Compression After Expansion: On Skill, Misattribution, and the Geometry of Work

Flyxion

Independent Researcher

April 13, 2026

Abstract

This essay develops a unified framework for understanding the systematic misperception of difficulty in human skill, technical production, and social representation. The central thesis is that apparent simplicity in completed artifacts—whether mathematical proofs, constructed objects, software systems, or natural language outputs—is the historical residue of prior expansion, iteration, and structural accumulation, and not an intrinsic property of the artifact itself. Contemporary technological and economic systems compound this misperception by presenting compressed endpoints as primary forms, thereby obscuring the generative processes that produced them. We formalize this asymmetry through two interlocking theoretical structures: a general account of expansion and compression as operations on configuration spaces, grounded in an entropic asymmetry between construction and destruction; and Aspect Relegation Theory, a cognitive model in which skill is understood as compressed structure encoded through the progressive withdrawal of attention from stabilized sub-processes. Together these frameworks explain a range of phenomena spanning cognitive science, labor economics, media theory, and epistemology, and ground a normative critique of representations that sever compressed outputs from their generative histories.

Contents

1. Introduction: The Illusion of Effortless Form

There is a persistent structural illusion at work in the way human beings perceive completed artifacts. A finished building, a published proof, a released piece of software, a polished essay—each presents itself as a coherent, apparently simple object whose internal structure invites no particular reflection on the process by which it came to be. Observers encountering such objects tend, with remarkable consistency, to draw inferences about the difficulty of producing them that are calibrated to the apparent simplicity of the result rather than to the generative complexity of the process. This inference is systematically false, and understanding why it is false is the primary aim of this essay.

The illusion has several reinforcing sources. Compression is a constitutive feature of representation: the point of producing a final form is precisely to eliminate from view the intermediate steps, failed attempts, and structural discoveries that preceded it. This is not a contingent feature of bad communication but a functional requirement of efficient representation. A proof that displayed every exploratory conjecture, failed lemma, and notational revision through which it was produced would be unreadable. A building that preserved every scaffolding stage would be uninhabitable. Compression is, in this sense, necessary. But it is also epistemically misleading when the compressed form is mistaken for the primary form, and when the prior expansion is forgotten or invisible.

A second reinforcing source is cognitive. Human beings develop skills through a process that renders previously effortful operations automatic, relocating them beneath the threshold of conscious attention. Expert performance therefore appears effortless because it is effortless—but this effortlessness is the product of accumulated labor, not its absence. The experienced constructor who walks through a derelict building and immediately perceives a tractable sequence of operations is not seeing something simple; she is perceiving a high-entropy configuration through a lens of internalized low-entropy structure. The novice who sees chaos is, in a precise sense, correct about the current state of the building. What they lack is the generative pathway.

A third source is institutional and economic. Modern media and commercial systems are organized around the representation of outcomes rather than processes. Startup narratives, viral content, minimalist design aesthetics, and compressed pedagogical formats all function by presenting relegated forms while suppressing the construction histories that produced them. This is not always deliberate de-

ception, but it reliably produces systematic distortions in how difficulty, effort, and value are perceived and distributed.

The core claim of this essay may be stated directly: *simplicity is not primitive but derived*. What appears simple is typically the result of prior, invisible complexity that has been successfully compressed. This compression has a structure—it presupposes expansion, it is achieved through stabilization, and it is fragile in ways that the compressed form does not reveal. Understanding this structure is not merely an epistemological exercise. It has consequences for how we understand labor, pedagogy, creativity, and the ethics of representation.

2. Expansion Before Compression: A Formal Treatment

2.1. Preliminaries

We introduce a minimal formal vocabulary sufficient to capture the structural asymmetry between construction and compression.

Definition 2.1. Let a *system* $\mathcal{S}$ be defined as a configuration of components subject to a constraint set $\mathcal{C}$. A configuration s is said to be *coherent* if and only if $s \models \mathcal{C}$.

Let Ω denote the space of all possible configurations of $\mathcal{S}$. Define the subset of valid configurations:

$$\Omega_{\text{valid}} = \{s \in \Omega \mid s \models \mathcal{C}\}$$

and its complement:

$$\Omega_{\text{invalid}} = \Omega \setminus \Omega_{\text{valid}}.$$

In all nontrivial systems of interest, the valid subset is vanishingly small relative to the full configuration space:

$$|\Omega_{\text{valid}}| \ll |\Omega|.$$

This inequality encodes the fundamental rarity of coherent structure. It is not a contingent empirical fact but a structural feature of constrained systems: the more constraints a system must satisfy, the smaller the fraction of configurations that qualify as valid.

2.2. Expansion as Exploration

Definition 2.2. *Expansion* is a process E acting on an initial configuration $s_0 \in \Omega$, generating a sequence:

$$E(s_0) = \{s_1, s_2, \ldots, s_n\}$$

where each s_i is obtained through transformation, trial, or variation.

Expansion is characterized by traversal of configuration space, explicit evaluation of constraint satisfaction, discovery of failure modes, and articulation of dependencies. It is therefore not merely generative but epistemic: expansion reveals the structure of Ω_{valid} through systematic interaction with $\mathcal{C}$. This is the process through which invariants are discovered, dependency structures are mapped, and failure modes are identified and resolved.

2.3. Compression as Projection

Definition 2.3. *Compression* is a mapping:

$$C : \Omega_{\text{valid}} \to \Sigma$$

where Σ is a lower-dimensional representational space.

Compression eliminates redundancy while preserving constraint satisfaction. It encodes the invariants discovered during expansion and suppresses intermediate derivations. Typical instances include refactored code, minimal mathematical proofs, concise explanatory prose, and completed architectural forms. Compression is therefore not discovery but representation: it presupposes that the structure it encodes has already been found.

2.4. The Dependency of Compression on Expansion

Proposition 2.4. *No valid compression mapping $C : \Omega_{\text{valid}} \to \Sigma$ can be constructed without prior identification of the invariant structure of Ω_{valid}, which in turn requires expansion.*

Proof sketch. A compression mapping removes components deemed redundant while preserving correctness. Determining redundancy requires knowledge of which elements are invariant under permissible transformations—that is, which structural features are necessary for constraint satisfaction and which are incidental.

Such invariants are not given a priori. They are identified through comparison of multiple candidate configurations, observation of which variations produce constraint violations, and stabilization of the recurring structural features that survive this variation.

All of these operations require traversal of Ω, with particular attention to the boundary between Ω_{valid} and Ω_{invalid}. This traversal constitutes expansion. Therefore, compression presupposes expansion. □

The practical corollary is direct: one cannot produce a valid compressed representation of something one has not first explored at greater length. This applies equally to mathematical exposition, software architecture, physical construction, and written argument. The simplified final form is not a starting point but an achievement.

2.5. Entropic Asymmetry

We now formalize the asymmetry between construction and disruption that underlies much of the argument of this essay.

Definition 2.5. Let *entropy* be defined informally as:

$$H(s) \propto \log |\Omega_{\text{reachable}}(s)|$$

where $\Omega_{\text{reachable}}(s)$ is the set of configurations accessible from state s under the available operations.

Then construction corresponds to the transition $\Omega \to \Omega_{\text{valid}}$, a movement toward lower entropy, while destruction corresponds to $\Omega_{\text{valid}} \to \Omega$, a movement toward higher entropy.

Proposition 2.6 (Asymmetry of Operations)**.** *The number of operations that increase entropy exceeds the number of operations that decrease it under any fixed constraint set $\mathcal{C}$.*

Corollary 2.7. *Destructive and evaluative processes are generically easier than constructive ones.*

The interpretation is direct. A coherent system occupies a narrow region of configuration space. Any perturbation that violates even a single constraint suffices to exit this region. By contrast, constructing or maintaining coherence requires selecting from a highly constrained subset of operations—those that preserve satisfaction of all constraints in $\mathcal{C}$ simultaneously. The combinatorial asymmetry is

not a matter of degree but of structural geometry: valid configurations form a low-dimensional manifold within a vastly higher-dimensional ambient space.

2.6. Fragility of Compressed Representations

Since compression encodes a narrow basin within Ω, compressed representations are structurally fragile in a precise sense. Small perturbations can produce invalid states; omitted intermediate structure prevents reconstruction; and interpretability depends on prior familiarity with the expansion process that generated the compressed form.

We may therefore state the central structural principle of this section:

> *Expansion is required to discover coherence; compression encodes it. Construction is constrained entropy reduction, while destruction is unconstrained entropy increase. The asymmetry between these two processes is not contingent but structural, and it underlies the systematic misperception of difficulty wherever compressed artifacts are encountered without their generative histories.*

3. Skill as Internalized Structure

3.1. The Standard Account and Its Inadequacy

The standard account of expertise treats skill as a matter of accumulated knowledge and trained performance. This is not wrong, but it is underspecified in a way that prevents it from explaining the most important features of expert cognition: the disappearance of effort, the compression of explanation, and the systematic gap between expert and novice perception of the same situation.

A more adequate account requires attention to what happens to the structure of a task as it is learned, not merely to the content of what is known.

3.2. Aspect Relegation Theory: A Formal Model of Cognitive Compression

We introduce Aspect Relegation Theory as a cognitive analogue of the expansion–compression framework developed in the previous section.

Definition 3.1. Let a *task* T be decomposable into a finite set of aspects:

$$T = \{a_1, a_2, \ldots, a_n\}$$

where each aspect a_i corresponds to a constraint, sub-operation, or dependency required for successful execution of T.

Definition 3.2. At any time t, the *active attention set* is:

$$A(t) \subseteq T$$

representing the subset of aspects currently under conscious monitoring. We assume a bounded attentional capacity: $|A(t)| \leq k$ for some fixed cognitive limit k.

Definition 3.3. An aspect a_i is said to be *stabilized* if repeated execution satisfies its associated constraints without requiring explicit monitoring.

Definition 3.4. The *relegation operator* R is defined as:

$$R : A(t) \to A(t+1)$$

such that:

$$A(t+1) = A(t) \setminus \{a_i \in A(t) \mid a_i \text{ is stabilized}\}$$

Relegation is thus the progressive withdrawal of attention from aspects that have achieved reliable, constraint-satisfying execution. The process defines learning as a trajectory:

$$A(0) = T \xrightarrow{\text{practice, correction, stabilization}} A(t) \to \emptyset.$$

Initial performance of a task requires maximal attentional engagement: $A(0) = T$. Over time, through repetition and correction, $A(t)$ contracts toward a minimal residual, which may approach the empty set or a small subset corresponding to higher-level coordination.

3.3. System 1 as Relegated System 2

This framework yields a direct reinterpretation of the distinction between deliberative and intuitive cognition.

Proposition 3.5. *What is commonly described as fast, intuitive cognition (System 1) is the residual execution of previously stabilized and relegated aspects of slow, deliberative cognition (System 2).*

The argument follows directly from the definitions. Early execution requires full decomposition and conscious monitoring of the task. Stabilized aspects are removed from attention via the relegation operator. Execution persists despite

reduced attentional involvement because the stabilized aspects continue to run, governed by internalized constraint structures rather than explicit monitoring.

The standard account, which treats System 1 and System 2 as distinct cognitive systems with different functional architectures, is therefore misleading. System 1 is not a different kind of thinking; it is System 2 at a later stage of its own development. The difference is temporal, not architectural.

Corollary 3.6. *System 1 cognition is not primitive but derived. What appears as immediate intuition is the residue of prior deliberation.*

3.4. Hierarchical Attention and Compositional Depth

Because attention is bounded, relegation does more than make tasks feel easier. It enables increasing compositional complexity by freeing attentional capacity for higher-level structure.

As lower-level aspects are removed from $A(t)$, capacity becomes available for higher-order coordination, abstraction, and integration across sub-tasks. This produces a hierarchical attentional structure: low-level operations become automated, mid-level structures are partially attended, and high-level planning is actively monitored. Skill is therefore not merely efficiency but depth of compositional control—the capacity to operate at a higher level of abstraction because lower levels have been successfully relegated.

3.5. Skill as Compressed Structure

The parallel with the formal framework of the previous section is now explicit.

Structural Process	**Cognitive Process**
Expansion	Learning and exploration
Compression	Relegation
Invariants	Stabilized aspects
Minimal form	Automatic execution

Relegation is cognitive compression. The stabilized aspects that have been relegated encode the structural invariants discovered through prior expansion. Skill is compressed structure encoded through stabilized execution. And the phenomenological experience of ease corresponds precisely to $|A(t)| \ll |T|$—not to any reduction in the underlying complexity of the task, but to the successful internalization of that complexity beneath the threshold of attention.

This has an important implication. Ease is a function of attentional load, not structural complexity. The fact that a task feels easy to an expert is not evidence that the task is simple. It is evidence that the expert has successfully completed the expansion phase, identified the structural invariants, and relegated them to automatic execution. The complexity has not disappeared; it has been internalized.

3.6. Irreversibility and Reconstructive Loss

Relegation introduces a partial loss of reconstructability. Once aspects have been removed from attention, intermediate steps may no longer be explicitly representable, explanations may become compressed or incomplete, and reconstruction of the full expansion history requires re-expansion. Expert knowledge is often operationally accessible but descriptively compressed: the expert can do the thing but may struggle to explain how.

This irreversibility is not a failure but a structural feature. It explains why experts often give advice that is correct for them but useless for novices—the relegated aspects are no longer visible to the expert, so their advice skips steps that are not yet automated for the novice. The compressed explanation presupposes an expansion the novice has not yet performed.

4. The Misreading of Outputs

4.1. Compressed Outputs and Shifted Burden

The structural analysis of the previous sections has a direct implication for the interpretation of compressed outputs in any domain. When a compressed representation is presented without its generative history, the interpretive burden shifts from the producer to the receiver. The producer has internalized the expansion; the receiver has not. The result is a systematic asymmetry in accessibility.

This asymmetry has different consequences depending on context. In mathematics, an elegant proof may be genuinely illuminating for a reader who has worked through the domain, because the compression encodes structure the reader can reconstruct. The same proof may be opaque to a reader who lacks this background. The compression has not failed; it has simply been misaddressed.

In commercial and media contexts, the situation is more troubling. Compressed outputs are frequently presented as primary forms, not because the audience has

the background to reconstruct them, but because the compression serves a rhetorical function: it makes production appear simple, and simplicity signals accessibility. Consider the genre of compressed instruction—"just learn to code," "just start writing," "just build the product"—which presents relegated forms as if they required no prior expansion. Each of these phrases describes a valid terminal state while suppressing the entire trajectory required to reach it.

4.2. Valid and Invalid Compression

There is an important distinction to be drawn between compression that is reconstructible and compression that is parasitic on undisclosed prior knowledge.

Valid compression encodes structure that an appropriately prepared receiver can, in principle, reconstruct. A minimal proof, a well-refactored codebase, a concise argument—these are examples of valid compression because the structural invariants they encode are recoverable, given sufficient background. The compression serves efficiency without destroying accessibility.

Invalid compression, by contrast, presents a compressed form that depends on unexplicated prior knowledge and is presented as if it required none. This is the structure of what might be called *epistemic outsourcing*: the appearance of simplicity is achieved not by encoding structural invariants efficiently but by silently delegating the expansion work to the receiver, who is not equipped to perform it and is not told that they must.

The distinction matters because invalid compression produces false impressions of difficulty distributions. If the production of something is presented in its relegated form, observers infer that production is as easy as consumption. This inference drives systematic misattributions of effort, value, and competence.

5. Fundamental Attribution Error and the Erasure of Process

5.1. The Attribution Framework

The misperception analyzed in the previous section is amplified by a well-documented feature of social cognition. Lee Ross's fundamental attribution error describes the systematic tendency to attribute observed outcomes to dispositional properties of agents—their talent, intelligence, or character—while underweighting situational and processual factors that contributed to those outcomes.

Applied to the domain of skill and production, the error takes a specific form. When we observe a successful outcome, we attribute it to the intrinsic qualities of the producer rather than to the structural conditions, accumulated labor, and sequential expansion that made the outcome possible. When we observe a failed or difficult attempt, we attribute it to the producer's inadequacy rather than to the genuine structural difficulty of the task.

This attribution pattern is not random. It is driven by the availability of information. Outcomes are visible; processes are not. The compressed artifact is present; the expansion history has been suppressed. Given this informational asymmetry, dispositional attribution is not irrational—it is a reasonable response to the available evidence. But the evidence is systematically incomplete, and the resulting attributions are systematically distorted.

5.2. Media Amplification

Modern media ecosystems amplify this bias in several reinforcing ways. Selection effects ensure that the outcomes most widely circulated are those that appear most dramatically discontinuous with ordinary effort—the overnight success, the brilliant solution, the effortlessly productive artist. Narrative conventions frame these outcomes as expressions of individual quality rather than structural process. And the compression inherent in all media representation strips away the expansion histories that would reveal the underlying structure.

The result is a collectively maintained illusion about the distribution of effort. Observed successes appear to require little work; invisible failures disappear from the distribution entirely. The sample of outcomes that reaches public attention is not a representative sample of the underlying process, and treating it as such produces badly miscalibrated expectations.

6. Probabilistic Asymmetry and Survivor Narratives

6.1. The Structure of Success Distributions

A related distortion operates at the level of population-level inference. In many domains of skilled production—entertainment, athletics, entrepreneurship, academic research, artistic practice—outcome distributions are heavily right-skewed. A small number of participants achieve highly visible success; the large majority do not achieve comparable recognition, regardless of comparable effort.

This structure has a straightforward consequence for inference from observed outcomes. The successes that are visible are a non-representative sample of the underlying population of attempts. They are selected precisely because they are extreme, and they are circulated precisely because they are legible as success. The process of selection and circulation ensures that what is visible is not what is typical.

6.2. Narrative Selection Bias

Philip K. Dick, whose eventual canonical status belies a career of economic precarity and repeated rejection, is one illustration of this structure among many. The canonical form of such a life, reconstructed retrospectively, reads as a narrative of difficult genius eventually recognized. This reconstruction suppresses the contingency, the selection effects, and the many equally capable individuals for whom recognition never arrived. The survivor's story is not a guide to the distribution from which it was selected.

The practical consequence is a form of visibility filtering that makes success appear more accessible and more meritocratic than it is. This is not simply an epistemological inconvenience. It functions as a structural mechanism for undervaluing labor that does not produce visible, scalable, or circulated outputs.

7. Labor, Essential Work, and Recognition Failure

7.1. The Visibility Gap

The preceding analysis converges on a structural asymmetry in how different categories of labor are perceived and rewarded. We may distinguish two broad categories.

Essential work is labor that is infrastructural, repeatable, and stabilizing. It includes teaching, maintenance, care, construction, food production, and the countless forms of activity that sustain the conditions under which visible production becomes possible. This work is typically invisible in the sense that its outputs are the absence of failure rather than the presence of success. Clean buildings, functioning infrastructure, educated children, maintained systems—these do not announce themselves as achievements. They are the background against which other achievements occur.

Visible work, by contrast, is rare, amplified, and rewarded in proportion to its

visibility. It includes performances, publications, launches, and other events whose outputs are legible as discrete, attributable achievements. This work benefits from all the amplification mechanisms described in previous sections.

7.2. Temporal Misalignment

Recognition is also subject to a temporal misalignment that compounds the visibility gap. The labor that stabilizes conditions for future achievement may be recognized—if at all—long after it was performed, and often only after some visible subsequent event makes its structural role legible. The teacher whose student later achieves prominence is recognized, if at all, through the lens of that student's success. The posthumous recognition of artists, theorists, and inventors follows the same structure: the expansion work is recognized only once its compression has become visible in subsequent outputs.

This temporal misalignment is not merely unfair. It is structurally distorting, because it feeds back into the attribution errors described earlier. If recognition consistently arrives late and derivatively, the cultural model of how achievement is produced is systematically biased toward the visible compression and away from the prior expansion.

8. Wealth, Abstraction, and Constraint Blindness

8.1. Delegation and Loss of Constraint Visibility

Wealth confers the capacity to delegate constraint management. The wealthy individual does not personally manage the supply chains, maintenance schedules, infrastructure dependencies, and coordination problems that structure everyday life; these are managed by others, on their behalf, and rendered invisible. This delegation is one of the primary functional benefits of wealth, and it is not inherently problematic. What is problematic is its epistemic consequence.

When constraint management is delegated, the experience of constraints disappears. The delegating agent interacts with a world that appears to present itself in already-compressed form: problems that appear to have simple solutions, processes that appear to require only a decision, systems that appear to be responsive to intention alone. The expansion phase has been performed by others and is invisible.

8.2. Constraint Blindness and Compressed Advice

The consequence is what we might call *constraint blindness*: a systematic incapacity to perceive the structural conditions that would make apparent solutions non-trivial. This produces the familiar genre of compressed advice from those who have not personally performed the expansion: "just do this," "simply approach it differently," "the solution is obvious." These statements are not always dishonest. They reflect a genuine perceptual state in which the constraints that make the task difficult are not visible to the speaker.

The advice is nonetheless harmful when treated as actionable guidance, because it presents a compressed form that presupposes an expansion the listener must perform but is not told about. The speaker has access to a world in which the constraints are managed; the listener does not.

9. Deontological Stability Under Skewed Distributions

9.1. Consequentialism and Distributional Distortion

The probabilistic asymmetries described in Section 6 create a specific problem for consequentialist ethics. If outcomes are distributed across heavily right-skewed distributions, and if only the extreme successes are visible, then the expected value calculation an agent might perform using observed outcomes is badly miscalibrated. The actions that appear, from the observed distribution, to produce the best outcomes are not the actions that produce the best expected outcomes across the full distribution.

This distortion is not a contingent feature of imperfect information. It is structural: the selection mechanisms that make certain outcomes visible are precisely the mechanisms that produce the distortion. A consequentialist agent reasoning from visible outcomes is therefore systematically misled by the very information environment that consequentialism requires them to use.

9.2. Kantian Stability

Immanuel Kant's categorical imperative—the requirement that one act only on maxims one could consistently will to be universal laws—provides a different foundation for practical decision-making, one that is structurally robust to the distributional distortions described above. Actions are evaluated not by their

consequences in a skewed observable distribution but by their compatibility with universal principles that abstract over distribution.

Applied to the domains of labor and practice, this produces a different orientation. Cleaning, studying, building, maintaining, and teaching are justified not because they produce visible rewards in the observed distribution but because they satisfy principles that one can consistently endorse regardless of where they fall in that distribution. This is not an ascetic recommendation but a structural observation: deontological commitments stabilize behavior under conditions where consequentialist reasoning is corrupted by selection effects.

10. Incomplete Systems and the Perception of Potential

10.1. Expert and Novice Perception

The divergence between expert and novice perception of incomplete or disordered systems is one of the most striking practical implications of the framework developed here. An expert surveying a derelict building perceives not chaos but a tractable sequence of operations: assess structural integrity, identify constraint violations, decompose into sequenced sub-tasks, estimate materials and labor, project a path to coherence. A novice surveying the same building perceives undifferentiated disorder whose resolution is not apparent.

Both perceptions are, in a precise sense, accurate. The building is disordered; the novice is correct that its current configuration is far from valid. What the novice lacks is not information about the building's current state but a generative pathway—a sequence of operations that traces a trajectory from the current configuration to one in Ω_{valid}.

10.2. Tractability versus Coherence

The divergence can be clarified through the entropy asymmetry introduced in Section 2. The observer who lacks a decomposition of the system encounters a high-entropy configuration: undifferentiated disorder, ambiguous dependencies, and indeterminate pathways to completion. The expert observer implicitly maps this configuration onto a sequence of entropy-reducing operations. What appears as chaos to one is already partially ordered to the other.

Crucially, the expert does not perceive the system as complete but as tractable. The difference lies not in optimism but in the availability of a generative pathway.

The novice lacks such a pathway and therefore experiences the system as irrecoverable. From the novice's perspective, this experience is not irrational—without a pathway, the system is practically irrecoverable, for them.

11. Construction, Destruction, and Entropic Asymmetry

11.1. A General Principle

The analysis of expert and novice perception leads naturally to a general principle that has been implicit throughout this essay but deserves explicit statement.

Proposition 11.1 (Asymmetry of Construction and Destruction)**.** *To construct a coherent system is to locate and maintain a low-entropy configuration within a high-dimensional space of possibilities. To break or evaluate a system is to identify any perturbation that moves it out of this narrow basin of stability. These two activities operate under fundamentally different constraints.*

Construction requires anticipating interactions, resolving dependencies, and preserving coherence across time. It must satisfy global coherence conditions. Destruction, by contrast, need only satisfy local failure conditions: identify any perturbation that moves the system from Ω_{valid} to Ω_{invalid}.

Corollary 11.2. *The ease of breaking or criticizing a system is not evidence of the ease of constructing it. The two activities sample from different regions of the space of operations.*

11.2. Misattribution from Evaluative Ease

The persistent misattribution described throughout this essay has a precise source in this asymmetry. Agents whose primary interaction with systems is evaluative—critics, reviewers, inspectors, users—encounter systems in their already-compressed forms and interact with them through the low-dimensional task of identifying failures. The apparent ease of this task is mistaken for evidence that the construction task was correspondingly easy. But construction sampled from the constrained manifold of coherence, while evaluation samples from the effectively unbounded space of disorder.

The experiential consequence is that evaluative agents may systematically underestimate the difficulty of generative work. This misalignment is not a moral failing but a structural consequence of the epistemic position from which evaluation is conducted. Understanding it does not eliminate it, but it does enable a more

accurate model of the distribution of labor.

12. Writing, Resistance, and the Generative Necessity of Expansion

12.1. Writing as Expansion

The analysis of this essay is not merely a theoretical observation about cognition and labor in general. It applies with particular directness to the activity of writing. A writer who attempts to produce the compressed form—the elegant, concise, finalized argument—without first performing the expansion will consistently fail to produce it, not because they lack talent but because compression is structurally dependent on prior expansion.

Drafting is expansion: the articulation of candidate structures, the discovery of dependencies, the identification of failure modes in an argument, the establishment of which elements are invariant and which are redundant. Revision is stabilization: the correction of constraint violations, the consolidation of surviving structure, the reduction of redundancy. The final form is compression: the encoding of invariants discovered through expansion and stabilization.

The common advice to "write simply" or to "get to the point" is not wrong as a description of the desirable endpoint, but it is misleading as an account of the process by which that endpoint is reached. One cannot begin with simplicity; one can only arrive at it.

12.2. The Time Dimension of Work

This has a general implication about the time dimension of generative work. There is no shortcut through expansion. The dependencies must be discovered, and discovery requires traversal. The constraints must be stabilized, and stabilization requires repetition and correction. The compression can only encode what expansion has found.

This is why work that appears to have been produced effortlessly almost never was. The effortlessness is a property of the final representation, not of the process that produced it. And this is why the instruction to "just do it" fails systematically as practical guidance: it describes a terminal state while providing no pathway through the expansion required to reach it.

13. The Ethics of Compression

13.1. Representations and Reconstructibility

The arguments developed in previous sections ground a normative conclusion about the ethics of representation. Compression is not inherently problematic; it is functionally necessary and often genuinely clarifying. The problem arises when compression is deployed in ways that suppress the expansion history in ways that mislead observers about the nature of the underlying work.

A compressed representation severs the output from its generative pathway. This severance is sometimes appropriate: the reader of a finished proof does not need a complete record of every exploratory conjecture. But it becomes ethically problematic when the severance is used to misrepresent the difficulty of the generative process, to obscure the labor of those who performed the expansion, or to create the impression that the compressed form could be produced without prior expansion.

Definition 13.1. A compressed representation is *ethically compliant* if it preserves, in principle, the reconstructibility of the generative pathway that produced it, or if it accurately represents the expansion process that preceded the compression.

13.2. Entropic Misrepresentation

An additional normative constraint emerges from the entropic asymmetry of construction and destruction. Representations that foreground ease while omitting the entropy-reducing processes required to achieve it systematically mislead observers about the nature of work. Such representations conflate the accessibility of failure—which is high, because valid configurations are a small subset of the full configuration space—with the accessibility of success.

Ethical compression must therefore preserve not only semantic content but processual recoverability. A compressed representation should, in principle, allow reconstruction of the generative pathway that produced it. Where this is not possible, compression functions not as clarification but as obfuscation: it presents the low-entropy output of constrained construction as if it were a natural or effortless state.

13.3. Criteria for Ethical Representation

We may summarize the normative position in three criteria. First, transparency of process: representations of skilled output should, where possible, disclose or gesture toward the expansion processes that preceded them, rather than presenting only the compressed endpoint. Second, recoverability of structure: compressed representations should preserve the structural invariants they encode in a form that allows, at least in principle, reconstruction by an appropriately prepared receiver. Third, accuracy of effort distribution: representations should not systematically misrepresent the probability distribution over effort required to produce similar outcomes.

These criteria are not absolute; context matters. But they provide a framework for evaluating the ethics of compression in specific cases, and they ground a substantive critique of the many representational practices—in marketing, pedagogy, media, and technology—that systematically violate them.

14. Conclusion: Toward a Theory of Visible Structure

The argument of this essay may be summarized in four related claims.

First, simplicity is earned, not given. What appears simple in any artifact—mathematical, architectural, written, computational, or social—is the result of prior expansion, iteration, and structural accumulation. The expansion phase discovers invariants; the compression phase encodes them. Neither can substitute for the other, and the compressed form does not reveal the process that produced it.

Second, difficulty is relational, not intrinsic. A task is difficult relative to the expansion work that has or has not been performed, the aspects that have or have not been stabilized, and the attentional resources that are or are not available. What is effortless for an expert is effortless because they have already performed the expansion; it is not effortless because it is simple.

Third, misattribution arises from compressed representations. When artifacts are encountered in their compressed forms, without access to their generative histories, observers systematically misattribute the difficulty of production. This misattribution is not a moral failing but a structural consequence of the epistemic position from which compressed outputs are encountered.

Fourth, the misattribution has consequences. It distorts how labor is valued, how advice is given, how achievement is attributed, and how expectations are formed. It generates a class of representations—commercial, pedagogical, and journalistic—that present compressed endpoints as primary forms and systematically mislead about the effort distributions from which they were produced.

The correction proposed here is not a recommendation to make things more complicated or to celebrate difficulty for its own sake. It is a recommendation for epistemic honesty about the structure of generative processes: an acknowledgment that what appears simple is the result of prior, invisible complexity that has been successfully relegated; that construction is constrained entropy reduction while destruction is unconstrained entropy increase; and that the compressed artifact, presented without its expansion history, is always a partial and potentially misleading representation of the work that produced it.

Understanding this structure does not make expansion easier. But it does make the misperception of it harder to sustain.

A. Summary of Formal Results

For convenience, we collect the principal formal results established in the body of the essay.

Proposition 2.4 No valid compression mapping can be constructed without prior identification of the invariant structure of the valid configuration space, which in turn requires expansion.

Proposition 2.6 The number of operations that increase entropy exceeds those that decrease it under any fixed constraint set. Consequently, destructive and evaluative processes are generically easier than constructive ones.

Proposition 3.5 What is described as intuitive cognition (System 1) is the residual execution of previously stabilized and relegated aspects of deliberative cognition (System 2).

Proposition 11.1 Construction satisfies global coherence conditions while destruction satisfies only local failure conditions. The two activities operate under fundamentally different constraints, and the ease of the latter is not evidence of the ease of the former.

B. Glossary of Key Terms

Aspect A constraint, sub-operation, or dependency constituting part of a task.

Active attention set $A(t)$ The subset of aspects under conscious monitoring at time t.

Coherent configuration A configuration satisfying all constraints in $\mathcal{C}$.

Compression A mapping from the valid configuration space to a lower-dimensional representational space.

Entropic asymmetry The structural fact that configurations in Ω_{invalid} vastly outnumber those in Ω_{valid}, making construction harder than destruction.

Expansion A process of configuration-space traversal through which invariant structure is discovered.

Relegation The withdrawal of attention from stabilized aspects via the relegation operator R.

Stabilization The condition of an aspect when repeated execution satisfies its constraints without explicit monitoring.

C. Note on Related Theoretical Frameworks

The formal structures developed in this essay intersect with several established frameworks. The expansion–compression distinction bears structural resemblance to the *minimum description length* principle in statistical learning theory, which formalizes the trade-off between model complexity and data fit. The entropic asymmetry is closely related to the second law of thermodynamics and to the statistical mechanics of constrained systems. The attention-bounded model of relegation connects to resource-rational accounts of cognition in contemporary cognitive science.

The broader theoretical program of which this essay forms a part—the Relativistic Scalar–Vector Plenum framework and associated work—provides a field-theoretic account of constraint, flow, and irreversibility that may eventually supply a more rigorous mathematical foundation for the arguments developed here. The present essay is intended as a conceptually accessible treatment of the core ideas, preceding rather than assuming that more technical apparatus.