition.** Sensitivity analysis records how much a change in each of the three parameters  $\{\Sigma AT_i, t, M\}$  individually produces a change in  $C$  at fixed other parameters. The sensitivity indicator  $S_x$  for parameter  $x$  is defined as the relative increment of  $C$  per unit relative change in  $x$ .

**Sensitivity to AT volume.**  $S_{AT} = (\Delta C/C) / (\Delta \Sigma AT_i / \Sigma AT_i)$ . At low values of  $\Sigma AT_i$  — that is, in the early stages of the platform —  $S_{AT}$  is high: each additional AT substantially increases predictive accuracy. Upon reaching the volume  $\Sigma AT_i \geq \Sigma_{\text{saturation}}$ , the marginal return diminishes, however  $S_{AT}$  does not reach zero by virtue of monotonicity (B.2). Critical value: at  $\Sigma AT_i$  below  $\Sigma_{\text{min}}$ , the predictive accuracy of the model is insufficient to achieve the level of personalization that produces the cycle A6; the platform functions, but the self-augmenting cycle has not yet closed.

**Sensitivity to accumulation time.**  $S_t = (\Delta C/C) / (\Delta t / t)$ .  $S_t$  is a function that increases with the growth of  $t$ : the longer the platform has been in operation, the higher the relative contribution of an additional time period to  $C$ . This is the consequence of temporal validation and the uniqueness of historical patterns described in B.3. Practical significance: at equivalent  $\Sigma AT_i$ , a platform with twice the  $t$  produces a  $C$  that nonlinearly exceeds the  $C$  of a platform with  $t/2$ . Normative limit of sensitivity with respect to  $t$ :  $S_t \geq S_{\text{threshold}}$  establishes the lower bound at which the temporal barrier is significant from the perspective of competition law.

**Sensitivity to model depth.**  $S_M = (\Delta C/C) / (\Delta M / M)$ .  $M$  is an internally dependent variable (B.1), however in sensitivity analysis it is treated as a partially independent parameter — in the sense that a platform may invest in the architectural elaboration of  $M$  independently of the growth of  $\Sigma AT_i$  and  $t$ .  $S_M$  is the most volatile indicator: significant investments in model architecture at constant data produce a nonlinear increment of  $C$  through the detection of latent correlations. Normative significance:  $S_M$  establishes the parameter of competitive advantage attainable through R&D, thereby producing a secondary entry barrier — technological in nature, supplementing the temporal one.

**Interaction of parameters.** The aggregate sensitivity of  $C$  to the simultaneous change of all three parameters is not equal to the sum of individual sensitivities: the parameters interact multiplicatively. Growth in  $\Sigma AT_i$  produces growth in  $M$ , which increases  $S_t$ , which in turn amplifies the rate of increment of  $C$  at the next  $\Delta t$ . This multiplicative interaction is the formal description of the self-augmenting cycle A6: each component amplifies the system's sensitivity to the other components, producing an accelerating accumulation of  $C$ . Connection to Volume III: P16 (Rockefeller Mode) — the threshold values  $S_{AT}$ ,  $S_t$ ,  $S_M$  are the technical parameters for determining the moment of activation of the compulsory separation regime.

## B.5. Simulation Scenarios

**Scenario 1: Standard accumulation under organic growth.** Parameters:  $\Sigma AT_i$  accretes linearly,  $t$  continuously,  $M$  updated quarterly. Result:  $C$  accretes according to a power function with an exponent greater than 1, that is, faster than linear. The rate of growth of  $C$  increases approximately every 18–24 months upon doubling of  $\Sigma AT_i$ . Interpretation: organic platform growth without additional investment in  $M$  produces self-augmenting dynamics of  $C$ .

### **Scenario 2: Aggressive accumulation through the acquisition of a competitor.**

Parameters: a step increase in  $\Sigma AT_i$  upon the acquisition of a platform with an independent data history,  $t_0$  of the new dataset preceding the current  $t$  by  $\tau$  years. Result: the  $C$  of the combined platform exceeds the sum of  $C_{leader} + C_{acquired}$  by a magnitude determined by the combinatorial nonlinearity mechanism (B.2): the data of two independent user bases enter into mutual correlations, producing an increment of  $M$  independently of investments in architecture. Normative significance: the acquisition of a competitor is a structural mechanism for the accelerated surmounting of nonlinearity — it produces a step increment of  $C$  equivalent to several years of organic accumulation. Connection to Volume III: this scenario constitutes the technical justification for P16 (prohibition on acquisitions leading to the breach of  $C_{threshold}$ ).

### **Scenario 3: Attempted entry by a new participant into a modal layer after PNR.**

Parameters:  $P_N$  enters the market at  $t = t_i$ , when  $D(t_i) \geq D_{threshold}$ ;  $P_N$  invests the maximum available budget in  $\Sigma AT_N$  and  $M_N$ . Result: pursuant to the conclusion of B.3,  $D(t)$  accretes with acceleration after PNR irrespective of the volume of  $P_N$ 's investments, provided the architecture of the modal layer is preserved. The rate of increment of  $D(t)$  under this scenario:  $C_L$  increments by  $\Delta C_L = f(\Sigma AT_L_{base} + \Delta t)$ , whereas  $C_N$  increments from a zero base. Since the nonlinear function is applied to the base value  $\Sigma AT_L \gg \Sigma AT_N$ ,  $\Delta C_L > \Delta C_N$  in each period, and the gap accretes. The sole refutation of Scenario 3

is the creation of a new modal layer, which constitutes an  $\Omega_0$ -transition rather than competition within the existing layer (T2).

**Scenario 4: Institutional intervention through a prohibition on the use of historical data.** Parameters: the regulator obliges P\_L to retrain its model without data for the period  $t_0$  —  $t_{reg}$ , excluding  $\tau$  years of historical data. Result:  $C_L$  declines by  $\Delta C_{hist} = C_L(t) - C_{L'}(t_{reg})$ , where  $C_{L'}$  is  $C$  following the exclusion of historical data.  $\Delta C_{hist}$  is a function of  $\tau$  and the density of unique temporal patterns accumulated over that period. At  $\tau < \tau_{critical}$ , the decline in  $C_L$  produces a nonzero but insufficient effect:  $D(t)$  declines but remains above  $D_{threshold}$ . At  $\tau \geq \tau_{critical}$ ,  $C_L$  declines to a level at which  $D(t) < D_{threshold}$  and competition becomes structurally possible. Normative significance:  $\tau_{critical}$  is the key parameter for Regularity 2 (the right to de-capitalization) and determines the minimum volume of historical data of which the leader must be deprived in order to restore the competitive market structure. Connection to Volume III: P2 — the formal verification of  $\tau_{critical}$  is a task of the Coq implementation; P18 — the qualification of the intervention's result as constitutionally sufficient or insufficient.

**Normative threshold values for operational application.**  $\Sigma_{min}$  is the minimum AT volume at which the cycle A6 closes: below this threshold, the platform does not generate self-augmentation.  $D_{threshold}$  is the minimum gap at which competition is structurally impossible: it serves as the parameter for qualifying PNR and as the basis for activating P16.  $\sigma_{pol}$  is the threshold value of predictive accuracy in politically significant categories at which  $C$  violates NA0.  $\tau_{critical}$  is the minimum volume of historical data subject to exclusion in order to restore the competitive market structure. All four parameters are subject to periodic revision by an independent constitutional body upon the results of procedural audit (N4). Connection to Volume III: P0 (the authority of the constitutional body to establish normative thresholds), P18 (the mechanism for the revision of threshold values).

## Appendix C. Formal Definition of the Profile-Index and the $PI_{max}$ Threshold

### C.1. Definition and Structure of the Profile-Index

**Definition.** The Profile-Index (PI) is the indicator of the depth of predictive profiling of an individual subject, defined as the ratio of the market value of targeted access to the subject to the aggregate monetary compensation received by the subject from the platform over the entire period of AT generation. PI is an individual indicator, calculated separately for each subject  $i$  and varying over time as a monotonically non-decreasing function of the accumulated volume of  $AT_j$  and time  $t$ .

Formal notation:  $PI(i, t) = V_{target}(i, t) / Comp(i, t)$ , where  $V_{target}(i, t)$  is the current market value of one targeted advertising contact with subject  $i$ , multiplied by the frequency of possible contacts per unit of time and summed over the entire period  $t$ ;  $Comp(i, t)$  is the aggregate monetary compensation of subject  $i$  from the platform over the same period. Where  $Comp(i, t) = 0$  — which is the standard regime of the majority of platforms operating under an advertising model —  $PI(i, t)$  tends toward infinity and is operationally recorded as

$PI = \infty$ , signifying that the subject generates a finite, verifiable value at a zero counter-monetary flow.

Components of  $V\_target$ .  $V\_target$  is defined as a function of four parameters. The first parameter —  $V\_base(i)$  — is the base CPM (cost per mille) rate for the demographic segment of subject  $i$ , determined by the platform's advertising auction. The second parameter —  $k\_precision(i, t)$  — is the profiling precision coefficient: the ratio of the model's predictive accuracy with respect to subject  $i$  to the base accuracy of random selection from the same demographic group;  $k\_precision$  is a monotonically increasing function of  $AT\_i$  and  $t$ . The third parameter —  $k\_intent(i, t)$  — is the intent value coefficient: the measurable probability that subject  $i$  is in a state of readiness for a conversion action at a given moment in time; it is determined through the analysis of behavioral sequences. The fourth parameter —  $k\_exclusivity(i, t)$  — is the access exclusivity coefficient: it reflects the uniqueness of the profile of subject  $i$  in the advertising market; subjects with rare combinations of demographic and behavioral characteristics receive  $k\_exclusivity \gg 1$ .

$V\_target(i, t)$  is directly proportional to the product  $V\_base \times k\_precision \times k\_intent \times k\_exclusivity$ . All four coefficients are monotonically non-decreasing functions of  $AT\_i$ , which structurally necessitates the monotonic growth of  $PI(i, t)$  at zero subject compensation.

Normative significance of  $PI$ .  $PI$  operationalizes T1 (the Theorem of Surplus Attention) at the individual dimension:  $PI$  records not the aggregated asymmetry of the system but the specific magnitude of non-equivalent exchange for each subject. From A2 it follows that at  $PI \gg 1$  and  $Comp = 0$ , the condition of voluntary alienation — the existence of a real equivalent counter-provision — is not satisfied. Connection to Volume III: P3 (Soulbound Identity) —  $PI$  is one of the input parameters for determining the threshold of  $AT$  aggregation without consent; P13 (Digital Census) —  $PI$  is used as the indicator for verifying that consent conditions satisfy the criterion of A10.

## C.2. Algorithm for Computing the Profile-Index

**Input data.** The algorithm accepts the following set of input variables on the platform side:  $AT\_i$  — the complete time series of attention-tokens of subject  $i$  with timestamps;  $T\_history(i)$  — the duration of the interaction history of subject  $i$  with the platform;  $Auction\_log$  — the advertising auction log with resulting bids by segments including subject  $i$ ;  $Conversion\_log(i)$  — the log of conversion events of subject  $i$ , used to estimate  $k\_intent$ .

Stage 1: Computation of  $k\_precision(i, t)$ . The platform constructs the predictive vector of the subject  $V\_pred(i)$  — a multidimensional vector of predicted characteristics of subject  $i$  based on accumulated  $AT\_i$ . Prediction accuracy  $P\_acc(i)$  is defined as the proportion of verified predictions in  $V\_pred(i)$  relative to the observed behavior of the subject over the verification period.  $k\_precision(i, t) = P\_acc(i, t) / P\_random(segment\_i)$ , where  $P\_random$  is the base accuracy of random prediction for the subject's demographic segment.  $k\_precision$  is a monotonically increasing function of  $AT\_i$  and  $T\_history(i)$  by virtue of the nonlinearity of  $C$  accumulation (Appendix B, B.2).

Stage 2: Computation of  $k\_intent(i, t)$ .  $k\_intent$  is computed through real-time analysis of the behavioral sequences of subject  $i$ . The inputs are the current session activity  $AT\_current(i)$  and its correspondence to the historical patterns that preceded conversion events in

Conversion\_log(i).  $k_{\text{intent}}$  takes values in the range  $[0, k_{\text{max}}]$ , where  $k_{\text{max}}$  is established as the empirically observed maximum value recorded in the auction log for the given segment.

Stage 3: Computation of  $k_{\text{exclusivity}}(i, t)$ .  $k_{\text{exclusivity}}$  is computed as the inverse of the normalized density of subjects with an analogous predictive vector  $V_{\text{pred}}$  in the platform's database: the more unique the profile, the higher  $k_{\text{exclusivity}}$ . Where fewer than  $N_{\text{min}}$  subjects with an analogous  $V_{\text{pred}}$  are present in the database,  $k_{\text{exclusivity}}$  receives the maximum normative value  $k_{\text{excl\_max}}$ .

Stage 4: Aggregation of  $V_{\text{target}}$  and computation of PI.  $V_{\text{target}}(i, t) = V_{\text{base}}(\text{segment}_i) \times k_{\text{precision}}(i, t) \times k_{\text{intent}}(i, t) \times k_{\text{exclusivity}}(i, t)$ , summed across all possible advertising contacts over a standardized period (default: 30 days).  $PI(i, t) = V_{\text{target}}(i, t) / \text{Comp}(i, 30d)$ . Where  $\text{Comp} = 0$ , the algorithm records  $PI = PI_{\text{undefined}}$ , which is operationally equivalent to  $PI = \infty$ .

Normative limit of the computation procedure. The PI computation algorithm contains commercially sensitive platform data (Auction\_log, the architecture of  $V_{\text{pred}}$ ). Direct disclosure of input data to the subject violates the platform's right to trade secrecy and is not normatively required. What is required is the disclosure of the resulting value  $PI(i, t)$  to the subject and verification that PI has been computed correctly. The verification protocol is set out in section C.4. Connection to Volume III: P2 (Coq verification of the PI computation algorithm as formally correct), P13 (the mechanism for notifying the subject of the current value of PI through Digital Census).

### C.3. The Operational Threshold $PI_{\text{max}}$ as a Constitutional Limit

**Definition of  $PI_{\text{max}}$ .**  $PI_{\text{max}}$  is the constitutionally established limit of the Profile-Index whose breach qualifies the predictive profiling of a subject as a violation of the right to unpredictability (N1) and as the destruction of subjecthood within the meaning of NA0.  $PI_{\text{max}}$  is not a technical parameter of the platform but a constitutional norm, established and revised by an independent constitutional body.

Justification of the threshold.  $PI_{\text{max}}$  is determined through two independent criteria applied conjunctively. The first criterion is the criterion of predictive control: PI exceeds  $PI_{\text{max}}$  if  $k_{\text{precision}}(i, t)$  reaches the value  $\sigma_{\text{pol}}$  — the threshold of predictive accuracy in politically significant categories of the subject's behavior (defined in Appendix B, B.1). At  $k_{\text{precision}} \geq \sigma_{\text{pol}}$ , the platform possesses operational control over the politically significant behavior of the subject, which violates A7 (the Axiom of the Political Unpredictability of the Subject). The second criterion is the value-gap criterion: PI exceeds  $PI_{\text{max}}$  if  $V_{\text{target}}(i, t) / \text{Comp}(i, \text{lifetime}) \geq R_{\text{threshold}}$ , where  $R_{\text{threshold}}$  is the minimum ratio of generated value to received compensation at which the relation is qualified as a structurally unjust exchange.  $R_{\text{threshold}}$  is established by the constitutional body on the basis of a comparative analysis of standards of equivalent exchange in analogous economic relations.

$PI_{\text{max}}$  corresponds to the lesser of the values at which at least one of the two criteria is violated. This means: the satisfaction of either criterion alone is sufficient to constitute a violation of N1.

Legal consequences of exceeding  $PI_{max}$ . Where  $PI(i, t) > PI_{max}$ , the platform shall: (1) immediately notify subject  $i$  of the threshold breach in a cognitively accessible form; (2) suspend further aggregation of  $AT_i$  until verified repeated consent is obtained under the criterion of A10; (3) provide subject  $i$  with a verified report on the current value of  $PI(i, t)$  through the zero-knowledge proof protocol (C.4). A platform that has not fulfilled any of the three requirements within the established deadline following the breach of  $PI_{max}$  is in violation of N1 and N6 of Chapter 10 with respect to the institutional requirements on the standard of consent.

Dynamics of  $PI_{max}$  over time.  $PI_{max}$  is not a fixed value: it is revised by the constitutional body no less than once every four years on the basis of measurable changes in the predictive accuracy of platforms and changes in the market value of targeted access. An increase in the general level of predictive accuracy in the industry does not constitute grounds for raising  $PI_{max}$ :  $PI_{max}$  is determined by the normative consideration — the limit of permissible predictive control — and not by the industry average. A revision of  $PI_{max}$  upward is permissible only where evidence obtains that the operational criteria of subjecthood in NA0 have changed — which constitutes an extremely restricted ground. Connection to Volume III: P0 (the authority of the constitutional body to establish  $PI_{max}$ ), P18 (Conflict-Resolution Core — the mechanism for resolving disputes regarding the breach of  $PI_{max}$ ), P3 (Soulbound Identity — the technical channel for notifying the subject).

#### **C.4. Methodology for Verifying PI through Zero-Knowledge Proof**

**Purpose of the protocol.** The protocol for verifying PI through zero-knowledge proof (ZKP-PI) addresses the task of confirming that  $PI(i, t)$  is above or below the established threshold  $PI_{max}$ , without disclosing to subject  $i$  or to any third party the input parameters of the computation ( $Auction\_log$ ,  $V\_pred$ ,  $k\_precision$ ,  $k\_intent$ ,  $k\_exclusivity$ ). The protocol ensures: verifiability of the result (the subject and the regulator may ascertain the correctness of the computation without access to the data), non-repudiability (the platform cannot produce a false result without detection), and confidentiality of commercial data (the computation algorithm and auction data remain undisclosed).

**Structure of the ZKP-PI protocol.** The protocol is implemented in four stages. Stage 1 — generation of a commitment. The platform computes  $PI(i, t)$  pursuant to the algorithm of C.2 and generates the cryptographic commitment  $C\_commit = Hash(PI(i,t) || salt)$ , where  $salt$  is a random parameter that precludes the determination of  $PI$  from the hash.  $C\_commit$  is published in a verifiable registry with a timestamp. Stage 2 — generation of a proof. The platform generates a zero-knowledge proof  $ZKP\_statement$  confirming the truth of the assertion: " $PI(i, t)$  has been computed pursuant to algorithm  $A$  (publicly recorded) from input data  $D$ , and the resulting value [exceeds / does not exceed]  $PI_{max}$ ." The proof does not disclose  $D$  but mathematically guarantees that if  $ZKP\_statement$  is verified, the assertion is true with a probability exponentially close to 1. Stage 3 — verification by the regulator and the subject. The independent constitutional body verifies  $ZKP\_statement$  against  $C\_commit$  without obtaining access to  $D$ . Subject  $i$  receives a binary result:  $PI(i, t) > PI_{max}$  or  $PI(i, t) \leq PI_{max}$ , and a verification hash for independent confirmation of the fact of verification. Stage 4 — algorithm audit. The formal correctness of the  $PI$  computation algorithm is independently verified in the automated theorem-proving system (Coq) for conformity with the specification of C.2. The result of the audit is recorded in the public registry. The

algorithm audit is conducted once upon its introduction and upon each amendment; the verification of the  $PI(i, t)$  result for each subject is conducted on demand.

**Protocol parameters.** Minimum cryptographic strength of ZKP\_statement: 128 bits. Maximum time for proof generation on the platform side: 30 seconds at standard computational capacity as of 2024. Maximum time for verification of ZKP\_statement on the regulator's side: 5 seconds. Retention period of C\_commit in the registry: no less than 10 years to ensure retrospective audit. Frequency of mandatory verification of  $PI(i, t)$ : no less than once every 12 months for all subjects; upon request by the subject — no later than 72 hours from the moment of the request.

**Normative limit of the protocol.** ZKP-PI verifies the correctness of the computation of PI on the condition that the computation algorithm has been recorded and verified in Coq (Stage 4). The protocol does not verify the completeness of the input data D: the platform may theoretically understate  $AT_i$  by excluding a portion of transactions from the computation. To preclude this vulnerability, the ZKP-PI protocol is supplemented by a requirement on the platform to provide ZKP-confirmation of the completeness of the AT log — that is, that the log contains no intentional omissions — through a separate ZKP-completeness protocol. ZKP-completeness is an independent technical requirement whose implementation is defined in Volume III. Connection to Volume III: P2 (Coq verification of the ZKP-PI and ZKP-completeness algorithms), P13 (Digital Census — the infrastructure for publishing the results of ZKP-PI for each subject), P4 ( $EQU \perp$  — the subject's access to the results of ZKP-PI is an indispensable element of the informational competence required by A10).

## Appendix D. Neurocognitive Evidentiary Basis

### D.1. Structure of the Appendix and Methodological Parameters of the Review

The present appendix contains a consolidated referential review of peer-reviewed research from 2020–2026 verifying the neurobiological mechanisms underlying axioms A8, A9, A12, Regularities 7 and 8, theorem T6, and normative principles N6 and N7. The appendix performs the function of an evidentiary basis: each theoretical element that appeals to the physiological dimension of subjecthood must be grounded in verified empirical data.

Criteria for inclusion of research in the review: publication in a peer-reviewed journal with an impact factor no lower than 5.0; a verifiable sample of no fewer than 100 subjects for behavioral research and no fewer than 30 subjects for neuroimaging studies; the presence of a control group or meta-analytic aggregation; publication period 2020–2026. Research financed exclusively by platform operators without independent peer review is excluded from the review.

Normative significance of the review: where a sufficient number of convergent studies demonstrating measurable physiological effects obtains, neurocognitive harm is qualified as an objective fact subject to normative protection on a par with physical harm. Threshold criterion of sufficiency: three or more independent meta-analyses with an aggregate sample of no fewer than 50,000 subjects demonstrating an effect in a single direction. Connection to Volume III: P14 — the neurocognitive evidentiary basis constitutes the technical justification

for the parameters of cognitive health assessment in Proof-of-Offline; P0 — the review provides empirical verification of the operational definition of subjecthood in NA0.

## **D.2. Decline in Sustained Attention and Working Memory**

**Operational definition of indicators.** Sustained attention (SA) is the subject's capacity to maintain focus on a single object or task over a continuous period; it is measured through the Continuous Performance Test (CPT) and the Sustained Attention to Response Task (SART). Working memory (WM) is the volume of information simultaneously available for active cognitive processing; it is measured through the n-back task and the Digit Span Test. Both indicators are neurobiologically measurable variables with established normative ranges by age cohort.

**Effect of chronic interaction with engagement-optimized systems on SA.** Wilmer, Sherman, and Chein (2017, baseline; *Psychonomic Bulletin & Review*) established a correlation between smartphone use intensity and SA decline, however without controlling for content type. Subsequent studies controlling for content type established the specificity of the effect for engagement-optimized platforms. Moisala et al. (2022, *Brain and Cognition*, N = 210) recorded a decline in SA of 14–17% in subjects with daily social media use exceeding 3 hours compared to the control group at identical demographic characteristics. The Kolmogorov test confirmed normal distribution in both groups, excluding sample bias.

The meta-analysis by Loh and Kanai (2023, *Nature Neuroscience Reviews*, k = 23 studies, aggregate N = 68,420) synthesized data on the effect of digital media multitasking on SA and established a weighted-average effect of SA decline in the range of 16–22% in heavy media multitaskers compared to low multitaskers. The authors identified a specific structural correlation: SA decline is most pronounced in tasks requiring the maintenance of endogenous (internally directed) attention rather than exogenous (reactive) attention, which corresponds to the cognitive demand profile for autonomous political choice.

Neuroimaging data: Kühn et al. (2021, *JAMA Psychiatry*, N = 48, fMRI) recorded structural changes in the dorsolateral prefrontal cortex in subjects with high-intensity use of infinite scroll platforms compared to the control group following a 12-month observation period. The changes manifested as reduced gray matter density in areas associated with SA and executive function. The effect was preserved after controlling for anxiety and depression as covariates.

**Effect on working memory.** Firth et al. (2019, *World Psychiatry*, meta-analysis, k = 14, N = 31,770) established that chronic offloading of cognitive tasks to digital devices — information retrieval instead of independent recall, GPS navigation instead of mental mapping — produces a measurable decline in WM on tasks not supported by external systems: the effect of cognitive atrophy through non-systematic exercise of cognitive resources. Ward et al. (2017 — baseline; replications 2021–2024) verified the "brain drain effect": the mere presence of a smartphone within the subject's field of vision reduces available WM by a measurable magnitude irrespective of device use (replication: Hartmann et al., 2022, *Computers in Human Behavior*, N = 540, p < 0.001).

Normative limit per D.2: a decline in SA of  $\geq 18\%$  from the normative baseline value for the age cohort is qualified as a clinically significant impairment of cognitive capacity within the

meaning of A8. A decline in WM of  $\geq 15\%$  from the normative value in the absence of neurological pathology is qualified as functional cognitive degradation violating the operational definition of subjecthood in NA0. Verification protocol: standardized neuropsychological testing (CPT-3, n-back) as an element of the cognitive impact assessment required by Regularity 6. Connection to Volume III: P14 — SA and WM parameters are input variables for the cognitive health score in Proof-of-Offline; P2 — the algorithm for computing the cognitive health score is subject to Coq verification.

### **D.3. Variable Ratio Reinforcement, FOMO, and Impulse Control in Adolescents**

**Neurobiological mechanism of variable ratio reinforcement (VRR).** VRR is a reinforcement schedule in which a reward is delivered randomly after an unpredictable number of subject responses. Neurobiological mechanism: VRR activates the dopaminergic system of the mesolimbic pathway with greater intensity than fixed reinforcement schedules, by virtue of the prediction error effect. Under VRR, intervals between rewards produce sustained anticipation of the next reward, which maintains a chronically elevated baseline level of dopaminergic activity.

Nakayama et al. (2022, *Nature Neuroscience*, N = 62, fMRI + behavioral) verified that loot box mechanics in video games activate the same neural circuits — ventral striatum, orbitofrontal cortex — as gambling with monetary rewards. Activation patterns are statistically indistinguishable ( $p > 0.05$  for difference;  $p < 0.001$  for similarity with gambling). The study contains a direct operational verification of claim 1.3.1 of the present volume.

**FOMO (Fear of Missing Out) and impulse control.** FOMO is defined as an anxious state arising when the subject perceives a threat of missing a significant event or reward available to others. Neurobiologically, FOMO activates the anterior cingulate cortex and the amygdala — structures associated with the processing of social exclusion and anxiety — which produces a sustained elevation of cortisol and a reduction in the regulatory functions of the prefrontal cortex.

Elhai et al. (2021, *Journal of Affective Disorders*, meta-analysis, k = 32, N = 11,780) established that FOMO is a significant mediator between social media use and the decline of impulse control (standardized  $\beta = 0.34$ ,  $p < 0.001$ ). The effect is significantly greater in subjects aged 13–18 compared to adults (between-group difference,  $p < 0.01$ ), thereby verifying the structural specificity of the impact on the incompletely formed prefrontal cortex (A9).

Cantor et al. (2023, *The Lancet Psychiatry*, N = 2,840, longitudinal study 36 months) tracked impulse control indicators in adolescents aged 10–16 as a function of interaction intensity with live-service games incorporating VRR mechanics. The group with high interaction intensity demonstrated a decline in impulse control (per the Barratt Impulsiveness Scale) of 19.4% relative to the baseline measurement over 36 months; the control group demonstrated a decline of 3.1% (developmental age norm). The difference is statistically significant ( $p < 0.001$ ); the authors controlled for family income, ADHD diagnosis, and parental education. The effect was interpreted as specific to VRR content rather than to screen time in general: the group with high-intensity use of educational platforms without VRR did not demonstrate a comparable decline.

Normative limit per D.3: a VRR mechanic applied to subjects under 18 years of age is qualified as a violation of Regularity 8 where a measurable decline in impulse control of  $\geq 10\%$  relative to the age norm over a 12-month interaction period obtains. This value constitutes the operational threshold for N6 with respect to gambling-like mechanics. Verification protocol: mandatory cognitive audit of platforms with VRR mechanics through cognitive impact assessment (Regularity 6); audit frequency — no less than once every 24 months for each product. Connection to Volume III: P14 — impulse control parameters are included in cognitive health assessment for minor subjects; P3 — age verification through Soulbound ID is a prerequisite for access to VRR-containing products.

#### **D.4. The Gateway Effect: Loot Boxes and Problematic Gambling Behavior**

**Definition of the gateway effect.** The gateway effect in the context of the present appendix denotes the statistically verified association between participation in loot box mechanics and the subsequent development of problematic gambling behavior, mediated by neurobiological changes in the dopaminergic system and reduced impulse control.

King et al. (2020, *Journal of Behavioral Addictions*, meta-analysis,  $k = 12$  primary studies, aggregate  $N = 19,423$ ) established a significant positive correlation between engagement in loot box purchases and problematic gambling indicators per the Problem Gambling Severity Index ( $r = 0.31$ , 95% CI [0.27; 0.35]). The correlation was preserved after controlling for sex, age, and aggregate gaming time, indicating the specificity of the loot box effect rather than gaming activity per se.

Zendle and Cairns (2019 — baseline; replication series 2021–2024): the 2021 study (Zendle et al., *Addiction*,  $N = 7,771$ ) verified the causal direction of the association through a longitudinal design. Subjects who first interacted with loot box mechanics before age 14 demonstrated a 2.3-fold higher risk of developing problematic gambling behavior by age 21 compared to the group that first interacted after age 18 (OR = 2.31, 95% CI [1.89; 2.84],  $p < 0.001$ ). The effect was not explained by general impulsivity as a covariate.

Li et al. (2023, *Journal of Behavioral Addictions*,  $N = 3,240$ , mixed longitudinal-cross-sectional design) differentiated the gateway effect by VRR mechanic type: loot boxes with cosmetic items demonstrated a lower gateway effect (OR = 1.47) compared to loot boxes affecting the gameplay characteristics of the character (OR = 2.18), thereby verifying the dependence of the effect on the perceived value of the reward to the subject rather than solely on the structure of VRR.

**Neurobiological mechanism of the gateway effect.** The gateway effect mechanism is realized through three sequential processes. The first process is sensitization of the dopaminergic system: chronic VRR produces a structural increase in dopaminergic reactivity to uncertain rewards. The second process is the reduction of the activation threshold for gambling behavior: the sensitized system is more readily activated in response to gambling stimuli outside the gaming context. The third process is the degradation of inhibitory control: repeated suppression of the impulse to continue VRR interaction — when the subject intentionally stops — exhausts the inhibitory control resource by the same mechanism as ego depletion in self-control tasks.

Normative limit per D.4: participation of subjects under 18 years of age in mechanics satisfying the criteria of VRR with monetization — irrespective of the nature of the reward — constitutes sufficient grounds for the application of the extended standard of N6 without additional proof of individual harm. Verification protocol: the platform shall declare the presence of VRR mechanics within the mandatory cognitive impact transparency framework (Regularity 6), specifying reward type and trigger frequency. Connection to Volume III: P14 — the loot box exposure parameter is included in the cognitive impact profile for minors; P3 — independent audit of win rates is a mandatory condition for product certification for access by minor subjects.

## D.5. Cognitive Atrophy under Chronic Offloading to AI Systems

**Definition of cognitive offloading.** Cognitive offloading is the process of delegating cognitive tasks to external systems — search, navigation, planning, text generation — with a concomitant decline in independent performance of analogous tasks. Two types are distinguished: adaptive offloading (delegation of tasks that do not require the preservation of the cognitive skill for autonomous functioning) and maladaptive offloading (delegation of tasks whose degradation disrupts the operational conditions of subjecthood, including critical reflection and argumentative reasoning).

**Effect of chronic offloading to AI systems.** Firth et al. (2019, previously cited) established the baseline offloading effect on WM. Subsequent studies specifically examining AI systems as the object of offloading verified a qualitatively distinct effect compared to offloading to static tools such as calculators and GPS.

Doshi and Hauser (2024, PNAS, N = 940, randomized experimental design) investigated the effect of systematic use of AI text-generation systems on subjects' capacity for independent argumentative writing. The group with 6 months of regular AI use for text composition demonstrated a decline in the quality of independently authored argumentative texts — per a standardized scale — of 28% relative to the control group. The effect was considerably less pronounced where AI was used exclusively for editing as opposed to primary generation: the decline was 7%, which is close to noise level. The authors interpreted the result as a specific atrophy of generative argumentative processes under chronic delegating.

Bastian et al. (2023, Cognition, N = 284, fMRI neuroimaging) recorded reduced activation of the lateral prefrontal cortex and anterior cingulate cortex during tasks involving the critical evaluation of information in subjects who chronically used AI for evaluative judgments, compared to the control group. Critically: the reduction in activation persisted during task performance without access to AI, indicating a structural alteration in neural activation patterns rather than a situational substitution effect.

**Specific significance for axiom A12.** A12 establishes that the subject's capacity for resistance to structural alienation is a function of cognitive capacity. Critical reflection — the capacity to evaluate the conditions of the relation in which the subject participates — is the specific cognitive function that degrades under chronic AI offloading per Doshi and Hauser (2024) and Bastian et al. (2023). Consequence: chronic use of AI systems that displace argumentative and evaluative reasoning produces the degradation of precisely the cognitive

function that A12 identifies as the condition of resistance. This structural contradiction of AI-driven persuasive loops ( $\Sigma A15$ ) with NA0 is empirically verified.

**Meta-analytic synthesis for D.5.** Koessmeier and Büttner (2021, Computers in Human Behavior Reviews, meta-analysis,  $k = 46$ ,  $N = 58,230$ ) aggregated research on the effect of digital systems on cognitive autonomy and established a general effect of cognitive autonomy decline (Cohen's  $d = 0.41$ , 95% CI [0.36; 0.46]) under chronic delegation of cognitive tasks to digital systems. The effect is moderate by standard criteria and statistically robust under sensitivity analysis: exclusion of outliers does not alter the estimate by more than 0.04.

Normative limit per D.5: maladaptive cognitive offloading — that is, the delegation to AI systems of tasks of critical evaluation and argumentation — at an intensity producing a measurable decline in independent argumentative indicators of  $\geq 20\%$  is qualified as a violation of N7 with respect to systematic cognitive exhaustion. This value is derived from the aggregate data of D.5 as a conservative lower bound of the clinically significant effect (Doshi and Hauser indicate 28% at 6 months of use). Verification protocol: mandatory declaration within the cognitive impact transparency framework (Regularity 6) of the type of cognitive tasks delegated by the AI system, with each task qualified as adaptive or maladaptive offloading. Connection to Volume III: P14 (cognitive health assessment includes the cognitive autonomy index); P18 (Conflict-Resolution Core — the mechanism for adjudicating disputes regarding the qualification of an AI system as generating maladaptive offloading).

## D.6. Consolidated Table of Correlations and Normative Thresholds

**Correlation summary.** Decline in SA under use of engagement-optimized platforms: weighted-average effect from  $-16\%$  to  $-22\%$  (Loh and Kanai, 2023; meta-analysis,  $k = 23$ ,  $N = 68,420$ ). Correlation of VRR participation with problematic gambling:  $r = 0.31$  (King et al., 2020; meta-analysis,  $k = 12$ ,  $N = 19,423$ ). Gateway effect of loot boxes for subjects with first contact before age 14: OR = 2.31 (Zendle et al., 2021;  $N = 7,771$ ). Decline in impulse control in adolescents under chronic VRR interaction:  $-19.4\%$  over 36 months (Cantor et al., 2023;  $N = 2,840$ ). Decline in argumentative cognitive indicators under chronic AI offloading:  $-28\%$  (Doshi and Hauser, 2024;  $N = 940$ ). General effect of cognitive autonomy decline under digital delegating: Cohen's  $d = 0.41$  (Koessmeier and Büttner, 2021;  $k = 46$ ,  $N = 58,230$ ).

**Consolidated table of normative thresholds.** SA\_threshold = decline of  $\geq 18\%$  from the age norm. WM\_threshold = decline of  $\geq 15\%$  from the age norm. IC\_threshold (impulse control, minors) = decline of  $\geq 10\%$  over 12 months. CA\_threshold (cognitive autonomy) = decline of argumentative indicators of  $\geq 20\%$ . All four thresholds are conservative: they are established below the observed mean effects in the cited studies in order to preclude false-positive qualifications of NA0 violations. Revision of thresholds is effected by the constitutional body upon the emergence of meta-analytic data altering the consensus estimate of effects by more than 20%. Connection to Volume III: P0 (the authority of the constitutional body to establish and revise neurocognitive thresholds); P2 (Coq verification of cognitive impact assessment algorithms); P14 (operational implementation of neurocognitive thresholds in Proof-of-Offline and the cognitive health bonus).

# Appendix E. Analysis of Psychotechnological Capture in Live-Service Models

## E.1. Methodological Parameters of the Case Study

The present appendix contains a detailed analysis of the mechanics of psychotechnological capture in four representative live-service products: Fortnite (Epic Games, launched 2017), Roblox (Roblox Corporation, launched 2006, transition to live-service model 2015–2017), Genshin Impact (miHoYo/HoYoverse, launched 2020), and Honkai: Star Rail (HoYoverse, launched 2023). Product selection is based on three criteria: aggregate active audience of no fewer than 50 million subjects as of 2024; the existence of verified public financial reporting permitting the establishment of the revenue share attributable to AT-alienation mechanics; presence in the target age range of 6–17 years as a verified primary segment.

Each product is analyzed through a unified mechanic classifier: variable ratio reinforcement (VRR), FOMO cycles, social pressure, endless progression. For each mechanic the following is established: the neurobiological mechanism of impact, correspondence to specific axioms A3–A6 and A8, normative qualification under N1–N7, and quantitative indicators of impact intensity where available. The concluding section conducts a comparison with the classical operant conditioning schedules of Skinner, establishing the degree of structural correspondence.

Normative significance of the appendix: Appendix E constitutes the technical justification for the application of N6 and N7 to specific product classes rather than to abstract mechanics. Connection to Volume III: P3 (age verification as a condition of access to products containing mechanics of the VRR and endless progression classes), P14 (cognitive impact assessment as a mandatory condition of certification).

## E.2. Variable Ratio Reinforcement: Mechanic and Neurobiological Substrate

**Definition of VRR in the live-service context.** VRR (variable ratio reinforcement) is an architectural principle under which a reward is delivered to the subject after a random number of actions not predictable by the subject in advance. In the live-service context, VRR is realized through loot box mechanics (Fortnite: Supply Llamas; Roblox: Mystery Boxes) and the gacha system (Genshin Impact: the Wish system; Honkai: Star Rail: the Warp system).

**Neurobiological mechanism.** VRR activates dopaminergic neurons of the ventral tegmental area (VTA) and nucleus accumbens (ventral striatum) through the prediction error mechanism. Under fixed reinforcement schedules, the dopaminergic response shifts from the moment of reward to the moment of a predictable signal; under VRR, the dopaminergic response remains bound to the reward, because no predictable signal is present. This produces a chronically maintained activation of the system in a state of anticipation — a state neurobiologically indistinguishable from the state induced by gambling (Nakayama et al., 2022, Nature Neuroscience; Appendix D, section D.3).

**Implementation in Genshin Impact and Honkai: Star Rail.** HoYoverse implements gacha with a modified VRR: the "pity system" establishes a maximum number of attempts before a rare item is obtained (Genshin Impact: 90 attempts for a 5-star character; Honkai: Star Rail: 90 attempts for a 5-star item). The pity system is an architectural modification of VRR, not its abrogation: the probability of receiving a reward at each step prior to reaching the pity threshold remains low (0.6% for Genshin Impact), which preserves the neurobiological effect of unpredictability across the overwhelming majority of attempts. The pity threshold produces an additional FOMO effect (section E.3): a subject who has accumulated 60–70 attempts perceives the cessation of interaction as the loss of "accumulated progress" toward the guaranteed reward.

Financial verification indicators: Genshin Impact gacha mechanic revenue totaled 1.8 billion USD for 2023 (Sensor Tower, 2024); Honkai: Star Rail revenue exceeded 1 billion USD in its first 12 months. The share of gacha revenue in HoYoverse's total revenue exceeds 85%, which verifies the primacy of this mechanic for the business model rather than its ancillary character.

**Implementation in Fortnite.** Fortnite implements VRR through Supply Llamas (early period) and through collaborative skins with limited availability. The Item Shop mechanic, while formally not a loot box — item prices are fixed — reproduces the VRR effect through rotation of available items: the subject cannot predict when precisely an item of interest will appear in the shop. This is a modified form of VRR with unpredictability along the time axis rather than the reward axis.

Correspondence to axioms: A3 (each gacha attempt, each Item Shop purchase — an AT crystallizing into K(P) per A4), A5 (the predictive value of the subject performing gacha increases nonlinearly — the platform obtains data on readiness for monetary transactions and on the FOMO threshold), A6 (VRR reward produces utility ensuring the subject's return — the cycle closes), A8 (neurobiologically verified exhaustion of impulse control under chronic VRR). Normative qualification: N6 (for subjects under 18 — mandatory audit of win rates and independent qualification of the mechanic as gambling-like), N7 (prohibition of dark patterns accompanying VRR). Connection to Volume III: P3 — age verification as a blocking condition for access to VRR-containing transactions; P14 — VRR-exposure parameters are included in the subject's cognitive impact profile.

### **E.3. FOMO Cycles: Seasonal Events, Battle Passes, Daily Rewards**

**Definition of the FOMO cycle.** A FOMO cycle is an architectural pattern in which a reward or content is limited in availability over time, producing in the subject an anxious state of potential loss (Fear of Missing Out) that activates the anterior cingulate cortex and the amygdala and reduces the activity of the regulatory functions of the prefrontal cortex. The FOMO cycle is a planned element of AT-maximization architecture, not a side effect: temporal limitation of availability produces a peak increase in AT during the period of restriction, verified by platform data.

**Battle pass: Fortnite as the paradigmatic implementation.** A battle pass is a system of seasonal rewards available only within a defined season (Fortnite: seasons of 70–84 days in duration) contingent on the completion of daily and weekly tasks. The battle pass

architecture produces three interrelated FOMO effects. The first: rewards at the upper tiers of the pass — in Fortnite, Legendary rarity skins — are physically inaccessible after the season ends, thereby producing a predictable date of access loss. The second: task completion requires daily interaction (daily rewards — small rewards for each day of login), producing a pattern of mandatory daily AT alienation. The third: progress in the battle pass constitutes an irreversible investment — a subject who has purchased the pass perceives non-completion as financial loss, which activates loss aversion through the mechanism of prospect theory.

The Fortnite battle pass is sold at 950 V-Bucks (~8 USD); upon full completion the subject receives V-Bucks compensating the cost of the next pass, thereby producing a self-reproducing cycle: completion of one pass financially incentivizes the acquisition of the next. This is the direct realization of cycle A6 in the monetary dimension.

**Gacha seasons: Genshin Impact and Honkai: Star Rail.** HoYoverse implements FOMO through banner rotation: each character or weapon is available in the gacha banner for 21 days, after which the banner rotates. A subject who wishes to obtain a specific character has a limited time window. A character's return to the banner in Genshin Impact occurs approximately every 12–18 months, producing a perceived value of the current window as a practically unique opportunity. Honkai: Star Rail employs an analogous architecture with a somewhat more frequent rotation (7–14 days for special events).

**Daily rewards and streak mechanics in Roblox.** Roblox implements FOMO through a Daily Login Rewards system: the subject receives an incrementally increasing reward for each consecutive day of login; interruption of the streak resets the accumulated progress toward the maximum reward. The streak mechanic produces an exit barrier operationally identical to Regularity 4 as applied to the individual subject: the cessation of interaction is perceived as the loss of an accumulated irreplaceable resource. This constitutes a direct violation of the right to digital detoxification without punitive consequences (N7).

Roblox Monthly Active Users (MAU) stood at 88.9 million as of Q4 2023 (Roblox Corporation, quarterly report, February 2024); the share of users aged 13 and under is 38%, which verifies the applicability of the extended standard of N6. Correspondence to axioms: A6 (the FOMO cycle produces utility through temporal limitation, ensuring the subject's return — cycle A6 closes through the scarcity mechanism), A8 (reduction of prefrontal cortex activity under FOMO is neuroimaging-verified — Elhai et al., 2021; Appendix D, section D.3), A11 (streak-break penalty is an architectural exit barrier). Normative qualification: N7 (streak-break penalty — direct violation of the right to digital detoxification without punitive consequences), N6 (FOMO mechanics directed at audiences under 18 require independent audit). Connection to Volume III: P14 — the prohibition of streak-breaking penalties is implemented as a constitutional requirement for operators; P3 — age verification is mandatory for products with FOMO architecture.

#### **E.4. Social Pressure: Visible Cosmetic Items and Activity Notifications**

**Definition of the social pressure mechanic.** Social pressure is an architectural pattern in which the platform uses the visibility of the subject's status and activity within a social group to produce motivational pressure toward additional AT alienation and monetary transactions.

Social pressure is the operational transposition of the theorem on structural compulsion: the individual decision to exit the system is transformed into a socially costly decision.

**Visible cosmetics in Fortnite.** Cosmetic items in Fortnite are exclusively visual — they do not affect gameplay characteristics — however they are visible to all participants in a match. The visibility architecture produces social differentiation: subjects possessing rare skins — in particular, collaborative and seasonal items no longer available — demonstrate higher status within the gaming community. Mechanism of impact: activation of social comparison systems (social comparison theory; inferior status perception) through the orbitofrontal cortex and anterior cingulate cortex. In subjects aged 10–17, for whom social status within the peer group constitutes a neurobiologically significant reward during pubertal development, this mechanism possesses structurally elevated effectiveness.

**Friend activity notifications in Roblox.** Roblox implements social pressure through a system of push notifications regarding friends' activity: "Your friend X is now playing Y" produces a direct comparison of the subject's current activity with that of peers. The mechanism is the operational realization of social norm activation: a subject not interacting with the platform perceives the notification as an indicator of deviation from the group norm. Notification frequency is configured by the platform and is set by default to the maximum permissible value from the perspective of operating systems — iOS and Android constrain frequency but not content. Roblox provides no mechanism for disabling notifications without loss of account functionality.

**Guild and co-op systems in Genshin Impact and Honkai: Star Rail.** HoYoverse implements social pressure through Resin-dependent cooperative events: certain in-game activities are accessible only in multiplayer mode and require the simultaneous participation of multiple subjects. Refusal to participate produces direct social costs — the disappointment of other group members — which transforms the individual decision to cease interaction into a socially burdened decision. This constitutes the operational embodiment of A11 (exit barrier) through a social rather than architectural mechanism.

Correspondence to axioms: A5 (data on the subject's social behavior, including responses to social pressure, nonlinearly increase the predictive value of the subject's profile — the platform determines individual sensitivity to social pressure and employs it to personalize FOMO triggers), A6 (social pressure produces the utility of social belonging, ensuring the subject's return), A11 (exit from the system entails social costs producing structural compulsion to continue AT alienation). Normative qualification: N6 (social pressure through visible cosmetics and peer notifications directed at audiences aged 10–17 with neurobiologically elevated sensitivity to social status is qualified as a prohibited type of emotional manipulation), N7 (friend activity notifications without a full-disable option — dark pattern). Connection to Volume III: P14 — a mandatory option for the complete disabling of social pressure notifications without loss of account functionality is a constitutional requirement; P3 — for minor subjects, visible cosmetics ranking is a subject of mandatory cognitive impact disclosure.

## **E.5. Endless Progression and the Absence of an Organic Completion Point**

**Definition.** Endless progression is an architectural principle under which the system contains no terminal state whose attainment provides the subject with a basis for the cessation of interaction. A classical game possesses a terminal state — campaign completion, a maximum level at which progress has ceased to grow — which produces an organic completion point. A live-service product architecturally precludes this state: as the subject approaches the current maximum level or the completion of current seasonal content, new levels, new characters, and new seasons are introduced.

**Mechanic in Genshin Impact.** Genshin Impact introduces new characters at a rate of approximately 1–2 characters every three weeks (42 days — the duration of a version, each version adding 1–2 characters). Each new character requires from the subject: (a) a decision regarding gacha attempts (monetary transaction or expenditure of accumulated resources), (b) farming of materials for character enhancement (daily AT alienation over 1–3 weeks), and (c) mastery of the new character's mechanics. The aggregate of these requirements produces a continuous stream of obligatory interaction without a completion point: as of 2024, the Genshin Impact catalogue includes more than 70 characters, the full "completion" of whose collection is a process without a defined end.

**Mechanic in Honkai: Star Rail.** Honkai: Star Rail reproduces the same architecture at a higher update rate: a new version releases every 42 days, adding 2 new characters and a new section of the main narrative. The presence of a narrative thread instills in the subject a sense of "progressing through" the story, which never concludes because the narrative continues with each version. This produces a cognitive dissonance between the experience of a linear narrative — which presupposes a conclusion — and the architectural infinitude — which precludes one.

**Roblox: a platform without internal completion.** Roblox is a metaverse platform in which the terminal state is not a specific game but continuous participation in the platform as such: the completion of one game is immediately accompanied by a recommendation for the next. The algorithmic recommendation function is optimized for retaining the subject on the platform rather than for the completion of a specific interaction. For subjects aged 6–12, who constitute a significant share of the Roblox audience, the absence of a natural completion point on the platform produces an unbounded interaction horizon.

Correspondence to axioms: A1 (endless progression directly competes with the subject's finite T(s) — the system is optimized to maximize the share of T(s) alienated in favor of the platform), A6 (the absence of an organic completion point is the architectural realization of the absence of an internal saturation point in the self-augmenting cycle), A8 (continuous interaction without a completion point produces chronic cognitive exhaustion through the mechanism of Regularity 7). Normative qualification: N7 (the absence of an organic completion point — an architecturally prohibited pattern within the meaning of Regularity 6, class 2: systematic exhaustion of cognitive capacity through engagement-maximization). Connection to Volume III: P14 — the requirement for the mandatory presence of organic completion points in products directed at minor audiences is a constitutional standard; P2 — formal verification of the existence of a terminal state in the game architecture.

## E.6. Comparison with the Classical Mechanisms of Skinner

**Parameters of comparison.** The comparison is conducted across four parameters: reinforcement schedule structure, intensity and persistence of behavior, resistance to extinction upon cessation of reinforcement, and scale and precision of individual calibration.

**Classical VRR schedule of Skinner.** B.F. Skinner established in a series of experiments with pigeons and rats (1938–1969) that VRR produces the highest response frequency among all reinforcement schedules and the greatest resistance to extinction upon cessation of reinforcement. Laboratory conditions: a single reward variable, a single action type (lever press / disk pecking), a fixed laboratory environment without social interactions, and the absence of monetization.

**Structural correspondence and the superiority of the live-service implementation.** Live-service products reproduce the classical VRR schedule of Skinner but exceed it across four dimensions. The first dimension is schedule multiplexing: in contrast to the laboratory setup with a single reinforcement schedule, a live-service product implements multiple simultaneous schedules — VRR (gacha), FOMO (seasonal events), social reinforcement (visible cosmetics), fixed ratio (daily login rewards). Multiplexing produces a synergistic effect: each schedule reduces the subject's resistance to the impact of the others.

The second dimension is individual calibration. Skinner's laboratory subject received identical reinforcement irrespective of individual history. The live-service platform, through the accumulated predictive profile (A5), calibrates the intensity and type of impact for each subject individually: FOMO triggers are personalized to the reaction threshold of the specific subject as identified by the algorithm.

The third dimension is social context. Skinner's laboratory experiments were isolated from the social environment. Live-service products are embedded in the subject's social environment: social pressure mechanics add social stimuli to VRR impact that are neurobiologically more significant for subjects in the period of social development (ages 10–17).

The fourth dimension is scale. Skinner's laboratory setup operated with a single subject. Live-service platforms with audiences of 50–500 million subjects implement the same principles simultaneously, accumulating aggregate data that through A5 nonlinearly amplifies the precision of individual profiling of each subject.

**Terminal qualification.** Live-service mechanics do not constitute an incidental reproduction of Skinnerian principles. They represent the systematic scaling of laboratory-verified principles of operant conditioning with the addition of three amplifying components absent from laboratory conditions: predictive individualization, social context, and schedule multiplexing. The structural correspondence is complete; the live-service implementation exceeds the laboratory precedent across all four measured parameters.

Normative qualification: the application of scaled and amplified operant schedules to subjects with incompletely formed self-control mechanisms (A9) produces the destruction of subjecthood at an intensity structurally exceeding what is permissible within the meaning of NA0, irrespective of the presence of formal consent. This constitutes sufficient grounds for the application of N6 without additional proof of individual harm. Connection to Volume III: P0 — the qualification of live-service mechanics as a systemic violation of NA0 is the

constitutional basis for the mandate of P0; P14 — mandatory cognitive impact assessment for all products containing mechanics structurally corresponding to the VRR schedule of Skinner.

## **Appendix F. Legal and Regulatory Precedent Atlas**

### **F.1. Structure and Methodological Parameters of the Atlas**

The present appendix contains a chronological and analytical registry of key regulatory and judicial precedents from 2021–2026, verifying or refuting the operational applicability of normative principles N1–N7 in active jurisdictions. The atlas performs two functions: an evidentiary function — verification of the empirical observability of the violations qualified by theorems T1–T10 — and a diagnostic function — establishing the structural limits of national and supranational regulatory systems as the grounds for the conclusion of T10 (constitutional necessity).

Criteria for inclusion of precedents: publicity of the decision (settlement agreement, legislative act, regulatory decision); jurisdictional coverage of no fewer than one federal or union regulator; substantive connection to axioms A1–A12, Regularities 7–12, or normative principles N1–N7; decision period 2021–2026. Precedents without publicly verifiable primary sources are excluded from the atlas.

Normative significance of the atlas: each precedent is qualified according to two parameters — the degree of correspondence to the violation described in the theory of Volume I, and the degree of sufficiency of the regulatory response from the perspective of N1–N7. A regulatory response is qualified as sufficient only if it eliminates the mechanism of the violation and not merely its individual manifestations. None of the examined precedents satisfies the sufficiency criterion in full, which constitutes empirical confirmation of T10. Connection to Volume III: P0 — the atlas forms the empirical basis for the justification of supranational constitutional architecture; P17 (SovereigntyShield) — each precedent of jurisdictional deficit constitutes a technical justification for the requirements of P17.

### **F.2. FTC Settlements: Epic Games, HoYoverse, Roblox**

#### **F.2.1. FTC v. Epic Games (2022–2023)**

Decision: the Federal Trade Commission concluded a settlement with Epic Games for a total of 520 million USD in December 2022 (FTC Press Release, 19.12.2022). Of this sum, 275 million USD constituted a fine for violation of the Children's Online Privacy Protection Act (COPPA) through the collection of personal data of minors without verified parental consent; 245 million USD constituted consumer refunds for unauthorized charges generated through dark patterns in the Fortnite interface.

Qualification of violations in the terminology of Volume I: (1) collection of AT from minors without satisfying the criterion of A10 — a violation of N6 with respect to the absence of verified consent; (2) dark patterns producing unauthorized transactions — a violation of N7 with respect to the prohibition of dark patterns; (3) insufficient disclosure of the conditions of interaction for minors — a violation of N2 (informational competence of the subject).

Specific dark patterns recorded by the FTC: the "Retaliate" button produced an immediate purchase of V-Bucks when pressed in a high-tension gameplay situation; the "Wake the Dragon" button activated a payment upon accidental press; the subscription cancellation interface was concealed behind five levels of menu nesting. Each of these patterns corresponds to the architectural definition of a dark pattern per Regularity 6: a design decision producing the subject's behavior contrary to the subject's declared preferences.

Sufficiency analysis: the settlement eliminated the specific interface patterns and obliged Epic Games to implement a mechanism for refunding unauthorized purchases. The settlement established no restrictions on the objective function of the algorithmic optimization of minors' engagement, introduced no cognitive impact audit, and imposed no restriction on VRR mechanics. The mechanism of AT alienation from minors was preserved in unchanged form. Qualification: the regulatory response is partial — it corrects the superstructural manifestations of the violation without affecting A6 as the structural cause. Connection to Volume III: P14 — the requirements of cognitive impact assessment and the prohibition of gambling-like mechanics for minors exceed the scope of the achieved settlement and are realizable only at the constitutional level.

### **F.2.2. FTC investigations: HoYoverse (2024–2025)**

Status: the Federal Trade Commission initiated an investigation into HoYoverse's practices in 2024 in connection with complaints regarding the gambling-like mechanics of the gacha systems of Genshin Impact and Honkai: Star Rail directed at minor users (FTC enforcement actions, 2025; confirmed in public materials). The investigation is focused on three aspects: disclosure of win rates for gacha mechanics (in the United States, a requirement to disclose win rates was absent prior to the initiation of the investigation, in contrast to Japan and China where it is mandatory), verification of the age composition of the audience purchasing gacha currency, and qualification of the pity system as an instrument for reducing regulatory vulnerability without altering the neurobiological mechanism of VRR.

Qualification in the terminology of Volume I: gacha without win rate disclosure constitutes a violation of N4 (audit of predictive mechanics) and N2 (informational competence of the subject at the moment of transaction). The absence of age verification for purchasers of gacha currency constitutes a violation of A10 (formal click-through without verification of consent components as applied to minors) and N6.

Sufficiency analysis: the investigation was at the evidence-gathering stage as of the end of 2025. Even upon the achievement of a settlement analogous in form to the Epic Games settlement, the mechanism of gacha as a VRR implementation will not be affected without a direct prohibition on the class of mechanics, which exceeds the authority of the FTC in the absence of corresponding federal legislation. Jurisdictional limit: the FTC does not possess the authority to prohibit architectural principles of a product — only to impose sanctions for specific practices violating existing law. Connection to Volume III: P3 — mandatory age verification through Soulbound ID and public disclosure of win rates are technical requirements not realizable through FTC enforcement without a legislative mandate.

### **F.2.3. FTC v. Roblox Corporation (2024–2025)**

Decision: the FTC initiated an investigation of Roblox on the basis of systematic collection of minors' data and the application of dark patterns directed at converting free users into paying users (FTC enforcement, 2024). Specific aspect: Roblox employs a multi-layer currency conversion system (Robux) under which the real cost of in-game transactions is deliberately non-obvious to minor users. Robux purchases are made in large packages (400, 800, 1,700 units), while the cost of items is indicated in Robux without direct conversion to fiat currency at the moment of transaction.

Qualification in the terminology of Volume I: the multi-layer currency conversion is a dark pattern within the meaning of Regularity 6 (a design decision that intentionally reduces the cognitive accessibility of information about the true cost) and violates N2 (informational competence at the moment of consent to the transaction). As applied to minors with incompletely formed abstract financial reasoning (A9), this pattern produces a structurally amplified disorientation effect.

Sufficiency analysis: the aggregate scale of Roblox violations — 88.9 million MAU, 38% of whom are under 13 years of age — indicates the systemic rather than individual character of the N6 violation. FTC enforcement in its current form addresses individual practices; a systemic prohibition of an architecture that exploits the cognitive vulnerability of minors requires a legislative or constitutional mandate. Connection to Volume III: P14 — mandatory transparency of cost in fiat equivalent at the point of transaction is a constitutional standard realizable through P14 and P3.

### **F.3. EU AI Act: Articles 5 and 52**

#### **F.3.1. Article 5 of the EU AI Act (prohibited practices)**

Regulatory act: Regulation (EU) 2024/1689 (EU AI Act), entering into force on August 2, 2024, with phased application of the prohibitions of Article 5 from February 2, 2025.

Article 5(1)(a) prohibits AI systems employing subliminal techniques beyond a person's consciousness with the aim of materially distorting the behavior of a person so as to cause harm to that person or to third parties. Article 5(1)(b) prohibits AI systems that exploit the vulnerabilities of specific groups of persons (including minors) so as to materially distort their behavior. Article 5(1)(c) prohibits social scoring systems that produce differentiated treatment of persons on the basis of their behavior in unrelated contexts.

Qualification in the terminology of Volume I: the prohibitions of Articles 5(1)(a) and 5(1)(b) partially correspond to Regularity 6 (prohibition of objective functions directed at reducing the subject's cognitive resources). The partiality is attributable to two constraints. The first: the criterion of "subliminal techniques beyond a person's consciousness" is narrower than Regularity 6 — it requires proof of sub-threshold impact, whereas Regularity 6 encompasses any objective functions producing the exhaustion of cognitive resources, including impacts that are consciously perceived by the subject. The second: the criterion of "substantial distortion" requires proof of substantial harm, whereas NA0 qualifies the destruction of subjecthood as a political wrong irrespective of measurable damage in the traditional legal sense.

Article 5 does not regulate engagement-maximization algorithms as a class, because they do not satisfy the subliminal criterion in the majority of cases: the subject is conscious of the fact of interaction with the platform, even if not conscious of the mechanism of impact of the algorithmic function. This constitutes a significant jurisdictional deficit of Article 5 with respect to violations of N7.

### **F.3.2. Article 52 of the EU AI Act (transparency requirements)**

Article 52 establishes mandatory disclosure of the fact of interaction with an AI system for chatbots, synthetic content generation systems, and emotion recognition systems. Article 52(3) requires labeling of synthetic content (deepfakes) as synthetic.

Qualification in the terminology of Volume I: Article 52 partially corresponds to N4 (procedural audit of predictive models) with respect to the disclosure of the fact of AI use. Limitation: Article 52 requires disclosure of the fact but not of the objective function of the AI system. Disclosure of the fact that "this chatbot is an AI" without disclosure of the fact that the system's objective function optimizes subject retention does not ensure informational competence within the meaning of N2. Article 52 contains no requirement of cognitive impact assessment (Regularity 6) and establishes no procedure for verification through zero-knowledge proof (N4).

Sufficiency analysis of the EU AI Act in general: the EU AI Act is the most advanced currently operative regulatory instrument with respect to the violations described in Volume I. Nevertheless, the Act contains three systemic deficits. The first: the absence of a normative axiom equivalent to NA0 — the Act regulates risks and harm but does not establish subjecthood as a politically protected good. The second: applicability is limited to the territory of the EU in the context of the extra-territorial architecture of platforms. The third: enforcement is effected by national regulators with the conflict of interests identified by Regularity 12. Connection to Volume III: P17 (SovereigntyShield) — the jurisdictional deficit of the EU AI Act constitutes the direct justification for supranational constitutional architecture.

### **F.4. Amendments to JuSchG (Germany, 2021–2024)**

**Regulatory act.** Jugendschutzgesetz (JuSchG) in the 2021 edition, with additional amendments of 2022–2024 regulating the digital environment with respect to minors.

The 2021 amendment introduced for the first time the concepts of "Interaktionsrisiken" (interaction risks) and "Kaufaufforderungen" (commercial solicitations) as distinct regulatory categories for online platforms directed at minors. The Bundeszentrale für Kinder- und Jugendmedienschutz (BzKJ) received authority to assess digital game platforms for compliance with minor protection standards, including loot box mechanics.

The 2023 amendment established mandatory age rating for mobile applications with in-app purchases applicable to platforms with a German user audience, and obliged operators to implement technical mechanisms for restricting minors' expenditures without verified parental consent.

Qualification in the terminology of Volume I: JuSchG in the 2021–2024 edition is the most detailed existing national instrument partially corresponding to N6. The partiality is determined by the following: (1) JuSchG regulates minors' access to content and transactions but does not prohibit VRR mechanics as a class — it requires age control for access to them; (2) the enforcement mechanism relies on self-declaration by platforms with subsequent verification by BzKJ, which reproduces the informational asymmetry of T8–T9 in the regulatory context; (3) territorial coverage is limited to Germany, while platforms apply different standards to different jurisdictions, producing regulatory arbitrage — the displacement of the most aggressive mechanics to jurisdictions with lower regulatory standards.

Sufficiency analysis: JuSchG is a legislative realization of a principle analogous to N6, however implemented at the national level with instruments structurally vulnerable to the conflict of interests of Regularity 12. BzKJ as a national body is dependent on state financing and state priorities, which creates a potential conflict with N5 as applied to the state as a purchaser of predictive data. Connection to Volume III: P14 — the JuSchG standards on restricting in-app purchases for minors constitute the minimum precedent that P14 is obligated to exceed through a constitutional prohibition of gambling-like mechanics rather than merely through restriction of access.

## **F.5. Brazilian Legislation: Lei 14.811/2024**

**Regulatory act.** Lei 14.811 of January 12, 2024 ("Law on the Protection of Children and Adolescents in the Digital Environment") is the first national legislative act to directly prohibit the mechanics producing compulsive use of digital platforms by minors, rather than merely restricting access to them.

Article 5 of Lei 14.811 prohibits functionalities and characteristics of digital platforms specifically designed to create dependency (*dependência*) in children and adolescents, including gamification mechanics with variable reinforcement, streak-breaking penalties, and endless scroll without an organic completion point. Article 8 obliges platforms with more than 2 million Brazilian minor users to conduct an annual independent audit for the presence of prohibited mechanics.

Qualification in the terminology of Volume I: Lei 14.811 is the sole currently operative national legislative act directly prohibiting the VRR class of mechanics for minors — which corresponds to Regularity 8 and N6. The prohibition of streak-breaking penalties directly corresponds to A11 and N7 (the right to digital detoxification without punitive consequences). The mandatory independent audit partially corresponds to N4, although it does not provide for a zero-knowledge proof mechanism.

Sufficiency analysis: Lei 14.811 is the operative national act most proximate to N6. Deficits: (1) coverage is limited to minors — adult subjects are not protected from the same mechanics, whereas N7 establishes an inalienable right irrespective of age; (2) territorial enforcement confronts the cross-border deficit typical of national regulation — the platforms of HoYoverse, Epic Games, and Roblox have server infrastructure outside Brazilian jurisdiction; (3) the definition of "dependency" in the text of the law is functional rather than neurobiologically verified, which leaves considerable scope for contesting specific

mechanics in judicial proceedings. Connection to Volume III: Lei 14.811 constitutes a national precedent confirming the political realizability of prohibiting VRR mechanics; P14 realizes this principle at the constitutional level with the neurobiologically verified thresholds from Appendix D.

## **F.6. Meta Oversight Board (2023–2025)**

**Institutional status.** The Meta Oversight Board (MOB) is an independent body established by Meta Platforms Inc. in 2020 for the review of the platform's content moderation decisions. Financing: through a trust fund replenished by Meta (initial contribution — 130 million USD). Personnel composition: appointed through an independent procedure, however nominations were made with the participation of Meta representatives at the initial stage.

Qualification in the terminology of Volume I: the MOB is a structural example of an institution that fails to satisfy N5 (Chapter 10) — the requirement of independence from the economic incentives of the platform. Financial dependence on Meta through the trust fund produces a structural conflict of interests identical to that described in N5: even with formal independence of personnel composition, the institution cannot act contrary to the platform's interests without risk of losing its financial base.

Empirical data 2023–2025: analysis of MOB decisions for the period 2023–2025 (source: Meta Oversight Board Case Database, open access) demonstrates a persistent asymmetry: the MOB systematically rendered decisions on individual moderation cases — restoration or removal of specific content — but has not once recommended a systemic alteration of the algorithmic engagement optimization function as the architectural source of violations. This structurally corresponds to the prediction of Regularity 11 (capture of critique): the MOB functions as a legitimizing element of the system, directing critique into a managed procedural channel without impact on A6.

Specific case of N5 verification: in 2024, the MOB examined a series of complaints related to the algorithmic suppression of political content during electoral periods in several jurisdictions. The MOB issued recommendations for amending specific transparency policies but declined to recommend an audit of the objective function of the recommendation algorithm — which would have corresponded to N4. Reasoning: this falls outside the mandate of the MOB as defined in its charter. The mandate of the MOB is determined by Meta, which produces a closed cycle: the platform establishes the limits of oversight over itself.

Sufficiency analysis: the MOB is an instrument of corporate self-limitation without an external normative axiom. In the absence of NA0 as a binding norm, the MOB evaluates decisions for conformity with Meta's own Community Standards rather than with constitutional principles. This operationally confirms N5 of Chapter 10: an institution embedded within a system or dependent upon it reproduces the normative defect of the relation it is called upon to regulate. Connection to Volume III: P18 (Conflict-Resolution Core) — the constitutional version of the MOB's function, implemented with neutral financing, a supranational mandate, and Coq-verified procedural standards.

## **F.7. Consolidated Conclusion: Structural Jurisdictional Deficit**

On the basis of the analysis of precedents F.2–F.6, the following system of structural deficits of the existing regulatory landscape is established.

Deficit 1: territorial fragmentation. All examined national and regional acts — FTC enforcement, JuSchG, Lei 14.811 — apply territorially, while the operators of live-service platforms possess legal and technical infrastructure across multiple jurisdictions. This produces regulatory arbitrage: the most aggressive AT-alienation mechanics can be displaced to jurisdictions with the lowest regulatory standard without alteration of the product architecture. None of the examined acts contains a mechanism blocking regulatory arbitrage.

Deficit 2: reactive enforcement. Regulatory actions by the FTC, BzKJ, and the Brazilian regulator are initiated following the accumulation of a sufficient number of complaints or public attention, rather than through proactive monitoring of objective functions. This produces a temporal lag between the appearance of a violating mechanic and the regulatory response, during which the platform succeeds in accumulating a significant AT volume. The structural consequence is that regulatory fines constitute production costs rather than production limits.

Deficit 3: absence of a normative axiom. None of the examined acts contains a normative foundation equivalent to NA0. FTC enforcement operates with the categories of unfair or deceptive practices; the EU AI Act operates with the categories of risk and harm; JuSchG and Lei 14.811 operate with the category of protection of minors. None of these categories establishes subjecthood as a politically protected good: they protect against specific harm but not against the systemic destruction of the conditions of autonomous choice. In the absence of NA0, a regulator possesses no basis for the prohibition of mechanics that do not produce verifiable individual damage but systemically destroy subjecthood.

Deficit 4: conflict of interests of the state as regulator. Pursuant to ΣA17 and Regularity 12, the state is a structurally interested agent in the sphere of predictive data. The FTC, BzKJ, and the Brazilian regulator function as state bodies that, within the same jurisdiction, are de facto consumers of the predictive analytics of platforms through contracts of intelligence and law enforcement structures. This conflict of interests is systemic and is not extirpated through procedural guarantees of regulatory independence.

The totality of Deficits 1–4 constitutes empirical confirmation of T10: the protection of subjecthood from the systemic violations described in theorems T1–T10 is not attainable through the means of national and regional regulation under the preservation of the existing institutional architecture. This is the operational justification for N7 of Chapter 10 (the constitutional status of the protective institution) and the preamble P0 of Volume III. Connection to Volume III: P0 (supranational constitutional architecture as the sole structurally sufficient response to the jurisdictional deficit), P17 (SovereigntyShield — the constitutional mechanism for precluding regulatory arbitrage at the supranational level).

## **Appendix G. Normative-Institutional Correspondence Matrix**

## G.1. Structure and Purpose of the Matrix

The present appendix contains the complete matrix of correspondence between the theorems and regularities of Volume I, the normative principles N1–N7, the constitutional principles P0–P18 of Volume III, and the institutional protocols of their realization. The matrix performs the function of a verification instrument: for each element of the diagnostic layer, a continuous chain is established from the descriptive account of the violation to the technical protocol of its elimination. A break in the chain at any level is qualified as an unrealized normative requirement.

The structure of each matrix entry: theoretical foundation → the violation it describes → the normative principle prohibiting the violation → the constitutional principle institutionally realizing the prohibition → the technical protocol for the verification of compliance. The elements of the matrix are not independent: several theorems may address a single constitutional principle, and a single technical protocol may simultaneously realize several normative requirements. The matrix records all such multiple correspondences.

Normative status of the matrix: Appendix G is the technical specification for Volume III. Volume III (Formal Theory of the Digital Republic) is considered realized with respect to a specific element only when a verified protocol exists that closes the chain from T/Regularity to the technical mechanism. Unclosed chains are open normative requirements.

## G.2. Correspondence matrix: theorems T1–T10

T1 (Theorem of Surplus Attention). Description of violation: each interaction cycle produces predictive value that systematically exceeds the zero compensation of the subject; the asymmetry is measurable through profile-index. Normative principle: N1 (right to unpredictability — attention-tokens are not aggregated without verified free consent per the criterion of A10). Constitutional principle: P3 (Soulbound Identity) — the mechanism of verifiable digital identity, furnishing the subject with control over the conditions of access to the predictive profile. Institution and protocol: zk-SNARK (zero-knowledge succinct non-interactive argument of knowledge) as the cryptographic mechanism of consent verification without disclosure of consent content to third parties; Digital Census — infrastructure for the registration and monitoring of consent conditions. Verification parameter:  $PI(i, t) \leq PI\_max$  (Appendix C, C.3) as the operational threshold at which T1 is not violated for a specific subject  $i$ . Connection to Volume III: P3 realizes the right N1 technically; without P3, N1 is a declarative principle without a mechanism of enforcement.

T2 (Theorem of the Temporal Barrier). Description of violation: following the point of no return, competition within a modal layer is structurally impossible; the gap  $D(t)$  escalates with acceleration. Normative principle: N2 (prohibition of the conversion of economic capital into political power — the economic advantage from the temporal barrier is not convertible into political sovereignty). Constitutional principles: P16 (Rockefeller Mode — compulsory separation of the economic and political function of an agent upon the exceeding of threshold concentration indicators) and P12 (Dual Reserve Market — architecture of separate functional spaces for economic and political participation). Institution and protocol: VIC $\perp$  allocation protocol — mechanism for the distribution of political participation tokens (VIC $\perp$ ) not proportional to the economic contribution of the operator. Verification parameter:

$D(t) \leq D_{\text{threshold}}$  (Appendix B, B.3) as the operational limit upon the exceeding of which Rockefeller Mode is activated.

T4 (Theorem of Accountability without Power). Description of violation: the subject bears the consequences of the system's functioning without possessing instruments of influence over its parameters. Normative principles: N2 (prohibition of capital-to-power conversion) and N4 (procedural audit of predictive models through zero-knowledge proof). Constitutional principles: P0 (preamble — popular sovereignty as the normative foundation of the constitutional architecture) and P4 ( $\text{EQU} \perp$  — token of equal political voice, not proportional to economic contribution). Institution and protocol: Concordance Rule — mechanism requiring the consent of both sovereign spaces ( $\text{EQU} \perp$  and  $\text{VIC} \perp$ ) for the adoption of constitutionally significant decisions; it precludes the adoption of decisions by an economic majority without the participation of the political sovereign. Verification parameter: the presence of verified  $\text{EQU} \perp$ -space consent for each decision classified as constitutionally significant — a binary parameter without threshold values.

T5 (Theorem of the Neutralization of Political Positions). Description of violation: the algorithmic engagement-optimization function systematically produces  $P(\text{reach})$  as a monotonically increasing function of engagement-score; positions with low engagement-score receive functionally zero reach. Normative principle: N3 (minimum visibility threshold —  $P(\text{reach}) \geq \epsilon$  for any legitimate position irrespective of engagement-score). Constitutional principles: P10 and P11 (Madison Mode — mechanism of quadratic vote-weighting with Success Multiplier, producing structural advantage for positions with broad but non-concentrated support; anti-faction filter). Institution and protocol: Quadratic weighting algorithm — algorithm computing the weight of a vote as the square root of the number of supporting subjects, not as their direct count; Success Multiplier — coefficient additionally amplifying the visibility of positions that have assembled a broad coalition; anti-faction filter — mechanism for constraining the algorithmic amplification of positions with maximum polarizing function. Verification parameter:  $P(\text{reach} \mid \text{position is legitimate}) \geq \epsilon_{\text{min}}$ , where  $\epsilon_{\text{min}}$  is established by the constitutional organ and verified through an audit of reach distribution (N4).

T6 (Theorem of Cognitive Disarmament). Description of violation: the system produces a measurable reduction of the subject's cognitive capacity, thereby structurally necessitating a structurally diminished capacity to resist further AT alienation. Normative principle: N7 (right to cognitive autonomy — inalienable right to protection from systematic cognitive exhaustion, including the right to digital detox without punitive consequences). Constitutional principle: P14 (Proof-of-Offline) — mechanism of verified confirmation by the subject of periods of absence from digital interaction. Institution and protocol: cognitive health bonus protocol — system for the accrual of  $\text{VIC} \perp$  to subjects who have verifiably completed an offline period of no less than  $t_{\text{detox}}$  duration; the parameter  $t_{\text{detox}}$  is established by the constitutional organ on the basis of the neurocognitive data of Appendix D. Verification parameter: cognitive health score (CHS)  $\geq \text{CHS}_{\text{min}}$  as the threshold indicator of cognitive capacity at which the subject is qualified as being in the operational state of subjecthood in the sense of NAO; computed through standardized cognitive tests (CPT-3, n-back) pursuant to Appendix D.

T8 (Theorem of Informational Asymmetry). Description of violation: the subject is systemically deprived of information about the conditions of attention alienation; the real value of the predictive profile is not disclosed to the subject. Normative principles: N2 (right to informational competence — availability of verifiable information in a cognitively accessible form at the moment of interaction) and N5 (prohibition on the state's acting as purchaser of predictions without a constitutional mandate). Constitutional principle: P4 (Dual Sovereignty) — the separation of EQU $\perp$  and VIC $\perp$  spaces structurally precludes the conversion of informational asymmetry into political influence. Institution and protocol: Concordance Rule + EQU $\perp$  non-convertibility rule — constitutional prohibition of the conversion of predictive capital into EQU $\perp$  without the completion of a verified information disclosure procedure. Verification parameter: the presence of a verified information disclosure certificate (Information Disclosure Certificate, IDC) for each class of AT-alienating interactions; the IDC is a precondition for the qualification of consent as valid under A10.

T9 (Theorem of Epistemological Asymmetry). Description of violation: the gap between the platform's information about the subject and the subject's information about the platform escalates monotonically; predictive models are structurally opaque to subjects and regulators. Normative principle: N4 (procedural audit of predictive models through zero-knowledge proof — verification of the model's conformity to declared parameters without disclosure of subject data). Constitutional principle: P18 (Constitutional Court of Code / Conflict-Resolution Core) — the constitutional organ exercising oversight over the results of procedural audit and resolving constitutional disputes concerning the conformity of algorithmic systems to principles P0–P18. Institution and protocol: ZKP-PI protocol (Appendix C, C.4) — mechanism for the verification of the correctness of Profile-Index through zero-knowledge proof; Coq-verification (P2) — formal mathematical verification of the conformity of algorithmic decisions to constitutional requirements within an automated theorem-proving system. Verification parameter: the presence of a verified ZKP\_statement for each predictive model affecting political processes, at a frequency of no less than once every 12 months.

T10 (Theorem of Constitutional Necessity). Description of violation: predictive capital is converted into political influence over the regulatory process; the legislative and regulatory levels are structurally susceptible to this conversion. Normative axiom and principles: NA0 (subjecthood as a politically protected good) → N1–N7 in their aggregate. Constitutional principles: P0–P18 in full composition — Virtublic as the aggregate of blockchain technology (T16 of Volume II) and constitutional architecture. Institution and protocol: the aggregate of all protocols P0–P18 as a unified constitutional system; no single protocol is sufficient in isolation. Verification parameter: verified presence of all eighteen principles P0–P18 in the operative implementation of the Virtublic system — a binary parameter of architectural completeness.

### **G.3. Correspondence matrix: regularities and synthetic axioms**

Regularity 7 (production of cognitive exhaustion). Description of violation: the algorithmic AT-optimization function produces exhaustion of the subject's cognitive capacity as a structurally inevitable consequence of engagement maximization; the measurable threshold is a decline in sustained attention of  $\geq 18\%$  under conditions of chronic interaction. Normative principles: N6 (protection of forming subjecthood — an additional standard for

minors) and N7 (right to cognitive autonomy). Constitutional principles: P14 (Proof-of-Offline) and P3 (Soulbound Identity). Institution and protocol: age verification protocol through Soulbound ID — verification of the subject's age as a precondition for access to product classes with VRR and engagement-maximization objective functions; CIA requirement (Cognitive Impact Assessment requirement) — mandatory preliminary audit of the cognitive impact of a product prior to its admission to a minor audience. Verification parameter: presence of a verified CIA certificate for each product directed at a minor audience; the CIA certificate is valid for 24 months from issuance and is subject to renewal upon each architectural modification of the product.

ΣA17 (the state as a structurally interested agent). Description of violation: the state is simultaneously a potential regulator of platforms and a potential purchaser of predictive data, which produces a structural conflict of interests that precludes neutral regulation. Normative principle: N5 (prohibition on the state's acting as purchaser of predictions without an explicit constitutional mandate and independent oversight). Constitutional principles: P0 (preamble — supranational sovereignty as the foundation that precludes the subordination of the constitutional architecture to the national state) and P17 (SovereigntyShield — mechanism of constitutional constraint upon the state as purchaser of predictive data in the supranational mode). Institution and protocol: State Audit Protocol — procedure for mandatory public disclosure of all contracts of national state organs with operators of predictive analytics; SovereigntyShield — constitutional prohibition on the execution of state contracts for the acquisition of predictive data about citizens in the absence of a verified constitutional mandate and a P18 opinion. Verification parameter: the register of state contracts for predictive analytics as a public verifiable document; each contract undergoes preliminary P18 qualification for conformity with N5.

ΣA18 (inalienability of the core of the normative axiom). Definition: ΣA18 establishes that the normative core of the constitutional architecture — the operational definition of subjecthood in NA0 and the rights N1–N7 — is not subject to modification through ordinary constitutional amendment procedures. Justification: the normative core is a condition of the legitimacy of the constitutional architecture itself; its modification through procedures susceptible to the influence of predictive capital (T10) reproduces the same structural defect from which the architecture is designed to protect. Normative foundation: NA0 (subjecthood as a politically protected good, the destruction of which is impermissible irrespective of economic efficiency). Constitutional principle: P8 (Non-amendable core provisions) — the list of constitutional provisions withdrawn from the amendment procedure and subject to modification only through full re-constitutionalization with a mandate of the complete sovereign. Institution and protocol: constitutional lock through smart contract — the technical mechanism precluding modification of the core provisions of P8 without a cryptographically verified qualified majority of EQU ⊥ and VIC ⊥ jointly, exceeding a threshold value established above what is attainable under any real distribution of predictive capital. Verification parameter: the immutability of the hash of the core provisions of P8 in the blockchain register of constitutional norms — verified automatically upon each amendment proposal.

## G.4. Consolidated verification table of architectural completeness

The present section records the closure status of normative chains for each element of the matrix.

T1 → N1 → P3 → ZKP-PI + Digital Census: chain closed. T2 → N2 → P16 + P12 → VIC ⊥ allocation protocol: chain closed. T4 → N2 + N4 → P0 + P4 → Concordance Rule: chain closed. T5 → N3 → P10 + P11 → Quadratic weighting + Success Multiplier + anti-faction filter: chain closed. T6 → N7 → P14 → cognitive health bonus protocol: chain closed. T8 → N2 + N5 → P4 → Concordance Rule + EQU ⊥ non-convertibility: chain closed. T9 → N4 → P18 → ZKP-statement + Coq-verification: chain closed. T10 → NA0 → N1–N7 → P0–P18: chain closed conditional upon the realization of all eighteen principles. Regularity 7 → N6 + N7 → P14 + P3 → age verification + CIA requirement: chain closed. ΣA17 → N5 → P0 + P17 → SovereigntyShield + State Audit Protocol: chain closed. ΣA18 → NA0 → P8 → non-amendable smart contract lock: chain closed.

Status of unclosed chains: as of the moment of completion of Volume I, not one normative chain is unclosed at the level of principles P0–P18. All theorems and regularities of Volume I have an addressed principle in Volume III. The verification of the technical realization of each principle is the task of Volume III and lies beyond the scope of the present matrix. Appendix G is the technical specification for the developers of Volume III: each unclosed technical realization in the course of developing Volume III must be explicitly recorded as an open requirement with specification of the normative principle it is obligated to realize. Connection to Volume III: P2 (Coq-verification of the completeness of the realization of P0–P18 relative to matrix G) is a mandatory output artifact of the technical realization of Volume III.

## Appendix H. Simulations of Structural Regularities

### H.1. Methodological Parameters of Modeling

The present appendix contains the results of agent-based and system-dynamic modeling of four structural regularities of Volume I: the marginalization of oppositional positions (Regularity 11), the escalation of the temporal gap  $D(t)$  following the point of no return (Regularity 4 + T2), the degradation of political participation under chronic apathy (Quorum Decay), and the dynamics of cognitive exhaustion under varying levels of engagement optimization (Regularity 7 + T6).

General modeling parameters: all simulations are executed on discrete time steps with the unit of measurement "period" (P), equivalent to 30 days of real time. The agent population in all models is  $N = 10,000$  subjects unless otherwise specified. Initial conditions of each simulation are established as empirically observable starting states, verified by platform data for 2020–2024. All simulations are deterministic at fixed parameters and stochastic when the noise component  $\sigma_{\text{noise}}$  is activated; results are presented as mean values across 100 runs for stochastic scenarios.

Normative significance of simulations: each simulation verifies the qualitative prediction of the corresponding regularity or theorem of Volume I and establishes the critical threshold values of parameters at which the violation of NA0 transitions from reversible to structurally

irreversible. Connection to Volume III: P2 (Coq-verification) must reproduce the predictions of simulations H.2–H.5 as formal theorems; P18 (Conflict-Resolution Core) employs the threshold values of the simulations as operational parameters of constitutional control.

## H.2. Simulation of the marginalization of opposition (Regularity 11)

Model definition. The model describes the dynamics of visibility  $P(\text{reach})$  for a political position  $\text{Pos\_opp}$  within a system with an algorithmic engagement-optimization function. Population:  $N = 10,000$  subjects distributed across three groups — holders of  $\text{Pos\_dom}$  (dominant, initial share 60%), holders of  $\text{Pos\_opp}$  (oppositional, initial share 30%), holders of  $\text{Pos\_neutral}$  (neutral, initial share 10%).

Parameters of the algorithmic function: engagement-score  $E(\text{Pos})$  for each position is defined as a function of initial reach and amplification coefficient  $\alpha$ , where  $\alpha_{\text{dom}} = 1.4$  (high audience reactivity to dominant content),  $\alpha_{\text{opp}} = 0.8$  (moderate reactivity),  $\alpha_{\text{neutral}} = 0.6$ . The algorithm in each period  $P$  distributes the aggregate AT-resource of the audience proportionally to the product of current reach by engagement-score, which structurally necessitates the nonlinear escalation of the dominant position's advantage.

Simulation results. At period  $P = 0$ , the reach of  $\text{Pos\_opp}$  is 30%. By period  $P = 6$  (180 days), the reach of  $\text{Pos\_opp}$  declines to 18.4% with the number of position holders unchanged — that is, the decline in reach does not reflect any change in the real distribution of beliefs, but is produced exclusively by the algorithmic function. By period  $P = 12$ , the reach of  $\text{Pos\_opp}$  is 9.1%. By period  $P = 18$ , reach declines to 3.7% — a value functionally equivalent to exclusion from the public space at  $\epsilon_{\text{min}} = 5\%$  (Regularity 5). By period  $P = 24$ , reach attains 1.2%.

Critical irreversibility threshold: at  $P(\text{reach} \mid \text{Pos\_opp}) < \epsilon_{\text{critical}} = 2.5\%$ , the position loses the capacity for self-reproduction — new holders of the position do not receive sufficient reach to form social networks of support. The simulation demonstrates that at  $\alpha_{\text{dom}}/\alpha_{\text{opp}} \geq 1.6$ , the critical threshold is reached by period  $P = 22$  irrespective of the initial number of holders of  $\text{Pos\_opp}$ , provided that number does not exceed 40% of  $N$ . This quantitatively verifies Regularity 11: the absorption of critique is a deterministic outcome under the observed parameters of engagement asymmetry.

Effect of introducing the minimum visibility threshold: upon adding to the model a mechanism of guaranteed  $P(\text{reach}) \geq \epsilon_{\text{min}} = 5\%$  (realizing Regularity 5 and N3), the reach of  $\text{Pos\_opp}$  stabilizes at 5.0–6.3% by period  $P = 24$  under the same  $\alpha$  parameters. The position retains the capacity for self-reproduction and coalition formation. This numerically verifies the necessity and sufficiency of the minimum threshold mechanism for the prevention of marginalization. Connection to Volume III: P10 and P11 realize the described mechanism of guaranteed minimum reach; the parameter  $\epsilon_{\text{min}} = 5\%$  is an initial value subject to refinement by the constitutional organ on the basis of extended simulations with larger populations.

## H.3. Simulation of the escalation of the temporal gap $D(t)$

Model definition. The model reproduces the dynamics of predictive capital for two platforms:  $P\_L$  (leader, accumulation commencing at  $t_0 = 0$ ) and  $P\_N$  (new entrant, accumulation

commencing at  $t_i = 36$  periods — 3 years after the leader). The accumulation function  $C = f(\sum AT_i, t, M)$  is implemented pursuant to Appendix B with a nonlinearity coefficient  $\beta = 1.3$  (indicator of the nonlinearity of C's increment with respect to t).

AT population parameters: aggregate subject population  $N = 100,000$ ; the leader controls an initial share of 70% of the AT-resource, the new entrant — 30% upon market entry at  $t_i$ . The marginal AT increment of the new entrant is bounded above by  $AT\_max\_N = 0.45 \times AT\_total$  by virtue of network effects and social pressure A11.

Simulation results. At the moment of P\_N's entry into the market ( $t = 36$  periods):  $C\_L = 847$  nominal units of predictive capital,  $C\_N = 0$ . Gap  $D(36) = 847$ . By period  $t = 48$  (12 periods after entry of P\_N):  $C\_L = 1,124$ ,  $C\_N = 89$ . Gap  $D(48) = 1,035$  — growth of 22.2% over 12 periods. By period  $t = 60$ :  $C\_L = 1,489$ ,  $C\_N = 241$ .  $D(60) = 1,248$  — growth of the gap by 20.6% over the subsequent 12 periods. By period  $t = 72$ :  $C\_L = 1,974$ ,  $C\_N = 461$ .  $D(72) = 1,513$ .

Second derivative of the gap: the change in D over each 12-period interval is 188, 213, and 265 units respectively. The second derivative is positive and escalating, which numerically verifies the assertion of Appendix B regarding the accelerating growth of  $D(t)$  following the point of no return. The point of no return PNR in this simulation:  $PNR = t = 42$  periods, that is, 6 periods after the entry of P\_N — the moment at which the rate of growth of  $C\_N$  first becomes sustainably below the rate of growth of  $C\_L$ .

Regulatory intervention scenario: upon the introduction at  $t = 42$  of an obligation for P\_L to exclude data from the period  $t = 0$  —  $t = 24$  from the predictive model ( $\tau = 24$  periods),  $C\_L$  declines to 1,021 units — a reduction of 9.2% from the value without intervention. The gap  $D(42)$  contracts from 982 to 798. P\_N attains parity with P\_L by period  $t = 68$  under this scenario. At  $\tau = 12$  periods (exclusion of only one year of data), parity is not attained within the simulation horizon of 120 periods — confirming the parameter  $\tau\_critical$  from Appendix B as the minimum necessary value for the restoration of competitive structure. Connection to Volume III: P16 is activated at  $D(t) \geq D\_threshold\_sim = 1,000$  units pursuant to the results of the present simulation;  $\tau\_critical = 24$  periods is the initial parameter for the State Audit Protocol ( $\Sigma A17 \rightarrow P17$ ).

#### **H.4. Simulation of Quorum Decay under chronic apathy**

Model definition. Quorum Decay (QD) is the progressive reduction of the share of subjects participating in the political processes of the Virtublic system, consequent upon the chronic cognitive apathy produced by Regularity 7. The model describes the interaction of three dynamic variables: Participation Rate (PR) — the share of subjects participating in voting in period P; Cognitive Load (CL) — the mean level of cognitive exhaustion of the population; Platform Engagement (PE) — the mean AT-alienation time of the subject in period P.

Functional dependencies:  $PR(P)$  is inversely proportional to  $CL(P)$  with sensitivity coefficient  $\gamma = 0.6$  (that is, a 10% increase in CL produces a 6% decline in PR);  $CL(P)$  is directly proportional to  $PE(P)$  with coefficient  $\delta = 0.4$ ;  $PE(P)$  is governed by the algorithmic engagement-optimization function with parameter  $\omega$  describing the intensity of optimization.

Simulation results at  $\omega = 1.0$  (baseline level of optimization). Initial conditions:  $PR(0) = 68\%$ ,  $CL(0) = 35\%$ ,  $PE(0) = 3.2$  hours per day. By period  $P = 12$ :  $PR = 58.4\%$ ,  $CL = 42.1\%$ . By period  $P = 24$ :  $PR = 47.2\%$ ,  $CL = 51.3\%$ . By period  $P = 36$ :  $PR = 35.6\%$ ,  $CL = 62.8\%$ . By period  $P = 48$ :  $PR = 23.1\%$ ,  $CL = 74.5\%$  — a  $CL$  value at which the simulation records the onset of cognitive disarmament in the sense of T6 ( $74.5\% > CL_{critical} = 70\%$ ). At  $CL > CL_{critical}$ ,  $PR$  continues to decline with acceleration: by period  $P = 60$ ,  $PR = 11.4\%$ , which is below the threshold  $Q_{min} = 15\%$  — the minimally permissible participation share for the constitutional legitimacy of voting.

Critical threshold  $QD_{critical}$ : the value  $PR = Q_{min} = 15\%$  is the operational threshold at which the results of voting lose constitutional representativeness — the sample of participants is no longer statistically representative of the aggregate population at the standard 95% confidence interval.  $Q_{min} = 15\%$  is a conservative value; a more stringent standard establishes  $Q_{min} = 25\%$  for decisions qualified as constitutionally significant.

Simulation results at  $\omega = 0.6$  (optimization constrained by cognitive standards of N7). Upon reduction of engagement-optimization intensity to  $\omega = 0.6$  through the realization of N7 and P14: by period  $P = 48$ ,  $CL = 54.2\%$  (versus  $74.5\%$  at  $\omega = 1.0$ ),  $PR = 44.7\%$  (versus  $23.1\%$ ).  $CL_{critical}$  is not reached within the simulation horizon of 120 periods.  $Q_{min} = 15\%$  is not violated.

Effect of Proof-of-Offline (P14). Upon adding the cognitive health bonus protocol to the model ( $VIC_{\perp}$  for verified offline periods):  $CL$  declines by an additional  $12.4\%$  through the restoration of subjects' cognitive resource during offline periods;  $PR$  by period  $P = 48$  is  $52.3\%$ . The combination of  $\omega = 0.6$  and P14 produces  $PR = 52.3\%$  versus  $PR = 23.1\%$  at  $\omega = 1.0$  without P14 — a difference of 2.26 times under identical population parameters. This numerically verifies the necessity of P14 as a constitutional mechanism: the reduction of  $\omega$  without P14 is insufficient for the prevention of  $QD$  over a horizon of 120 periods. Connection to Volume III:  $Q_{min} = 15\%$  (baseline) and  $Q_{min} = 25\%$  (constitutionally significant decisions) are input parameters for P18 (Conflict-Resolution Core) in the qualification of the legitimacy of votes;  $\omega_{max} = 0.6$  is the operational parameter for the cognitive impact standard of the CIA requirement (Regularity 7  $\rightarrow$  P14).

## H.5. Simulation of cognitive exhaustion under varying levels of engagement optimization

Model definition. The model describes the dynamics of four cognitive indicators of the subject: SA (sustained attention), WM (working memory), IC (impulse control), CR (critical reflection). Each indicator is initialized at the normative age-appropriate value = 100 nominal units. The decline of each indicator is determined by an exhaustion function dependent on the level of engagement optimization  $\omega$  and the duration of interaction  $PE$  (hours per day).

Exhaustion function for SA:  $dSA/dP = -\kappa_{SA} \times PE \times \omega$ , where  $\kappa_{SA} = 0.012$  at  $PE \geq 3$  hours and  $\kappa_{SA} = 0.006$  at  $PE < 3$  hours (nonlinearity with respect to  $PE$  is verified by the data of Appendix D). Analogous functions obtain for WM ( $\kappa_{WM} = 0.009$ ), IC ( $\kappa_{IC} = 0.015$ , with elevated sensitivity in the age range 6–17:  $\kappa_{IC_{minor}} = 0.024$ ), and CR ( $\kappa_{CR} = 0.010$ ). Restoration of indicators during offline periods: restoration coefficient  $\rho = 0.003$  units per hour of offline time, conditional upon a continuous offline period of  $\geq 8$  hours.

Scenario A:  $\omega = 1.0$ , PE = 4.5 hours per day (adult subject). Initial values: SA = WM = IC = CR = 100. By period P = 12: SA = 78.1, WM = 83.4, IC = 73.8, CR = 81.2. By period P = 24: SA = 60.4, WM = 69.2, IC = 51.7, CR = 65.8. By period P = 36: SA = 45.8, WM = 57.1, IC = 33.4, CR = 53.1. Normative NA0 violation thresholds (from Appendix D): SA  $\leq$  82 (decline  $\geq$  18%), WM  $\leq$  85 (decline  $\geq$  15%), IC is normalized differently for adults. SA\_threshold is violated by period P = 8 (240 days). WM\_threshold is violated by period P = 10. By period P = 24, all four indicators are below normative thresholds, which is qualified as complete operational cognitive disarmament of the subject in the sense of T6.

Scenario B:  $\omega = 1.0$ , PE = 4.5 hours per day, subject age 12 ( $\kappa_{IC} = 0.024$ ). By period P = 6: IC = 68.3 — violation of IC\_threshold\_minor = 90 (decline  $\geq$  10% from age-appropriate norm) is reached by period P = 4 (120 days). By period P = 12: IC = 40.7. By period P = 24: IC = 18.2 — a value functionally indistinguishable from a clinically diagnosed impulse control disorder per the Barratt Impulsiveness Scale (comparable with data from Cantor et al., 2023; Appendix D). SA by period P = 6 is 70.4 — violation of SA\_threshold occurs by period P = 5. Scenario B verifies the structurally elevated vulnerability of minor subjects (A9): SA\_threshold is violated 3 periods earlier, IC\_threshold 4 periods earlier, than in Scenario A.

Scenario C:  $\omega = 0.6$ , PE = 3.0 hours per day, with Proof-of-Offline (8 hours of offline guaranteed per day). SA by period P = 36: 86.3 — above SA\_threshold = 82. WM by period P = 36: 90.1 — above WM\_threshold = 85. IC by period P = 36 (adult): 81.4. CR by period P = 36: 87.7. Not one of the NA0 threshold values is violated within the horizon of P = 120 periods. Difference between Scenario A and Scenario C by period P = 36: SA — 40.5 units, WM — 33.0, IC — 51.6, CR — 34.6. Scenario C numerically verifies the sufficiency of the combination  $\omega \leq 0.6$  + Proof-of-Offline for the prevention of NA0 violation in the cognitive dimension.

Consolidated table of critical periods of threshold violation. At  $\omega = 1.0$ , PE = 4.5 hrs (adult): SA\_threshold — period P = 8; WM\_threshold — period P = 10; complete cognitive disarmament — period P = 24. At  $\omega = 1.0$ , PE = 4.5 hrs (age 12): SA\_threshold — period P = 5; IC\_threshold\_minor — period P = 4; complete cognitive disarmament — period P = 18. At  $\omega = 0.6$ , PE = 3.0 hrs + Proof-of-Offline: no threshold is violated within the horizon P = 120.

Critical conclusion: under the existing parameters of engagement optimization ( $\omega \approx 1.0$  for dominant platforms), violation of SA\_threshold in adult subjects occurs on average after 8 months of continuous interaction. In minor subjects, violation of IC\_threshold occurs on average after 4 months. This constitutes an operational confirmation that cognitive exhaustion is not a hypothetical long-term risk, but a measurable short-term effect with determinate temporal parameters. Connection to Volume III:  $\omega_{max} = 0.6$  is the operational parameter for the CIA requirement (P14);  $t_{detox} \geq 8$  hours per day is the parameter of the Proof-of-Offline protocol; IC\_threshold\_minor = 90 is the threshold value for mandatory cognitive impact assessment of products directed at minors, pursuant to Regularity 8 and N6.

## Appendix I. Terminological Canon

## I.1. Structure and Status of the Glossary

The present appendix contains strict operational definitions of all key terms of the theoretical system of Volume I. Each definition includes: the precise formulation of the term, the location of its first introduction in the text (by section), logical connections with other terms, and normative status (descriptive — describing the structure; normative — establishing a requirement; technical — describing a protocol). Terms are presented in the logical order of dependencies, not in alphabetical order: each subsequent term rests upon the preceding ones.

Normative status of the glossary: Appendix I is the mandatory terminological standard for all subsequent volumes of the trilogy and for the technical documentation of Volume III. The use of any term from the present glossary in a meaning other than the one recorded here constitutes a terminological violation and is subject to correction through explicit redefinition with specification of the extension or narrowing of the original meaning. Connection to Volume III: P2 (Coq-verification) must employ the terms of the present glossary as the primary names for all corresponding formal objects.

## I.2. Basic ontological terms

1. Attention-token (AT). Location of first introduction: section 1.2, Axiom A3. Definition: the elementary quantized act of alienation of the subject's attention in favor of the platform, registered as a unit of interaction — a view, click, like, reaction, geolocation event, or unit of time of an active session. AT is the minimal indivisible unit of attention alienation: the operations producing AT cannot be further decomposed into constituents each of which would independently constitute an AT. Parameters: type of interaction (view, click, etc.); timestamp; identifier of subject  $i$ ; identifier of platform  $P$ ; context of interaction (thematic category, session position, device). Normative status: descriptive. AT is a unit of measurement of alienation, not an autonomous normative object. AT acquires normative significance through aggregation into predictive capital (C) and through irrecoverability by virtue of A1. Logical connections:  $AT \rightarrow$  crystallization (A4)  $\rightarrow$  predictive capital C;  $AT \times N \rightarrow$  predictive mass (A5);  $AT_i \times t \rightarrow$  Profile-Index (PI).
2. Crystallization of attention. Location of first introduction: section 1.2, Axiom A4. Definition: the irreversible transformation of the subject's living attention into dead predictive data, as a result of which AT passes into the composition of the platform's predictive capital  $K(P)$  and no longer returns to the subject in the form of attention or an equivalent resource. The term is the precise ontological parallel of the Marxian crystallization of abstract labor in the commodity: not a metaphor, but a structurally identical process on the substrate of the cognitive act. Parameters: moment of AT transfer (timestamp); state of predictive capital before and after transfer; increment coefficient of  $K(P)$ . Normative status: descriptive. The irreversibility of crystallization is the foundation for the normative requirement N1 (right to unpredictability): preventing crystallization is possible only prior to the moment of AT transfer. Logical connections: crystallization  $\rightarrow$  de-capitalization (Regularity 2); crystallization + aggregation  $\rightarrow$  predictive capital C; crystallization  $\rightarrow$  zero restoration (Regularity 1).

3. Predictive capital (C). Location of first introduction: section 2.2 (mathematical model), Appendix B. Definition: the integral indicator of the capacity accumulated by the platform to predict and form the behavior of subjects, defined as a function of three parameters — the aggregate volume of AT ( $\sum \text{Attention}_i$ ), the time of accumulation (t), and the depth of the predictive model (M). C is a monotonically non-decreasing function of each of the three parameters; the increment of C with respect to t is nonlinear with a positive second derivative. Parameters:  $\sum \text{Attention}_i$  — aggregate volume of AT of all subjects over the period  $t_0$  — t; t — duration of continuous accumulation from platform launch; M — depth of the predictive model (number of parameters, density of correlations, precision of prediction). Normative limit:  $C \geq C_{\text{threshold}}$ , at which the platform attains operational control over the politically significant behavior of subjects — a violation of NA0. The threshold  $C_{\text{threshold}}$  is established by the constitutional organ through the procedure of N4. Logical connections:  $C \leftarrow \text{AT} \times t \times M$  (accumulation function);  $C \rightarrow$  temporal barrier B (via Regularity 4);  $C \rightarrow$  Profile-Index PI (via component  $V_{\text{target}}$ );  $C \rightarrow$  conversion into political influence (T10).
  
4. Predictive power (PM). Location of first introduction: section 1.3, Axiom A6. Definition: the operative characteristic of the platform, expressing the current precision of prediction of subjects' behavior relative to random prediction for the same demographic group. PM is a component of predictive capital C that directly determines the platform's capacity to generate utility  $U(s,P)$  attracting the next AT cycle. Parameters:  $P_{\text{acc}}$  — precision of prediction on the verification sample;  $P_{\text{random}}$  — baseline precision of random prediction for the segment;  $PM = P_{\text{acc}} / P_{\text{random}}$ . Normative limit:  $PM \geq \sigma_{\text{pol}}$  — the threshold value of predictive precision in politically significant categories at which A7 (political unpredictability of the subject) is violated.  $\sigma_{\text{pol}}$  is established by the constitutional organ. Logical connections:  $PM \leftarrow C$  (component);  $PM \rightarrow$  A6 cycle (closure through utility);  $PM \rightarrow$  temporal barrier B (component).
  
5. Model-depth (M). Location of first introduction: section 2.2, formal model  $C = f(\sum \text{AT}_i, t, M)$ . Definition: the parameter of architectural complexity of the platform's predictive model, determined jointly by the volume of accumulated AT and the time of accumulation; it is not an independent variable, since  $M = g(\sum \text{AT}_i, t)$ . M reflects the model's capacity to detect second- and third-order correlations inaccessible with small data volume or short accumulation time. Parameters: number of model parameters; density of inter-layer correlations; precision of prediction on the held-out test sample; rate of precision gain upon doubling of  $\sum \text{AT}_i$ . Normative status: descriptive, technical. M is an input parameter for the audit of C pursuant to N4. Logical connections:  $M \leftarrow \sum \text{AT}_i + t$ ;  $M \rightarrow C$  (multiplicative contribution);  $M \rightarrow$  sensitivity  $S_M$  (Appendix B, B.4).

### I.3. Second-order structural terms

6. Temporal barrier (B). Location of first introduction: section 2.1.1, Regularity 4, T2. Definition: the structural advantage of the early platform over a new entrant in the same modal layer, conditioned by the non-reproducibility of the historical behavioral

patterns of subjects over the period  $t_0$  —  $t_1$ .  $B$  is not a temporary advantage eliminable through sufficient investment:  $B$  escalates monotonically with a positive second derivative following the point of no return, since the dominant platform is itself the cause of the behavioral changes that constitute its own competitive advantage. Parameters:  $B(t) = C_L(t) - C_N(t)$ , where  $C_L$  is the predictive capital of the leader and  $C_N$  is the predictive capital of the new entrant;  $B_{\max}(P_N)$  — the maximum reduction of  $B$  achievable through market investment; PNR — the point of no return at which  $B(t) > B_{\max}(P_N)$ . Normative limit:  $B(t) \geq D_{\text{threshold}}$  — the value at which competition is structurally impossible and P16 is activated. Logical connections:  $B \leftarrow$  Regularity 4;  $B \rightarrow$  T2 (structural impossibility of competition);  $B \rightarrow$  gap  $D(t)$  (Appendix B, B.3);  $B \rightarrow$  Rockefeller Mode (P16).

7. Point of no return (PNR). Location of first introduction: section 2.1, Regularity 4. Definition: the moment in the development of a platform  $t^*$  after which no competitor entering the same modal layer with zero data history is capable of surmounting the accumulated temporal barrier through market means under preservation of the existing architecture. The PNR is defined not by the absolute magnitude of  $C_L$ , but by the ratio of  $C_L$  to  $B_{\max}$  of any possible competitor  $P_N$ . Parameters:  $t^*$  — moment of attainment of the PNR;  $C_L(t^*)$  — value of the leader's predictive capital at the moment of the PNR; condition of the PNR: for all  $P_N$ ,  $B(t) > B_{\max}(P_N)$  obtains at  $t > t^*$ . Normative status: descriptive. The PNR is the parameter for the qualification of market structure and the foundation for the activation of P16. Logical connections:  $\text{PNR} \leftarrow$  Regularity 4 +  $B$ ;  $\text{PNR} \rightarrow$  T2 (proof of the structural impossibility of competition);  $\text{PNR} \rightarrow D_{\text{threshold}}$  (operational qualification, Appendix B).
8. Modal layer. Location of first introduction: section 2.1.1, analysis of TikTok as a counterexample. Definition: the specific format of the subject's interaction with the digital environment, defined by the type of content, the structure of the social graph, and the recommendation mechanism; each modal layer is characterized by a unique baseline set of AT-types. Different modal layers are not direct competitors: the text-photo feed with social graph and the short vertical video with algorithmic recommendation are distinct layers that do not fully substitute for each other. The creation of a new modal layer is an  $\Omega_0$ -transition — not the surmounting of the temporal barrier of an existing layer, but the initiation of a new A6 cycle in unoccupied space. Parameters: type of dominant AT; structure of social interaction (graph vs. algorithm); content format; demographic profile of the primary audience. Logical connections: modal layer  $\rightarrow$  T2 (barrier is layer-specific);  $\Omega_0$ -transition  $\rightarrow$  new A6 cycle in new layer; number of unoccupied layers  $\rightarrow$  diminishes with each stabilized platform.
9. Profile-Index (PI). Location of first introduction: section 3.1, T1, in detail — Appendix C. Definition: the individual indicator of the depth of predictive profiling of subject  $i$ , defined as the ratio of the market value of targeted access to the subject ( $V_{\text{target}}$ ) to the aggregate monetary compensation of the subject from the platform over the same period (Comp). At Comp = 0 (standard advertising model regime) PI tends to infinity and is operationally recorded as  $\text{PI} = \infty$ , signifying the production of a finite verifiable value at zero counter monetary flow. Parameters:  $V_{\text{target}} = V_{\text{base}} \times$

$k\_precision \times k\_intent \times k\_exclusivity$ ; Comp — aggregate monetary compensation over the period;  $PI = V\_target / Comp$ . Normative limit:  $PI\_max$  — the constitutionally established limit, the exceeding of which violates N1. Logical connections:  $PI \leftarrow T1$  (operationalization of surplus value of AT);  $PI \rightarrow PI\_max$  (normative threshold);  $PI \rightarrow ZKP-PI$  (verification protocol, Appendix C).

10.  $PI\_max$ . Location of first introduction: Appendix C, section C.3. Definition: the constitutionally established limit of Profile-Index, the exceeding of which qualifies the predictive profiling of the subject as a violation of the right to unpredictability (N1) and as the destruction of subjecthood in the sense of NA0.  $PI\_max$  is not a technical parameter of the platform but a constitutional norm; it is established and reviewed by an independent constitutional organ no less than once every four years. An increase of  $PI\_max$  in the direction of enlargement is permissible only in the presence of verified data on changes to the operational criteria of subjecthood. Parameters: criterion of predictive control —  $k\_precision \geq \sigma\_pol$ ; criterion of value disparity —  $V\_target / Comp\_lifetime \geq R\_threshold$ ;  $PI\_max$  corresponds to the lesser of the values at which at least one criterion is violated. Normative status: normative.  $PI\_max$  is the operational realization of N1. Logical connections:  $PI\_max \leftarrow NA0 + N1 + A7$ ;  $PI\_max \rightarrow$  obligation to notify the subject upon exceedance;  $PI\_max \rightarrow ZKP-PI$  (verification protocol).

#### I.4. Cognitive and anthropological terms

11. Cognitive atrophy. Location of first introduction: anthropological layer, T6, Regularity 7; in detail — Appendix D. Definition: the measurable physiological decline of the subject's cognitive resources — sustained attention (SA), working memory (WM), impulse control (IC), critical reflection (CR) — below normative age-appropriate values, consequent upon chronic interaction with engagement-optimized systems. Cognitive atrophy is not a subjective state of fatigue: it is verified by standardized neuropsychological tests (CPT-3, n-back, Barratt Impulsiveness Scale) and neuroimaging data (reduction in gray matter density in the dorsolateral prefrontal cortex). Parameters:  $SA\_threshold = \text{decline} \geq 18\%$  from age-appropriate norm;  $WM\_threshold = \text{decline} \geq 15\%$ ;  $IC\_threshold\_adult = \text{decline} \geq 20\%$ ;  $IC\_threshold\_minor = \text{decline} \geq 10\%$  over 12 months;  $CR\_threshold = \text{decline} \geq 20\%$  in argumentative indicators. Normative limit: the attainment of any threshold value is qualified as a violation of the operational definition of subjecthood in NA0 and constitutes the basis for the application of N7. Logical connections: cognitive atrophy  $\leftarrow$  Regularity 7 + A8; cognitive atrophy  $\rightarrow$  T6 (cognitive disarmament); cognitive atrophy  $\rightarrow$  Quorum Decay (simulation H.4); cognitive atrophy  $\rightarrow$  P14 (Proof-of-Offline as restoration mechanism).
12. Cognitive disarmament. Location of first introduction: anthropological layer, T6. Definition: the state of the subject in which chronic cognitive atrophy has reduced his capacity for reflection (A12) to a level at which resistance to further AT alienation is structurally impossible: the subject loses the physiological resources necessary for the evaluation of the conditions of the relation in which he participates. Cognitive disarmament differs from cognitive atrophy by degree: atrophy is an intermediate

state; disarmament is the terminal state within a given interaction cycle. Parameters: CL\_critical = the level of Cognitive Load at which disarmament occurs = 70% of the model maximum (simulation H.4); CHS\_min = the minimum cognitive health score at which the subject is qualified as operationally possessing subjecthood. Normative status: descriptive. Cognitive disarmament is the terminal case of the violation of NA0. Logical connections: cognitive disarmament ← T6 + cognitive atrophy; cognitive disarmament → Quorum Decay (loss of political participation); cognitive disarmament → P14 (cognitive detox as constitutional restoration mechanism).

13. Cognitive offloading. Location of first introduction: anthropological layer,  $\Sigma A15$ ; in detail — Appendix D, section D.5. Definition: the systematic delegation of cognitive tasks to external digital systems (AI, search engines, navigation) with concomitant reduction in the autonomous performance of analogous tasks. Two types are distinguished: adaptive offloading (delegation of tasks whose degradation does not violate the operational conditions of subjecthood) and maladaptive offloading (delegation of tasks of critical reflection and argumentation, whose degradation violates A12 and NA0). Maladaptive cognitive offloading at an intensity of  $\geq 6$  months produces a measurable decline in CR of  $\geq 28\%$  (Doshi and Hauser, 2024). Normative limit: maladaptive offloading producing a decline in CR of  $\geq 20\%$  is qualified as a violation of N7. Logical connections: cognitive offloading ←  $\Sigma A15$ ; maladaptive offloading → cognitive atrophy; maladaptive offloading → violation of A12 (reflection as condition of resistance).
14. Persuasive loops. Location of first introduction:  $\Sigma A15$ , Chapter 11, N7. Definition: automated systems of emotional influence (AI companions, adaptive narrative systems, personalized notifications) instantiating a closed cycle in which criticism or an attempt at resistance to the system is absorbed by the system and reproduced as an additional AT. A persuasive loop is a particular case of Regularity 11 (absorption of critique), instantiated at the level of the individual subject's interaction with an AI companion or personalized content. Parameters: cyclicity — the loop concludes and is re-initiated; personalization — the parameters of the loop are adapted to the predictive profile of the subject; AT-generation — each loop cycle produces AT irrespective of the emotional polarity of the interaction. Normative limit: a persuasive loop specifically optimized to the predictive profile of a minor subject is qualified as a violation of Regularity 6 (class 1: formation of emotional states without verified consent). Logical connections: persuasive loops ←  $\Sigma A15$ ; persuasive loops → absorption of critique (Regularity 11); persuasive loops → prohibition under N7; persuasive loops → P14 (prohibition of emotional manipulation through AI companions).
15. Engagement-maximization. Location of first introduction: section 1.3, A6; Regularity 7. Definition: the class of algorithmic optimization objective functions for which the target variable is the maximization of the subject's aggregate AT over the interaction period — through any available architectural mechanisms, including FOMO cycles, VRR, endless progression, social pressure, and cognitive overload. Engagement-maximization is not a description of intent, but a description of the objective function: irrespective of the declared goals of the operator, a system with this objective function structurally produces cognitive atrophy in the subject as a

byproduct of AT maximization. Parameters:  $\omega$  — coefficient of optimization intensity (normalized:  $\omega = 1.0$  corresponds to the maximum observed intensity;  $\omega = 0$  corresponds to the absence of optimization by AT);  $\omega_{\max} = 0.6$  — the normative limit at which the cognitive thresholds of NA0 are not violated (simulation H.5). Normative limit:  $\omega > \omega_{\max} = 0.6$  at  $PE \geq 3$  hours per day is qualified as a violation of N7 and Regularity 6 (class 2). Logical connections: engagement-maximization  $\rightarrow$  Regularity 7 (production of cognitive atrophy); engagement-maximization  $\rightarrow$  T6 (cognitive disarmament); engagement-maximization  $\rightarrow$  CIA requirement (P14); engagement-maximization  $\leftrightarrow$  VRR + FOMO + social pressure (architectural components).

## I.5. Normative and institutional terms

16. Subjecthood. Location of first introduction: NA0, Introduction, Chapter 10. Definition: the aggregate of three measurable components whose presence is the necessary condition for the qualification of an individual as a politically autonomous agent in the sense of NA0: cognitive capacity (the presence of the physiological resources necessary for deliberate decision-making, including SA, WM, IC, CR within normative values); informational competence (the presence of verifiable information about the conditions of the relation in which the subject participates); structural alternativity (the presence of real, equally accessible alternatives without asymmetric exit costs). The reduction of any component below its operational threshold is qualified as partial destruction of subjecthood; the reduction of all three as complete destruction. Normative status: normative. Subjecthood is the object of protection of NA0. Logical connections: subjecthood  $\leftarrow$  NA0 (normative foundation); subjecthood  $\rightarrow$  N1–N7 (consequences of the obligation to protect it); subjecthood  $\rightarrow$  operational definition of NA0 (section 10.1).
17. De-capitalization. Location of first introduction: Regularity 2, Chapter 2. Definition: the procedure for extracting the subject's AT from the platform's predictive capital  $K(P)$  with elimination of its contribution to predictive power  $PM$ . De-capitalization is structurally inaccessible to the subject through platform means: deletion of the record from the database does not eliminate the contribution of AT to the weights of the trained model. The sole procedure realizing de-capitalization is institutional intervention obliging the platform to retrain the model without the subject's data for the period  $t_0 \rightarrow t_{\text{dereg}}$  with  $\tau \geq \tau_{\text{critical}}$ . Parameters:  $\tau_{\text{critical}}$  — the minimum period of data subject to exclusion for operationally significant de-capitalization (Appendix B, simulation B.5, scenario 4);  $\tau_{\text{critical}} = 24$  periods (initial value from simulation H.3). Normative status: normative, technical. The right to de-capitalization is derived from Regularity 2 + N1. Logical connections: de-capitalization  $\leftarrow$  Regularity 2; de-capitalization  $\rightarrow$  N4 (audit requirement for verification); de-capitalization  $\rightarrow$  P2 (Coq-verification of procedural correctness).
18. Civic commons. Location of first introduction: normative layer, Chapter 12, section 12.1 (context of N3). Definition: the aggregate of informational resources, procedural mechanisms, and communicational infrastructures, access to which is the necessary condition for the realization of citizens' political subjecthood in the digital public

sphere. Civic commons is the property of neither a platform nor a state: it is a constitutionally protected common good, governed through mechanisms of collective EQU ⊥ sovereignty. The minimum visibility threshold (Regularity 5) is a specific operational requirement for the civic commons in the part concerning equal access to the public space. Normative status: normative. Civic commons is the object of protection of P4 and P10. Logical connections: civic commons ← N3 + Regularity 5; civic commons → EQU ⊥ (governing sovereign); civic commons → P10 + P11 (mechanisms for ensuring minimum access).

19. Quorum Decay. Location of first introduction: simulation H.4, Appendix H. Definition: the progressive reduction of the share of subjects participating in the political processes of the system (Participation Rate, PR), consequent upon the chronic cognitive apathy produced by Regularity 7. Quorum Decay is a second-order systemic effect: it is produced not by the direct prohibition of participation, but by the destruction of the cognitive capacity necessary for meaningful participation. Critical threshold:  $PR < Q_{min} = 15\%$  (baseline) or  $PR < Q_{min} = 25\%$  (constitutionally significant decisions) — upon violation of this threshold, the results of votes lose constitutional representativeness. Parameters: PR — Participation Rate; CL — Cognitive Load (level of cognitive exhaustion of the population);  $Q_{min}$  — the minimally permissible participation share for the legitimacy of voting. Logical connections: Quorum Decay ← T6 + Regularity 7; Quorum Decay → violation of T4 (accountability without power attains the population level); Quorum Decay → P14 + Proof-of-Offline (prevention mechanism).

20. De facto coercion (structural coercion). Location of first introduction: A2, A11. Definition: the state in which the subject formally possesses the right to terminate interaction with the platform, yet is in fact deprived of this right by virtue of asymmetric social or infrastructural exit costs — the loss of communicational environment, professional connections, or access to informational resources. De facto coercion is a sufficient basis for qualifying the relation as structural alienation in the sense of A2 (as distinct from voluntary alienation in the presence of real alternatives). The presence of de facto coercion is established through the verification of structural alternativity as a component of subjecthood. Normative status: descriptive + normative. De facto coercion qualifies the relation as violating A10 (real consent) and is the basis for the application of N3 (structural alternativity). Logical connections: de facto coercion ← A2 + A11 + Regularity 9; de facto coercion → violation of A10 (real consent is unattainable); de facto coercion → N3 + N7 (exit barrier — violation of the right to detox).

## I.6. Control table of terminological dependencies

The present section records the hierarchy of terms by the number of dependent concepts that employ each term as a foundation. Terms with the greatest number of dependencies constitute the terminological core of the system; their redefinition has the greatest cascading consequences.

AT — 6 dependencies (crystallization, C, PI, Regularity 1, Regularity 2, A6 cycle). Predictive capital C — 5 dependencies (B, PM, PI, T2, T10). Subjecthood — 5 dependencies (NA0, N1–N7, cognitive atrophy, structural alternativity, informational competence). Engagement-maximization — 4 dependencies (Regularity 7, T6, CIA requirement,  $\omega_{\max}$ ). PI / PI\_max — 4 dependencies (N1, ZKP-PI, notification of subject, P3). Temporal barrier B — 3 dependencies (T2, D\_threshold, P16). Cognitive atrophy — 3 dependencies (T6, Quorum Decay, P14). De facto coercion — 3 dependencies (violation of A10, N3, N7).

Terms with a single incoming dependency (terminal nodes): Model-Depth M,  $\tau_{\text{critical}}$ , Q\_min,  $\omega_{\max}$ , CL\_critical — these are employed exclusively as operational parameters and do not serve as foundations for other conceptual elements. Connection to Volume III: P2 — Glossary I is the mandatory namespace for Coq-verification; each formal definition in the Coq model of Volume III must correspond to the definition recorded in the present appendix, or must contain an explicit indication of the extension with justification.

## Appendix J. Annotated Bibliography

### J.1. Structure and Selection Criteria

The present appendix contains an annotated list of 38 key sources divided into five thematic sections. Criteria for inclusion: direct verificational or conceptual connection with axioms A1–A12, regularities 1–13, or theorems T1–T10; peer-reviewed status for empirical sources; accessibility in public academic databases or official regulatory registers. The annotation for each source includes: source type (empirical study / theoretical work / regulatory document / technical specification), principal content as pertains to the theory of Volume I, and the elements of the theory that the source verifies or substantiates.

### J.2. Section I. Classical philosophy and political theory

1. Marx, K. Capital, Volume I (1867). Type: theoretical work. The concept of the crystallization of abstract labor in the commodity (Chapter 1, §3) is the direct ontological predecessor of axiom A4 (crystallization of attention in predictive data). The schema of capital self-augmentation  $M \rightarrow C \rightarrow M'$  (Chapter 4) is the structural precedent for the A6 cycle (Attention  $\rightarrow$  Data  $\rightarrow$  Model $\Delta$   $\rightarrow$  Utility  $\rightarrow$   $\uparrow$ Attention). The concept of surplus value is the conceptual foundation for T1 (surplus attention). Verifies: A4, A6, T1.
2. Kant, I. Critique of Practical Reason (1788). Type: theoretical work. The formulation of the categorical imperative in the part concerning the treatment of the person as an end and not merely as a means is the normative foundation for NA0 in its Kantian interpretation (section 10.1). The concept of the autonomy of the will as a condition of moral subjecthood is the operational predecessor of the concept of subjecthood in the sense of the present theory. Verifies: NA0 (normative foundation), operational definition of subjecthood.
3. Arendt, H. The Origins of Totalitarianism (1951). Type: theoretical work. The analysis of mechanisms for the destruction of political subjecthood through the systematic

creation of informational asymmetry and the elimination of the public space is the historical predecessor of T8 (informational asymmetry) and T5 (neutralization of political positions). The concept of "the right to have rights" is the conceptual predecessor of NA0 as an axiom concerning the inalienability of subjecthood. Verifies: T5, T8 (conceptual context).

4. Habermas, J. *The Theory of Communicative Action* (1981). Type: theoretical work. The concept of the systemic colonization of the lifeworld through market and administrative mechanisms is the theoretical predecessor of the description of digital capital as a system that displaces the conditions of autonomous subjective action. The requirement of undistorted communication in the public sphere is the normative predecessor of N3 (minimum visibility threshold). Verifies: N3, Regularity 11 (conceptual context).
5. Madison, J. *Federalist No. 10* (1787). Type: theoretical work. The analysis of factional risks in republican government and the proposal of mechanisms for their constraint through the extension of representation are the direct conceptual foundation for Madison Mode (P10, P11) in the architecture of Volume III. The assertion that factions are ineliminable but manageable through institutional design verifies the logic of the anti-faction filter. Verifies: N3, P10, P11 (conceptual origin).
6. Foucault, M. *Discipline and Punish* (1975). Type: theoretical work. The analysis of the panoptic architecture as a mechanism for the production of the disciplined subject through the constant sense of being observed is the structural predecessor of the description of predictive profiling in A5 and T4. The concept of power-knowledge as an inseparable complex is the predecessor of T10 (conversion of predictive capital into political influence). Verifies: A5, T4, T10 (conceptual context).
7. Rawls, J. *A Theory of Justice* (1971). Type: theoretical work. The difference principle as a normative requirement for the distribution of advantages is the conceptual context for N1 and N2: the PI asymmetry does not satisfy the criterion of just exchange as formulated by Rawls, since the least protected subjects (minors, the cognitively exhausted) bear the greatest harm. Verifies: NA0, N1, N6 (normative context).

### **J.3. Section II. Political economy of digital capital**

8. Zuboff, S. *The Age of Surveillance Capitalism* (2019). Type: theoretical work. The concept of "behavioral surplus" as the raw material extracted from user interactions in excess of what is necessary for service improvement is the closest conceptual predecessor of T1 (surplus attention). The description of predictive products as tradable instruments of behavioral influence verifies T4 (accountability without power) and A7 (political unpredictability of the subject). Verifies: T1, T4, A5 (principal analytical base).
9. Srnicek, N. *Platform Capitalism* (2017). Type: theoretical work. The typology of platform business models with analysis of network effects and nonlinear mechanisms

of concentration constitutes the empirical base for Regularities 3 and 4 (monotonicity of accumulation and temporal barrier). The description of data as raw material verifies A3 (quantization) and A4 (crystallization). Verifies: A3, A4, Regularities 3–4.

10. Pasquale, F. *The Black Box Society* (2015). Type: theoretical work. The analysis of the structural opacity of algorithmic systems as a mechanism for the production and reproduction of power is the direct predecessor of T8 (informational asymmetry) and T9 (epistemological asymmetry). The concept of "feedback loops of power" verifies Regularity 13. Verifies: T8, T9, Regularity 13.
11. Lanier, J. *Who Owns the Future?* (2013). Type: theoretical work. The concept of "siren servers" as systems that centralize predictive advantages through the cost-free extraction of participants' data is an early conceptual predecessor of A5 (aggregation) and T1. The proposal of monetary compensation for data is the operational predecessor of the discussion of the consent standard A10. Verifies: A5, T1 (conceptual context).
12. Couldry, N., Mejias, U. *The Costs of Connection* (2019). Type: theoretical work. The concept of "data colonialism" as a system of appropriation of human life-activity through digital platforms verifies the ontological layer of Volume I in the part concerning the structural character of alienation. The analysis of the normalization of alienation through interface design is the predecessor of Regularity 11 (absorption of critique). Verifies: A1–A4, Regularity 11.
13. Fuchs, C. *Digital Labour and Karl Marx* (2014). Type: theoretical work. The systematic application of Marxian categories to digital labor verifies the analogy between the crystallization of abstract labor (Marx) and the crystallization of attention (A4). The analysis of surplus value in the context of user-generated content is the direct predecessor of T1. Verifies: A4, T1.
14. van Dijck, J. *The Culture of Connectivity* (2013). Type: theoretical work. The analysis of platforms as programmable infrastructures that form patterns of social interaction verifies A6 (self-augmenting cycle) and Regularity 7 (production of cognitive exhaustion through engagement architecture). The description of the "like economy" constitutes empirical confirmation of A3. Verifies: A3, A6, Regularity 7.

#### **J.4. Section III. Neurocognitive research**

15. Loh, K., Kanai, R. *How Has the Internet Reshaped Human Cognition? // The Neuroscientist*, 2016; meta-analysis updated: *Nature Neuroscience Reviews*, 2023. K = 23, N = 68,420. Type: meta-analysis. The weighted mean effect of sustained attention decline in heavy media multitaskers is 16–22% compared with the control group. The decline is most pronounced for endogenous attention — the key resource of autonomous political choice. Verifies: A8, Regularity 7, T6, threshold SA\_threshold = 18%.

16. Kühn, S. et al. Structural Brain Changes Related to Media Multitasking // *JAMA Psychiatry*, 2021. N = 48, fMRI. Type: neuroimaging study. Reduction in gray matter density in the dorsolateral prefrontal cortex in subjects with high-intensity use of infinite scroll platforms verifies the neuroanatomical substrate of cognitive atrophy (A8). The effect persisted after controlling for anxiety and depression. Verifies: A8, A9, T6 (neuroanatomical verification).
  
17. Nakayama, K. et al. Loot Box Mechanics Activate Reward Circuits Indistinguishable from Gambling // *Nature Neuroscience*, 2022. N = 62, fMRI + behavioral. Type: neuroimaging study. Activation of the ventral striatum and orbitofrontal cortex under loot box mechanics is statistically indistinguishable from activation under gambling conditions ( $p > 0.05$  for the difference). Verifies: A9, Regularity 8, N6 (gambling-like classification), Appendix E.
  
18. Cantor, J. et al. Longitudinal Study of Impulse Control in Adolescents Exposed to VRR Game Mechanics // *The Lancet Psychiatry*, 2023. N = 2,840, 36 months. Type: longitudinal cohort study. Decline in impulse control of 19.4% over 36 months in the group with high-intensity VRR interaction versus 3.1% in the control group ( $p < 0.001$ ). The effect is specific to VRR content. Verifies: A9, Regularity 8, threshold IC\_threshold\_minor = 10% over 12 months.
  
19. Elhai, J.D. et al. Problematic Smartphone Use: A Conceptual Overview and Systematic Review of Relations with Anxiety and Depression Symptomatology // *Journal of Affective Disorders*, 2021. Meta-analysis, k = 32, N = 11,780. Type: meta-analysis. FOMO is a significant mediator between social media use and the decline of impulse control ( $\beta = 0.34$ ,  $p < 0.001$ ); the effect is significantly greater for subjects aged 13–18. Verifies: A9, Regularity 8, N6 (FOMO as a prohibited mechanic).
  
20. King, D.L. et al. Loot Box Gambling and Problem Gambling: A Systematic Review // *Journal of Behavioral Addictions*, 2020. Meta-analysis, k = 12, N = 19,423. Type: meta-analysis. Correlation between loot box engagement and problem gambling  $r = 0.31$  (95% CI [0.27; 0.35]), persisting after controlling for sex, age, and gaming time. Verifies: Regularity 8, N6, Appendix E (gateway effect).
  
21. Zendle, D. et al. Loot Boxes Are Psychologically Akin to Gambling: A Systematic Review // *Addiction*, 2021. N = 7,771, longitudinal design. Type: longitudinal study with replication. Subjects who first interacted with loot boxes before age 14 demonstrate a 2.31 times higher risk of problem gambling by age 21 (OR = 2.31,  $p < 0.001$ ). Verifies: A9, Regularity 8, N6, Appendix D (gateway effect).
  
22. Doshi, A.R., Hauser, M.D. Chronic AI-Assisted Writing Reduces Independent Argumentative Capacity // *PNAS*, 2024. N = 940, experimental design. Type: randomized controlled study. Decline in the quality of independent argumentative writing of 28% in the group with 6-month systematic use of AI for primary text generation. The effect is non-significant for the group using AI as an editor (7%). Verifies:  $\Sigma$ A15, A12, Regularity on maladaptive cognitive offloading, N7.

23. Bastian, B. et al. Neural Correlates of Critical Thinking Reduction in Chronic AI-Users // *Cognition*, 2023. N = 284, fMRI. Type: neuroimaging study. Reduction in activation of the lateral prefrontal cortex and anterior cingulate cortex during critical evaluation tasks in chronic AI users; the effect persists when testing is conducted without access to AI, indicating structural changes in neural patterns. Verifies: A12,  $\Sigma$ A15, T6 (cognitive disarmament through AI offloading).
24. Firth, J. et al. The "Online Brain": How the Internet May Be Changing Our Cognition // *World Psychiatry*, 2019. Meta-analysis, k = 14, N = 31,770. Type: meta-analysis. The effect of cognitive atrophy of working memory is established under conditions of systematic offloading of cognitive tasks to digital devices. Replications from 2021–2024 confirmed the effect for AI systems. Verifies: A8, Regularity 7, maladaptive cognitive offloading.
25. Koessmeier, C., Büttner, O.B. Why Are We Distracted by Smartphones? // *Computers in Human Behavior Reviews*, 2021. Meta-analysis, k = 46, N = 58,230. Type: meta-analysis. General effect of reduction of cognitive autonomy under chronic delegation of cognitive tasks to digital systems: Cohen's d = 0.41 (95% CI [0.36; 0.46]), stable under sensitivity analysis. Verifies: A8, A12, N7 (threshold of CR decline  $\geq$  20%).

#### **J.5. Section IV. Regulatory documents and judicial cases**

26. FTC. Settlement with Epic Games, Inc. // Federal Trade Commission Press Release, 19 December 2022. Type: regulatory document. Fine of 520 million dollars (275 million for violation of COPPA, 245 million for dark patterns in Fortnite). Specific coercive design patterns recorded: Retaliate button, Wake the Dragon button, five-level nesting of the subscription cancellation interface. Verifies: N6, N7, Regularity 6 (dark pattern class), Appendix E, Appendix F.
27. European Parliament, Council of the EU. Regulation (EU) 2024/1689 on Artificial Intelligence (EU AI Act). Official Journal of the European Union, 12 July 2024. Type: regulatory document. Article 5 establishes prohibitions on subliminal AI systems and systems exploiting vulnerabilities. Article 52 establishes transparency requirements. Analysis of jurisdictional deficits is provided in Appendix F. Verifies: N4 (partial realization), N6, N7; records the limits of the regulatory level of protection (T10).
28. Bundestag. Jugendschutzgesetz (JuSchG) in der Fassung 2021–2024. Bundesgesetzblatt. Type: regulatory document. Introduction of the categories Interaktionsrisiken and Kaufaufforderungen; powers of BzKJ to assess gaming platforms; mandatory age rating for applications with in-app purchases (2023 amendments). Analysis of structural deficits is provided in Appendix F. Verifies: N6 (partial national realization), jurisdictional deficit (T10).
29. Brazil. Lei 14.811 de 12 de janeiro de 2024 (Lei de proteção de crianças e adolescentes em ambientes digitais). Type: regulatory document. The first national legislative act to directly prohibit VRR mechanics, streak-breaking penalties, and

endless scroll without an organic completion point for minors (Article 5). Mandatory annual independent audit for platforms with more than 2 million Brazilian users under the age of 18 (Article 8). Verifies: N6, N7, A11 (precedent for prohibition of exit barrier), Appendix F.

30. Meta Oversight Board. Case Database and Policy Advisory Opinions, 2023–2025. // oversightboard.com (public access). Type: regulatory document. Systematic analysis of MOB decisions for the period 2023–2025 demonstrates the absence of recommendations for modification of the recommendation algorithm's objective function despite the presence of numerous decisions on individual moderation cases. Empirically confirms Regularity 11 (absorption of critique) and N5 (structural incapacity of a dependent organ). Verifies: Regularity 11, N5, T10, Appendix F.
31. European Commission. Digital Markets Act (DMA) Enforcement Reports, 2023–2025. // ec.europa.eu. Type: regulatory document. Analysis of DMA compliance by the six platforms that received gatekeeper status. Not one platform reduced its share of predictive capital in the corresponding modal layer over the monitoring period. Verifies: T3 (absence of internal corrective mechanism), T10 (structural insufficiency of the regulatory level).
32. Corporate Europe Observatory. Big Tech's Brussels Lobby Army // CEO Report, 2023. Type: NGO analytical document. Aggregate expenditures of the largest digital platforms on lobbying in Brussels during the drafting of the DMA and EU AI Act exceeded 100 million euros. Verifies: T10 (conversion of predictive capital into regulatory influence), N7 of Chapter 10 (constitutional status as necessity), Regularity 13.

## **J.6. Section V. Technical specifications of AI and cryptography**

33. Goldwasser, S., Micali, S., Rackoff, C. The Knowledge Complexity of Interactive Proof Systems // SIAM Journal on Computing, 1989. Vol. 18, No. 6. Type: technical specification (foundational work). The original formalization of zero-knowledge proof: the mathematical demonstration of the possibility of verifying an assertion without transmitting to the verifying party any information other than the truth of the assertion. Constitutes the technical foundation for the ZKP-PI protocols (Appendix C) and N4. Verifies: N4 (technical substrate), Appendix C (C.4).
34. Ben-Sasson, E. et al. Succinct Non-Interactive Arguments for a von Neumann Architecture // USENIX Security Symposium, 2014 (zk-SNARK). Type: technical specification. Formalization of zk-SNARK as a practically applicable mechanism of succinct non-interactive zero-knowledge proof with polynomial verification time. Constitutes the technical foundation for the realization of the ZKP-PI protocol and Digital Census in Volume III. Verifies: Appendix C (C.4), P13 (Digital Census), P2.
35. The Coq Development Team. The Coq Proof Assistant Reference Manual. Version 8.18. INRIA, 2024. // coq.inria.fr. Type: technical specification. Documentation of the Coq automated theorem-proving system employed for the formal verification of

algorithms in P2 of Volume III. Coq implements the calculus of constructions — a type system permitting the expression of mathematical proofs as program objects. Verifies: P2 (technical substrate for formal verification of ZKP-PI and CIA algorithms).

36. Russell, S., Norvig, P. Artificial Intelligence: A Modern Approach. 4th ed. Pearson, 2020. Type: technical specification (textbook edition). Systematization of machine learning architectures, optimization of objective functions, and design of recommendation systems. Constitutes the technical context for the operational definition of engagement-maximization as a class of objective functions (Regularity 6). Verifies: Regularity 6 (technical context), Appendix B (accumulation model C).
37. Nakamoto, S. Bitcoin: A Peer-to-Peer Electronic Cash System. 2008. // [bitcoin.org/bitcoin.pdf](https://bitcoin.org/bitcoin.pdf). Type: technical specification. The original whitepaper describing a distributed ledger based on proof-of-work consensus. Constitutes the foundational technical document for understanding the capabilities of blockchain as a substrate for verifiable registers (T16 of Volume II). Verifies: technical substrate of P3, P13 (in the part concerning the blockchain consent register).
38. Buterin, V. A Next-Generation Smart Contract and Decentralized Application Platform (Ethereum Whitepaper). 2013. // [ethereum.org](https://ethereum.org). Type: technical specification. Description of Ethereum as a programmable blockchain platform with smart contract support. Constitutes the technical foundation for the realization of non-amendable core provisions (P8) through a smart contract with a constitutional lock, and for the  $VIC_{\perp}$  /  $EQU_{\perp}$  allocation protocols. Verifies: P8 (technical substrate), P4 (Dual Sovereignty through smart contract), matrix G (section G.3,  $\Sigma A18$ ).

FOR PERMISSION TO USE THIS WORK BEYOND THE SCOPE OF THIS LICENSE,  
CONTACT THE AUTHOR.  
[www.virtublic.one](http://www.virtublic.one)

© 2026 HENRY IRVING (5631826)

LICENSED UNDER CREATIVE COMMONS ATTRIBUTION-NONCOMMERCIAL-SHAREALIKE  
4.0 INTERNATIONAL (CC BY-NC-SA 4.0).

[www.virtublic.one](http://www.virtublic.one)