

Autobiographical Erasure in Scientific Writing

Interpretive Drift, Compression, and the Search for Invariants

Flyxion

July 2026

Abstract

Scientific writing is conventionally understood to suppress the personal in the name of objectivity: methods replace motivations, conclusions replace curiosity, and the scientist disappears behind the structures they discover. This essay argues that this suppression cannot be explained by methodological norms alone. It emerges from a deeper and more general phenomenon – the tendency of communication, under conditions of scale, ambiguity, and repeated reinterpretation, to favor representations that remain recognizable under transformation over representations that are merely true to experience. Beginning from the observation that even two people cannot guarantee shared understanding, the essay traces how interpretive divergence compounds combinatorially with audience size, how curricula, disclaimers, and formal notation function as devices for constraining that divergence, and how autobiography itself becomes difficult to transmit once a narrative implicates other people, leaks unintended information, or must survive transport through heterogeneous publics. The essay further argues that pseudonyms, scientific authorship, and even the discovery of mathematical invariants are instances of the same underlying strategy: the construction of structures whose identity persists despite repeated projection through different observers. A formal appendix develops these claims as propositions concerning projection, mutual information, channel capacity, and transformation order. The essay closes by arguing that autobiography is not eliminated by this process but relocated – from the chronology of events to the geometry of constraints – and that the essay itself, in making this argument, becomes evidence for its own thesis.

Contents

1	Introduction: The Disappearance of the Personal	3
2	Interpretive Drift and the Geometry of Misunderstanding	4
3	From Social Exhaustion to Formal Abstraction	6
4	Becoming a Symbol: Projection, Identity, and the Refusal of Interpretation	8
5	Interpretation Begins with Two People	10
6	The Combinatorics of Public Interpretation	10
7	Curricula, Disclaimers, and the Regulation of Interpretive Variance	11
8	Adler’s Problem of Understanding	12
9	The Narrative Self as an Interface	14
10	Public Intellectuals as Compressed Interfaces	15
11	Public Communication as Semantic Compression	16
12	Audience Constraints and the Compression of Meaning	17
13	Disclaimers as Interpretive Operators	19
14	The Search for Semantic Invariants	20
15	Autobiographical Erasure as a Communication Limit	22
16	The Observer Cannot Be Removed	23
17	A Brief Failure of Autobiography	24
18	Narrative Engineering and the Protection of Others	25
19	Inference, Metadata, and the Accidental Autobiography	27
20	Erasure and Exposure: Autobiography as Negotiated Projection	29

21 Names, Pseudonyms, and the Limits of Attribution	30
22 Pseudo-Identity and Collision Tolerance	32
23 Identity Through Transformation Order	33
24 Authorship as an Inference Problem	34
25 Authorship Attribution as Collision Detection	36
26 The Autobiography of Constraints	37
27 Why Theorists Drift Toward Geometry	38
28 Admissible Self-Projection	39
29 The Autobiography Was Never Missing	41
A Mathematical Formalism of Autobiographical Erasure	43
B Inference, Metadata, and Attribution	45
C Erasure, Exposure, and the Trace Principle	46
D Invariants, Geometry, and the Drift Toward Formalism	48
E Curricula, Disclaimers, and Frame Operators	49
F Narrative Engineering as Equivalence-Class Construction	50
G The Narrative Chain and the Erosion of Information	52
H Public Persona and Channel Capacity	52
I Admissibility and the Geometry of Self-Projection	54
References	55

1 Introduction: The Disappearance of the Personal

The conventional explanation for the impersonal character of scientific writing is that objectivity requires the removal of subjective experience. Scientific claims must be publicly verifiable, reproducible, and independent of the particular individuals who formulate them. The conventions of scientific prose therefore discourage autobiographical reflection, emotional narration, and personal testimony in favor of operational definitions, explicit methodologies, and formally stated conclusions. The resulting literature often appears detached from the lives of the people who produce it. Discoveries emerge without discoverers. Questions appear without motivations. Entire research programs seem to materialize from a sequence of logical deductions rather than from years of curiosity, frustration, uncertainty, and revision.

While this explanation captures an important aspect of scientific communication, it is incomplete. The suppression of the personal is not solely a methodological requirement imposed by the norms of science. It is also the result of a broader convergence of social, linguistic, technological, and cognitive pressures. Scientific writing occupies a peculiar position within contemporary communication ecosystems because it provides a refuge from interpretive instability. In environments where statements are frequently detached from their intended context, reinterpreted through political frameworks, amplified through outrage dynamics, or associated with controversies that bear little relation to their original meaning, formal discourse offers a mechanism for constraining ambiguity and reducing the space of admissible interpretations.

This pressure has intensified in the age of networked communication. Social media systems reward engagement rather than precision. Statements are extracted from their original contexts, redistributed to audiences with different assumptions, and interpreted according to local ideological priors. A remark intended as a technical observation may be interpreted as a political statement. An attempt at abstraction may be treated as advocacy. A descriptive claim may be understood as a normative endorsement. The resulting instability creates incentives to avoid personal expression altogether, particularly for individuals whose interests lie in theoretical, scientific, or technical domains. Under such conditions, the movement toward formal language is not merely an aesthetic preference but an adaptive response to interpretive uncertainty.

At the same time, natural language itself presents persistent difficulties. Human communication relies heavily upon ambiguity, contextual inference, metaphorical extension, and underspecified reference. These properties make language extraordinarily flexible, but they also generate a proliferation of possible interpretations. Every statement exists within a manifold of potential meanings, many of which may diverge substantially from the intentions of the speaker. As communication networks increase in scale, the probability that a statement will encounter an audience operating under different assumptions approaches certainty. The speaker consequently loses control over the interpretive trajectory of the utterance.

These pressures encourage a search for representational forms that exhibit greater stability under transformation. The vocabulary developed throughout this essay – channels,

noise, signal, redundancy, capacity – is borrowed directly from Shannon’s mathematical theory of communication [10], in which a message is understood not as a transparent vessel for meaning but as a signal that must survive transmission through a channel that degrades it. Mathematics, formal logic, and scientific notation can be understood as attempts to identify structures that remain invariant across observers, contexts, and interpretive communities. Rather than communicating through narratives grounded in personal experience, formal systems seek descriptions whose validity depends as little as possible upon the identity of the participant. The attraction of scientific writing therefore extends beyond methodological rigor. It reflects a deeper cognitive preference for structures that preserve meaning under projection.

This essay argues that autobiographical erasure in scientific writing should not be understood merely as a stylistic convention. It emerges from the interaction of multiple forces: the politicization of public discourse, the amplification dynamics of digital communication, the ambiguity of natural language, the accumulation of experiences in which innocuous statements generate unintended reactions, and the cognitive search for geometric invariants capable of surviving repeated transformations of interpretation. The disappearance of the personal is therefore not simply a consequence of scientific norms. It is the visible surface of a broader process through which inquiry seeks refuge in structures that remain stable when everything else becomes contingent.

2 Interpretive Drift and the Geometry of Misunderstanding

A central assumption underlying ordinary communication is that the meaning intended by a speaker is approximately the meaning received by a listener. Everyday conversation functions because this assumption is usually true enough to permit coordination. Participants share cultural backgrounds, situational context, conversational history, and a common understanding of what kinds of interpretations are reasonable. Under these conditions, the space of possible meanings is constrained, and most ambiguities are resolved automatically through contextual inference.

The situation changes dramatically as communication scales. Once a statement is separated from its original context and distributed across heterogeneous audiences, the number of admissible interpretations expands. Different readers bring different assumptions, experiences, conceptual frameworks, and expectations. The utterance no longer occupies a single interpretive environment but enters a collection of partially overlapping semantic spaces. Meaning begins to drift.

This phenomenon can be understood geometrically. Consider an utterance as a point within a high-dimensional semantic manifold. The intended meaning corresponds not merely to a sequence of words but to an entire local neighborhood consisting of assumptions, references, goals, and contextual constraints. Every act of interpretation projects this neighborhood into another conceptual space occupied by the listener. Because no projection is perfectly information preserving, portions of the original structure are inevitably lost. Different listeners therefore construct different images of the same statement. The resulting interpretations may remain close to one another, but they may also diverge substantially.

The effect becomes particularly pronounced in environments characterized by rapid redistribution and limited context. A statement may be quoted without surrounding qualifications. A technical remark may be encountered by an audience unfamiliar with the underlying domain. An observation about a phenomenon may be interpreted as a judgment regarding that phenomenon. An attempt to describe a structure may be mistaken for an endorsement of its consequences. In each case the semantic trajectory of the utterance departs from the trajectory originally intended by the speaker.

Importantly, these distortions need not arise from malice. They emerge naturally from the geometry of communication itself. Every listener reconstructs meaning using local information. Since different individuals possess different conceptual resources, the reconstruction process produces different results. Misunderstanding is therefore not an anomaly superimposed upon communication. It is a fundamental consequence of finite observation and incomplete access to context – close to what Eco, writing on the limits of interpretation, identifies as the tension between a text's openness to indefinite reading and the existence of structural constraints that nonetheless rule out certain readings as inadmissible [4].

The modern communication environment amplifies these tendencies. Social platforms reward statements that are rapidly interpretable by large audiences. Nuance imposes cognitive costs. Qualification consumes attention. Context requires effort to reconstruct. The economic incentives governing visibility therefore favor utterances that can be projected into many interpretive frameworks with minimal effort. Unfortunately, such utterances are often precisely those most vulnerable to distortion. The statement survives distribution by becoming increasingly detached from the circumstances that originally gave it meaning.

Over time, repeated exposure to these dynamics alters communicative behavior. Individuals learn that even carefully phrased remarks may acquire interpretations they neither intended nor anticipated. They observe innocuous statements becoming controversial, descriptive observations becoming political symbols, and technical discussions becoming absorbed into broader cultural conflicts. The resulting experience produces a form of communicative caution. One becomes increasingly aware not only of what a statement means but of the manifold of meanings it might acquire once released into a larger interpretive ecosystem.

This accumulation of experiences generates pressure toward representational forms that exhibit lower interpretive volatility. Formal definitions, mathematical expressions, logical derivations, and explicitly specified assumptions all function as mechanisms for reducing the dimensionality of the interpretive space. Their purpose is not merely precision in the ordinary sense. Their deeper function is to constrain the set of admissible projections. A theorem is attractive because it remains recognizable after repeated transformations. A mathematical structure survives transportation across contexts more effectively than an informal narrative because its meaning depends upon fewer unstated assumptions.

The movement toward scientific writing can therefore be interpreted as a response to interpretive drift. Scientific discourse attempts to construct linguistic objects whose semantic trajectories remain stable under repeated projection. Definitions replace implications. Operational procedures replace intuitions. Formal relationships replace rhetorical associations. The goal is not simply clarity but invariance. Scientific language seeks structures that

preserve their identity as they move through increasingly diverse interpretive environments.

Viewed from this perspective, autobiographical erasure acquires a new significance. Personal narratives are among the most context-dependent forms of communication. Their meaning depends heavily upon shared background knowledge, emotional resonance, historical circumstances, and tacit assumptions regarding intention. Because these elements are difficult to formalize, they are particularly vulnerable to distortion. The removal of the personal from scientific discourse is therefore not merely a rejection of subjectivity. It is an attempt to minimize semantic curvature by eliminating dimensions along which interpretive divergence can accumulate.

The scientist disappears from the text not because personal experience lacks importance, but because personal experience is difficult to transport without deformation. The more communication is expected to survive projection across audiences, institutions, and generations, the greater the incentive to replace autobiographical trajectories with structures that remain approximately invariant under transformation.

3 From Social Exhaustion to Formal Abstraction

The transition from personal expression to formal abstraction is often described as an intellectual choice motivated by a desire for rigor. While this explanation contains an element of truth, it overlooks the extent to which abstraction can emerge as an adaptive response to repeated communicative failures. The movement toward scientific and mathematical language is not always driven by an attraction to formal systems alone. In many cases it is equally driven by fatigue with the instability of ordinary discourse.

Every individual develops expectations regarding communication through accumulated experience. Early interactions typically create an implicit model in which words function as relatively reliable vehicles for transmitting intentions. Over time, however, this model is repeatedly challenged. Statements that appear straightforward produce unexpected reactions. Explanations that seem precise are interpreted in incompatible ways. Attempts at clarification generate further misunderstandings. Observations intended to be descriptive are received as normative claims, while abstract inquiries are interpreted through political, ideological, or personal frameworks that bear little resemblance to their original purpose.

These experiences do not merely produce frustration. They gradually alter one's understanding of language itself. Communication ceases to appear as a process of transferring meanings from one mind to another and instead begins to resemble a process of navigation through a landscape of competing interpretations. One becomes increasingly aware that every utterance occupies multiple semantic trajectories simultaneously. What matters is not only what is said but how it may be reconstructed by observers operating under different assumptions.

The contemporary communication environment intensifies this realization. Public discourse increasingly rewards rapid categorization. Statements are often interpreted according to their perceived alignment with existing narratives rather than according to their internal structure. Questions become associated with political identities. Curiosity becomes

confused with advocacy. Analysis becomes conflated with allegiance. Under such conditions, many subjects acquire a kind of semantic gravity that pulls nearby discussions into preexisting interpretive attractors regardless of the intentions of the participants.

For individuals primarily interested in understanding structures rather than participating in ideological conflicts, this process can become deeply discouraging. The difficulty lies not merely in disagreement. Scientific inquiry depends upon disagreement. The difficulty arises when discussions cease to be organized around the internal properties of arguments and instead become organized around external associations. At that point, the content of a statement becomes secondary to the narratives within which it is placed.

One consequence is a growing preference for representational systems that are less vulnerable to such transformations. Mathematics offers perhaps the clearest example. A mathematical object possesses properties that remain independent of political affiliation, social identity, or rhetorical framing. A proof may be accepted or rejected, generalized or criticized, but its evaluation depends primarily upon its structural relationships rather than upon the personal characteristics of the individual who produced it. Formal systems are therefore attractive not merely because they are precise but because they provide a degree of insulation from interpretive turbulence.

Scientific writing inherits many of these advantages. The conventions of the genre systematically remove features that are likely to trigger extraneous interpretations. Personal motivations are minimized. Emotional language is discouraged. Claims are linked to explicit evidence. Definitions are introduced before use. Assumptions are stated directly. Each of these conventions narrows the range of admissible reconstructions available to the reader. The resulting text becomes less expressive in some respects but more stable under transmission.

This process can be understood as a form of cognitive resource allocation. Every misunderstanding imposes a cost. Clarifications require effort. Context must be reconstructed. Misinterpretations must be corrected. When these costs accumulate over sufficiently long periods, individuals naturally seek communicative strategies that reduce the frequency of such failures. Formal abstraction emerges as one such strategy. Rather than continuously negotiating meanings within unstable semantic environments, one attempts to construct representations whose interpretation depends upon explicit structure rather than contextual inference.

The movement toward abstraction is therefore not simply an escape from the personal. It is often the result of repeated encounters with the limitations of personal communication. The scientist, mathematician, or theorist does not necessarily begin by rejecting narrative and experience. More commonly, narrative and experience reveal themselves to be difficult media through which stable understanding can be achieved. Formalism becomes attractive because it appears to offer a domain in which meaning can persist despite differences in perspective.

Paradoxically, however, this very process contributes to autobiographical erasure. The experiences that motivate the search for formal invariants are precisely the experiences least likely to appear in the resulting formal structures. The frustrations, misunderstandings, interpretive conflicts, and communicative failures that encourage abstraction disappear

once abstraction succeeds. The final theory preserves the invariant but not the trajectory through which the invariant was discovered. As a consequence, the literature records the destination while obscuring the path that made the destination necessary.

4 Becoming a Symbol: Projection, Identity, and the Refusal of Interpretation

The pressures that encourage autobiographical erasure are not unique to science. Similar dynamics appear whenever an individual becomes the object of large-scale interpretation. Artists, writers, public intellectuals, scientists, and political figures frequently discover that the meanings attributed to their words and actions begin to diverge from their own understanding of those words and actions. The phenomenon is sufficiently widespread that it appears less as a peculiarity of particular individuals than as a structural feature of communication itself.

The problem emerges whenever a person becomes embedded within a sufficiently large interpretive ecosystem. At that point, observers no longer encounter the individual directly. Instead, they encounter a representation constructed from fragments of prior encounters, media portrayals, cultural narratives, ideological commitments, and personal expectations. The individual becomes a symbolic object occupying a position within the conceptual landscape of others. Communication increasingly concerns the symbol rather than the person.

This transformation can be understood as a consequence of semantic projection. Every observer reconstructs an individual using locally available information. Because different observers possess different histories and assumptions, the resulting reconstructions differ. A single person therefore gives rise to a collection of projected identities distributed across multiple interpretive spaces. Some of these projections may resemble the individual closely. Others may bear only a tenuous relationship to the person from whom they originated. Nevertheless, once established, these projections acquire a degree of autonomy. They begin to participate in social discourse independently of the individual who generated them.

The consequences are particularly visible in public life. An artist may discover that audiences regard them as the embodiment of a political movement. A scientist may become associated with positions they never advocated. A writer may find that a narrow aspect of their work comes to dominate public perceptions of their entire intellectual identity. In each case, the individual encounters a widening gap between self-description and public representation. Communication increasingly takes place through layers of projection that neither side fully controls.

A recurring artistic example of this dynamic is Todd Haynes's film *I'm Not There*, which represents Bob Dylan through six different performers rather than one. The film's central tension – the public's insistence on Dylan as spokesman, prophet, or coherent narrative, against his refusal to remain any one of those things – dramatizes exactly the structural problem under discussion. The film's response is artistic multiplicity: it multiplies masks rather than settling on one. The retreat into formal abstraction explored in this essay is

a different response to the same problem. Instead of multiplying identities, the theorist narrows them. Instead of producing more narratives, they produce fewer, attempting to construct invariant structures less vulnerable to projection. In both cases the underlying difficulty is the same: once communication enters a large interpretive ecosystem, observers respond less to what was said than to what the statement represents within their own conceptual frameworks.

The recurring response to this condition is therefore not always confrontation. Often it is withdrawal. Some individuals retreat from public visibility altogether. Others cultivate ambiguity. Still others continually reinvent themselves in order to escape stabilization within a single interpretive framework. The common element is a resistance to becoming reducible to any particular projection.

The significance of this phenomenon for scientific communication is easily overlooked. Scientists are often assumed to occupy a purely technical domain insulated from such concerns. Yet the same pressures operate there as well. Scientific ideas circulate through institutions, media systems, academic communities, and public discussions. Researchers discover that explanations are reinterpreted, motivations are inferred, and intellectual positions become associated with broader narratives. Even when controversy is absent, the process of repeated interpretation gradually transforms the relationship between the individual and their work.

Under these conditions, abstraction acquires an additional function. Formalism does not merely increase precision. It provides a partial defense against symbolic capture. By directing attention toward structures rather than personalities, formal discourse attempts to reduce the degree to which ideas depend upon representations of their authors. Equations, proofs, models, and formal definitions can be evaluated independently of the narratives attached to the individuals who produced them. They are not immune to interpretation, but they possess a greater degree of stability than autobiographical narratives.

The movement toward scientific writing can therefore be understood as one manifestation of a broader phenomenon: the refusal to become reducible to projection. The scientist who removes personal experience from a paper, the mathematician who prefers formal derivations to anecdotal explanation, and the artist who continually escapes categorization may be responding to the same underlying condition. Each confronts a world in which meanings proliferate through interpretation and seeks a mode of expression capable of preserving identity under repeated transformation.

Autobiographical erasure emerges from this perspective as a paradoxical act of preservation. By removing the self from the text, one attempts to protect the work from becoming merely another representation of the self. The personal disappears not because it lacks significance, but because it has become difficult to communicate without being transformed into something else. What remains is an effort to construct forms that survive projection, structures that retain their identity despite the multiplicity of interpretations through which they must pass.

5 Interpretation Begins with Two People

It is tempting to regard interpretive instability as a phenomenon that emerges only in large communication systems. Social media, mass journalism, political discourse, and public controversy appear to provide obvious examples of statements acquiring meanings far removed from their original intentions. Yet the underlying phenomenon does not begin at the scale of societies. It begins at the scale of a single conversation.

Whenever one person communicates with another, meaning must cross a boundary separating two distinct cognitive systems. The speaker and listener possess different experiences, memories, associations, conceptual frameworks, emotional histories, and linguistic habits. Even when both parties use the same vocabulary, there is no guarantee that identical structures are being activated within their respective conceptual spaces. Communication therefore cannot be understood as the transfer of meaning itself. At best, it is the transmission of signals capable of inducing approximately similar constructions within different minds.

This observation implies that misunderstanding is not an anomaly requiring explanation. Perfect understanding is the anomaly. Most communication succeeds because the distortions introduced during interpretation remain small enough to permit coordination. Shared culture, shared language, shared experiences, and repeated interaction all help constrain the resulting divergence. Nevertheless, the possibility of interpretive drift is present from the very beginning.

The implications are profound. If interpretation is already uncertain when only two individuals are involved, then the challenge facing scientific communication cannot be reduced to problems of mass media or contemporary politics. The difficulty is inherent to communication itself. Every statement exists within a space of possible reconstructions, and every act of interpretation selects one trajectory from among many alternatives.

Scientific writing can therefore be understood as a response to a universal problem rather than a specialized academic convention. Its procedures seek to constrain interpretive freedom at the smallest possible scale, reducing ambiguity even before communication encounters larger audiences. The formalism of scientific discourse reflects an attempt to stabilize meaning in the presence of an unavoidable gap between minds.

6 The Combinatorics of Public Interpretation

Although interpretive instability originates in individual interactions, its growth under scaling is not linear. Each additional participant introduces a new interpretive space through which a statement may be reconstructed. As the number of observers increases, the number of potential semantic trajectories expands dramatically.

A statement delivered to a single listener generates one reconstruction. Delivered to ten listeners, it generates ten reconstructions. Yet the situation is more complex than a simple multiplication. Listeners discuss interpretations with one another, reinterpret statements through secondary reports, and react not only to the original utterance but also

to the interpretations already circulating within the community. Communication becomes recursive. Interpretations acquire interpretations of their own.

The resulting process resembles a branching system rather than a transmission channel. Meanings proliferate as they propagate. A statement may gradually acquire associations that were absent from its original context. Particular interpretations may become amplified because they are emotionally salient, politically useful, or socially rewarding. Others may disappear entirely. Over time, the relationship between the original utterance and its public representations becomes increasingly indirect.

This dynamic helps explain why communication strategies that function effectively in small groups often fail in large audiences. Informal conversation relies heavily upon shared context and real-time correction. Public communication lacks these mechanisms. The speaker must therefore anticipate a much larger space of potential interpretations. The challenge is not merely to communicate a message but to communicate a message that remains recognizable after repeated projection through heterogeneous audiences.

From this perspective, the conventions of scientific discourse appear less as stylistic preferences than as scaling technologies. Definitions, citations, methodological descriptions, and formal notation function as mechanisms for limiting the proliferation of semantic trajectories. They attempt to ensure that a statement retains its identity even after passing through numerous interpretive transformations. The appendix makes this scaling claim precise: once an audience must be told not merely the original statement but how to interpret its own divergent interpretations, the burden of explanation grows quadratically rather than linearly in audience size.

7 Curricula, Disclaimers, and the Regulation of Interpretive Variance

Many seemingly unrelated institutions can be understood as responses to the same fundamental problem: the tendency of interpretation to diverge from intention. Educational systems, legal frameworks, technical standards, public speaking conventions, and even comedy all contain mechanisms designed to regulate interpretive variance.

The educational curriculum provides a particularly clear example. A curriculum is often described as a sequence of topics to be covered, but its deeper function is to constrain the interpretive manifold available to students. Definitions establish common reference points. Exercises reveal misunderstandings. Examinations test whether particular interpretations have been acquired. The teacher's task is not merely to present information but to repeatedly reduce divergence between intended and reconstructed meanings.

Public speaking operates under similar constraints. Effective speakers rarely rely upon spontaneous exposition alone. Instead, they employ carefully structured narratives, repeated themes, familiar examples, and explicit framing devices. These structures function as stabilizing mechanisms that increase the probability that audience members will construct compatible interpretations. The larger the audience becomes, the more important such mechanisms become.

Comedy offers an especially revealing case because comedians often make their interpretive strategies explicit. Phrases such as “this is just a joke” do not add content to the joke itself. Rather, they establish a frame within which subsequent statements should be interpreted – precisely the operation Bateson identified in his account of play and meta-communication, where a signal such as “this is play” does not describe the content of an exchange but governs how every subsequent signal within it should be read [3]. The disclaimer attempts to position the audience within a particular semantic region before the material is encountered. Humor frequently depends less upon the literal content of a statement than upon the interpretive context in which that statement is received.

A similar transformation is now visible in computational systems. Earlier traditions of software engineering emphasized verification, correctness, and deterministic behavior. Programs were expected to conform closely to specifications. Contemporary generative systems increasingly employ a different strategy. Rather than guaranteeing correctness, they frequently acknowledge uncertainty through explicit disclaimers indicating that outputs may contain errors, hallucinations, or unreliable information.

These two approaches represent distinct methods for managing interpretive uncertainty. Verification seeks to reduce the space of admissible interpretations through formal constraints. Disclaimers acknowledge that uncertainty remains and transfer part of the interpretive burden to the user. Both strategies address the same underlying problem, namely the impossibility of guaranteeing that an intended meaning will be reconstructed exactly as intended.

Scientific writing belongs primarily to the tradition of constraint rather than disclaimer. Its characteristic features are attempts to narrow the interpretive manifold before misunderstanding occurs. Yet the existence of disclaimers in contemporary computational systems serves as a reminder that complete stabilization may be unattainable. Interpretation remains an open-ended process. The goal of formal communication is not to eliminate ambiguity entirely but to reduce its growth sufficiently that meaning can survive repeated transformations across observers, institutions, and time.

8 Adler’s Problem of Understanding

Long before the emergence of social media and large-scale digital communication, Mortimer Adler observed that understanding could not be assumed simply because words had been exchanged. Communication required the construction of similar conceptual structures within different minds, and the same text could generate substantially different interpretations in different readers. Much of what appears to be disagreement, on Adler’s account, is really a failure of semantic alignment: two people may use identical words while referring to different concepts, or different words while referring to the same concept. Adler further noted that the problem intensifies with abstraction. Concrete objects can often be pointed to; abstract concepts such as justice, freedom, intelligence, truth, or consciousness cannot. As abstraction increases, interpretive divergence becomes increasingly likely. The problem therefore existed prior to mass communication and prior to contemporary concerns regarding misinformation or political polarization.

The present argument extends Adler's observation by treating interpretation as a geometric process. Misunderstanding is not merely an occasional failure of communication but a natural consequence of projection between distinct conceptual spaces. Every act of communication maps an intended semantic structure into a different cognitive environment. Interpretive divergence therefore exists even between two participants. What changes under scaling is not the existence of the phenomenon but the number of possible trajectories through which meaning may evolve.

From this perspective, educational curricula, scientific writing, formal logic, legal language, and mathematical notation may all be understood as technologies for reducing interpretive variance. Their function is not simply the transmission of information but the construction of semantic invariants capable of surviving repeated projection across observers.

A deeper reading of Adler suggests that communication contains a hidden asymmetry. The speaker has access to the entire trajectory that produced a thought, while the listener receives only a compressed representation of that trajectory. A sentence may be the endpoint of years of observation, reflection, revision, and experience. The listener encounters only a small collection of symbols and must reconstruct the missing structure from local information. The scientist writes $T \rightarrow S$, where T is a long intellectual trajectory and S is the final statement. The reader must perform the inverse operation $S \rightarrow \hat{T}$, constructing an estimated trajectory \hat{T} from the statement alone. Communication succeeds only when $\hat{T} \approx T$. Yet there is no guarantee that this approximation will hold. Different readers construct different trajectories; different audiences reconstruct different motivations; different communities generate different interpretations.

What scientific writing does is peculiar. Instead of attempting to communicate the trajectory T , it increasingly focuses on properties of S that can be independently verified. The paper does not say, in effect, "here is how I arrived at this idea." It says, "here is a structure that remains true regardless of how I arrived at it." In a sense, scientific writing abandons the reconstruction problem entirely. It ceases trying to communicate the history of thought and instead attempts to communicate invariants that remain stable even when the history is forgotten.

That is why scientific papers often feel strangely impersonal. The author's actual intellectual life is hidden behind the final structure. The irony is that the omitted trajectory is often exactly what young researchers most need to see. They already possess the published theorem, proof, model, or experiment. What they lack is an understanding of how such things emerge from confusion, failure, curiosity, and persistence.

This suggests that autobiographical erasure is not merely a feature of scientific writing but a feature of successful knowledge itself. As knowledge becomes more universal, its origins disappear. Nobody learning algebra today learns the personal frustrations of al-Khwārizmī. Nobody using calculus routinely encounters the intellectual struggles of Newton or Leibniz. Nobody applying information theory needs to reconstruct the personal circumstances of Claude Shannon. The more universal an idea becomes, the less dependent it is on the person who discovered it. Science deliberately optimizes for this outcome.

The central tension of this essay is therefore not merely that science removes the personal.

It is that science succeeds by removing the personal. The very mechanism that allows ideas to survive across cultures, centuries, and disciplines is the same mechanism that obscures the human trajectories that produced them. Autobiographical erasure is not a mistake in scientific writing. It is the price paid for universality.

9 The Narrative Self as an Interface

A common assumption underlying discussions of autobiography is that personal narratives reveal the individual who produces them. Closer examination suggests the opposite. Personal narratives are already highly compressed social artifacts constructed for the purpose of communication – close to Ricoeur’s account of narrative identity, in which selfhood is not given prior to the story but constituted through the act of emplotment, so that the narrated self is always already an artifact rather than a transparent report [8].

No individual can communicate the totality of their experiences, memories, motivations, conceptual structures, and internal states. The informational requirements would be effectively unbounded. Consequently, social interaction depends upon simplified representations that allow rapid coordination.

The narrative self functions as one such representation. Statements such as “I am a scientist,” “I am a parent,” “I am an artist,” or “I am an entrepreneur” do not identify the person itself. They identify socially useful projections through which interaction can proceed.

In this sense, the narrative self resembles a name tag, a uniform, or a sign reading “staff.” Such markers do not describe the full individual. Rather, they provide a low-dimensional interface through which other participants can orient their expectations and organize their interactions – the same low-dimensional interface Goffman described under the heading of self-presentation, in which social life is structured around a performed front rather than direct access to the performer [6].

Autobiographical communication therefore begins with an act of compression. Before any listener encounters the person, they encounter a representation of the person. The self presented in communication is already a managed artifact designed to facilitate interpretation.

This observation alters the meaning of autobiographical erasure. Scientific writing does not remove the individual directly. It removes a communicative layer that was itself a projection. The scientific paper strips away not merely personal experience but also the socially optimized interface through which personal experience is ordinarily communicated.

What remains is not the person but a structure. Equations, definitions, models, and proofs survive. The narrative interface through which others would ordinarily identify the author is systematically reduced.

The resulting text therefore represents a peculiar limit case of communication. Rather than presenting a compressed representation of a person, it attempts to present a compressed representation of a structure whose validity is intended to remain independent of the person

who produced it. The progression can be summarized as a sequence of reductions:

Person → Narrative Self → Scientific Author → Formal Structure,

where each step removes dimensions of interpretation in order to make communication more stable across observers. The endpoint is not the self at all. The endpoint is an invariant – and this aligns with the broader theme developed throughout this essay, that cognition often appears to search for structures that survive projection, transport, and reinterpretation.

10 Public Intellectuals as Compressed Interfaces

The pressures that produce autobiographical erasure become especially visible in the case of public intellectuals. Large audiences impose constraints that differ fundamentally from those governing communication among specialists.

Consider a scientist communicating with professional peers. The conversation may contain technical terminology, mathematical formalism, unresolved questions, speculative proposals, and highly qualified claims. Such communication assumes substantial shared background knowledge and therefore permits a relatively large interpretive bandwidth.

As audience size increases, this bandwidth contracts. The communicator must increasingly rely upon concepts that remain stable across diverse interpretive communities. Technical vocabulary is simplified. Mathematical formalism is reduced or eliminated. Context-dependent arguments are replaced by portable narratives. Controversial topics are often avoided because they generate large numbers of competing interpretive trajectories.

The result is the emergence of a public persona that functions as a compressed communication interface. The public intellectual becomes associated with a relatively small collection of stable messages that can survive repeated projection across heterogeneous audiences.

Importantly, this compressed interface should not be confused with the individual. It is an adaptive artifact generated by the requirements of large-scale communication. The public figure may possess a much richer and more nuanced conceptual landscape than the one visible in their public representations. A theoretical physicist who becomes a science communicator, for instance, is typically forbidden – by the constraints of the medium rather than by any editorial decree – from using advanced mathematics, exploring unresolved controversies, or voicing tentative philosophical positions. The resulting public persona can come to resemble a caricature of the underlying discipline, not because the communicator's views are shallow, but because the channel through which they must speak admits only a narrow band of stable signal. The same communicator may rely on broad public support – subscriptions, sponsorships, crowdfunding – which further selects for content legible to the widest possible audience.

This dynamic creates a paradox. Success in public communication often requires the reduction of complexity, while intellectual inquiry frequently depends upon the exploration of complexity. The more widely an individual is heard, the greater the pressure to communicate through simplified structures capable of surviving transport across large and diverse interpretive environments.

Public visibility therefore tends to transform individuals into symbolic interfaces. Audiences interact not with the full cognitive system but with a stabilized projection optimized for communication. The public intellectual becomes recognizable precisely to the extent that complexity has been compressed into a portable form.

11 Public Communication as Semantic Compression

The dynamics described above become especially visible in the case of public communicators generally. Scientists, educators, authors, journalists, and public intellectuals are often regarded as individuals who have successfully amplified their voices to reach larger audiences. Yet a more revealing interpretation is possible. Public communication is not simply communication at larger scale. It is communication subjected to progressively stronger compression constraints.

Every increase in audience size expands the diversity of interpretive contexts through which a message must pass. The specialist audience of a research seminar possesses a shared technical vocabulary, common assumptions, and a relatively coherent conceptual framework. Such audiences can tolerate ambiguity because participants possess the resources necessary to reconstruct intended meanings. As communication moves outward toward broader publics, these shared structures become increasingly attenuated. The communicator must therefore adapt the message to survive transportation through a wider range of interpretive environments.

The result is a process of semantic compression. Technical vocabulary is replaced by common language. Mathematical formalisms are translated into metaphors. Qualifications are shortened. Caveats are reduced. Context is omitted. Nuance is compressed into memorable formulations capable of surviving redistribution through interviews, headlines, social media excerpts, and public discussion.

This transformation should not be understood as intellectual dishonesty. It is a consequence of channel capacity. The amount of information that can be reliably transmitted decreases as audience heterogeneity increases. Public communication therefore favors structures that remain recognizable after repeated projection through diverse conceptual spaces.

The consequences for the communicator are substantial. Over time, the public persona becomes increasingly distinct from the underlying intellectual trajectory. Audiences encounter a stabilized projection optimized for transmission rather than the full complexity of the individual. The scientist becomes associated with a handful of recurring themes. The philosopher becomes identified with a small collection of arguments. The writer becomes reduced to a few representative positions. The person gradually disappears behind a communication interface designed for large-scale interpretive stability.

This phenomenon is often interpreted as a failure of public understanding. Yet the process is more fundamental than simple misunderstanding. The public persona is not an inaccurate representation of the individual. Rather, it is a low-dimensional approximation constructed under severe communicative constraints. It functions much like a map. A map

is not the territory, but neither is it arbitrary. Its purpose is to preserve a subset of structures while discarding others – the same logic of legibility that Scott identifies in administrative projections of complex social terrain, where a state does not see a forest or a peasantry directly but only through simplified, standardized abstractions engineered to be governable at the cost of nearly everything else [9].

The same pressures that produce scientific abstraction therefore also produce public identities. Both are attempts to preserve certain invariants under conditions of repeated projection. The difference lies only in what is being preserved. Scientific writing attempts to preserve formal structures. Public communication attempts to preserve recognizable narratives. In both cases, a great deal of complexity is necessarily discarded.

The paradox is that successful communication often requires precisely the elimination of those aspects of experience that originally motivated communication. The more broadly an idea is distributed, the less of its original context can accompany it. The trajectory that produced the idea gradually disappears, leaving only those elements capable of surviving transport through large populations.

The scientific paper represents one endpoint of this process. The public persona represents another. Both are artifacts generated by the same underlying pressure: the need to construct communicative forms capable of surviving repeated transformations while retaining a recognizable identity. Neither is the individual. Both are compressed representations optimized for transmission through a world of imperfect interpreters.

12 Audience Constraints and the Compression of Meaning

The phenomenon of autobiographical erasure becomes easier to understand when viewed as a special case of a more general principle. Every act of communication is constrained by the capacities of its intended audience. The communicator is never free to transmit arbitrary complexity. Rather, communication must be adapted to the interpretive resources available to those receiving it.

A children's book provides a simple illustration. The author may possess extensive knowledge regarding psychology, history, biology, or social dynamics, yet only a small fraction of that knowledge can be communicated to an infant. Vocabulary must be restricted. Sentences must be shortened. Abstract concepts must be replaced by concrete examples. Ambiguity must be minimized. The resulting text is not false. Rather, it is a compressed representation optimized for a particular interpretive environment.

The same process occurs at every level of communication. Literature written for adolescents differs from literature written for specialists. Introductory textbooks differ from research monographs. Public lectures differ from graduate seminars. In each case the underlying subject matter may remain largely unchanged while the representation undergoes substantial transformation.

Consider the contrast between a popular science book and a technical research paper. The paper may contain mathematical derivations, specialized terminology, extensive qualifications, and unresolved ambiguities. A popular treatment of the same subject often

removes most of these elements. Equations become metaphors. Technical distinctions become narratives. Formal structures become intuitions. The goal is not merely simplification but adaptation to a different interpretive manifold.

The same phenomenon appears in fiction. Authors writing for younger audiences frequently avoid complex political questions, unresolved moral ambiguities, extensive historical context, or intricate philosophical arguments. These omissions do not necessarily reflect the author's beliefs. Rather, they arise from constraints imposed by the communicative channel. The narrative must remain navigable within the conceptual capacities of the intended audience.

From this perspective, communication resembles the design of a map. A map intended for a pedestrian differs from a map intended for an airline pilot. Neither map is complete. Both omit enormous quantities of information. Their value derives not from completeness but from preserving the structures relevant to particular navigational tasks.

Autobiographical narratives function similarly. When individuals describe themselves, they do not present the entirety of their experiences, memories, motivations, and conceptual structures. Such a description would be impossible. Instead, they construct audience-specific maps of themselves. The narrative offered to a colleague differs from the narrative offered to a child. The narrative offered during a job interview differs from the narrative offered to a close friend. Each represents a different projection of the same underlying structure.

Scientific writing occupies one extreme of this spectrum. Its audience is presumed to possess substantial technical competence. Consequently, many ordinary narrative devices become unnecessary or even counterproductive. Personal histories, emotional motivations, and autobiographical context are removed in favor of structures that can be evaluated independently of the individual who produced them. The resulting text is therefore not merely less personal. It is adapted to an audience whose primary concern is the stability of formal relationships rather than the reconstruction of personal trajectories.

The broader lesson is that communication always involves selective preservation. The question is never whether information will be lost; information loss is unavoidable. The question is which structures survive compression. Children's literature preserves concrete narratives. Public communication preserves memorable themes. Scientific writing preserves formal relations. Mathematical notation preserves invariants. Each communicative form can therefore be understood as a different answer to the same problem: how to maximize the persistence of meaning under the constraints imposed by a particular audience. The implication runs deeper still: the "self" may itself be an audience-dependent compression algorithm. The version of a person that appears in a grant proposal, a social media profile, a research paper, and a conversation with a child are not different people. They are different projections of the same underlying structure onto different interpretive manifolds.

13 Disclaimers as Interpretive Operators

The role of disclaimers in contemporary media reveals an important feature of communication that is often overlooked. A disclaimer does not alter the content of a message. Instead, it attempts to alter the interpretive context within which the message is received.

Consider the warning frequently displayed before a television series containing graphic violence, drug use, or mature themes: *viewer discretion is advised*. The disclaimer does not remove any scenes from the series. It does not modify the script. It does not alter the motivations of the characters or the events of the narrative. The informational content of the work remains unchanged. Yet broadcasters regard the disclaimer as important because it influences how viewers are expected to interpret what follows.

From an informational perspective, the disclaimer functions as a precondition on interpretation. It establishes a semantic frame before the content itself is encountered. The viewer is informed that subsequent events should not be understood as ordinary entertainment alone but as depictions involving themes that some audiences may find disturbing or emotionally significant.

This observation reveals a broader principle. Communication frequently depends not only upon the message but also upon the interpretive operators applied before the message arrives. Educational prerequisites, legal definitions, scientific methodologies, trigger warnings, content advisories, and introductory remarks all perform related functions. They constrain the space of admissible interpretations available to the audience.

The existence of such mechanisms reflects a recognition that meaning does not reside exclusively within content itself. Meaning emerges from the interaction between content and interpretive context. The same sequence of words, images, or events may produce radically different effects depending upon the conceptual framework through which it is encountered.

In this sense, disclaimers function as a distinct strategy for managing interpretive variance. Rather than reducing complexity within the message itself, they attempt to reduce divergence by shaping expectations. The communicator acknowledges that multiple interpretations are possible and intervenes by suggesting a preferred region of the semantic space.

This differs significantly from the strategy employed by scientific writing. Scientific discourse generally seeks to constrain interpretation through formalization. Definitions, equations, citations, and methodological descriptions are embedded directly within the content. The reader is guided toward particular interpretations through the internal structure of the document itself.

Disclaimers represent a complementary approach. Instead of modifying the structure of the message, they modify the conditions under which the message is received. They act as external operators on interpretation rather than internal constraints on expression.

The proliferation of disclaimers in contemporary culture may therefore be understood as evidence of increasing awareness regarding interpretive uncertainty. As communicative environments become more heterogeneous and audiences become more diverse, speakers

can no longer assume common interpretive frameworks. Additional mechanisms become necessary to orient listeners before communication begins.

Viewed from this perspective, the disclaimer preceding a television series, the legal disclaimer attached to a contract, the warning accompanying a machine learning system, and the methodological section of a scientific paper all address the same fundamental problem. Each attempts to reduce the divergence between intended and reconstructed meanings. They differ only in where the intervention occurs. Some modify the message itself. Others modify the interpretive field surrounding the message.

The disclaimer is therefore not merely a warning. It is an admission that communication alone is insufficient. Before meaning can be transmitted, the audience must first be positioned within an appropriate interpretive manifold. The warning that a generative system “may hallucinate” serves almost exactly this function: it does not change the output, but it changes the interpretive stance the audience is expected to adopt toward the output – exactly as a content advisory changes the stance of a viewer before a violent scene begins.

14 The Search for Semantic Invariants

The recurring examples examined throughout this essay appear at first to belong to unrelated domains. Scientific papers eliminate autobiography. Teachers follow curricula. Public speakers rely upon rehearsed narratives. Comedians frame jokes. Television programs employ disclaimers. Software systems warn users that outputs may be unreliable. Public intellectuals simplify their language for broader audiences. Yet beneath these diverse practices lies a common structural problem.

All communication occurs in the presence of transformation. A thought originates within one cognitive system and must pass through language, media, institutions, audiences, cultures, and historical periods before reaching another cognitive system. At every stage, distortions are introduced. Context is lost. Assumptions change. Associations shift. Interpretive frameworks evolve. Communication therefore faces a challenge analogous to the transmission of signals through a noisy channel.

The central question becomes: what survives?

Certain structures prove remarkably fragile. Emotional states, personal experiences, subtle motivations, local references, and autobiographical details often degrade rapidly as communication travels through larger interpretive systems. These elements depend heavily upon context and shared background knowledge. Once detached from their original environments, they become increasingly vulnerable to reinterpretation.

Other structures display greater persistence. Mathematical relationships survive translation between languages. Logical inferences remain recognizable across cultures. Geometric properties remain invariant under transformations that radically alter appearance. Scientific laws often outlive the societies that discovered them. Such structures possess an unusual capacity to maintain identity despite repeated projection through different interpretive spaces.

This asymmetry suggests that cognition itself may be oriented toward the discovery of

invariants. Human beings frequently seek descriptions that remain stable despite changes in perspective. The attraction of mathematics may derive not merely from precision but from the existence of structures that survive transformation. A theorem remains a theorem regardless of the personal characteristics of the individual presenting it. A geometric relationship remains unchanged when represented in different coordinate systems. Such objects exhibit a form of communicative robustness that ordinary narratives rarely achieve.

From this perspective, the movement toward abstraction appears less mysterious. Formal systems emerge because they contain unusually persistent structures. They provide a means of communicating relationships that remain recognizable even after substantial contextual loss. The appeal of abstraction is therefore not necessarily a rejection of lived experience. Rather, it reflects an attempt to identify features of reality capable of surviving transmission.

The same process can be observed in ordinary social life. Names, professional titles, national identities, institutional affiliations, and personal narratives all function as simplified representations designed to persist across interactions. They are communicative invariants of a sort. Their purpose is not to describe individuals completely but to provide sufficiently stable points of reference that coordination becomes possible.

Yet these invariants are purchased at a cost. Every act of stabilization requires compression. Every reduction of interpretive variance eliminates information. The map gains portability by sacrificing detail. The scientific paper gains reproducibility by sacrificing autobiography. The public persona gains recognizability by sacrificing complexity. The theorem gains universality by sacrificing the history of its discovery.

Autobiographical erasure therefore belongs to a broader family of phenomena. It is one instance of the tension between richness and stability. Rich descriptions preserve more information but travel poorly. Stable descriptions travel farther but preserve less. Communication continuously negotiates this trade-off, selecting structures that can survive the interpretive environments through which they must pass.

The search for semantic invariants may thus be understood as one of the deepest motivations underlying science, mathematics, education, and communication itself. Beneath the diversity of methods and disciplines lies a common aspiration: to discover forms capable of retaining their identity despite transformation. Such forms are valuable not because they eliminate interpretation but because they remain recognizable even when interpretation cannot be avoided.

Scientific writing represents a particularly refined expression of this aspiration. Its characteristic suppression of the personal reflects a deliberate attempt to privilege those aspects of knowledge that survive transport across observers, institutions, cultures, and generations. The resulting text may appear impersonal, but its impersonal character is not an accident. It is the consequence of selecting for structures whose persistence exceeds that of the individuals who first discovered them.

15 Autobiographical Erasure as a Communication Limit

The conventional interpretation of scientific objectivity assumes that personal experience is removed because it introduces bias. While partially correct, this explanation understates the depth of the phenomenon. The preceding analysis suggests a more general interpretation. Autobiographical erasure emerges whenever communication approaches its scaling limits.

As audiences become larger, more heterogeneous, and more temporally distant from the original speaker, the space of possible interpretations expands. The burden of maintaining semantic coherence increases accordingly. Communicators respond by relying upon structures that exhibit greater stability under projection. Informal narratives give way to formal descriptions. Context-dependent meanings are replaced by explicit definitions. Personal trajectories are replaced by public artifacts.

This progression is visible across many domains. Local stories become myths. Myths become doctrines. Doctrines become formal institutions. Informal practices become written procedures. Individual experiences become scientific observations. In each case the same underlying dynamic is present. The communicative object gradually detaches itself from the circumstances of its origin in order to survive transportation through larger systems.

Autobiographical erasure therefore appears not as a peculiar feature of academic culture but as a limiting behavior of communication itself. The more broadly an idea must travel, the greater the pressure to remove those aspects of the idea that depend upon local context. The resulting artifact becomes increasingly independent of the person who produced it.

The irony is that this process often obscures the very experiences that motivated inquiry. Curiosity, confusion, frustration, fascination, and persistence frequently disappear from the final representation. The reader encounters the endpoint without the trajectory. The map survives while the journey vanishes.

Yet this disappearance should not be understood solely as a loss. The same process that obscures the individual also permits knowledge to transcend the individual. Scientific ideas remain accessible long after their creators are gone. Mathematical structures survive the collapse of institutions and civilizations. Formal relationships can be rediscovered, translated, and reconstructed across centuries.

Autobiographical erasure is therefore simultaneously destructive and generative. It removes the personal history of discovery while enabling the discovered structure to achieve a degree of independence from its creator. Knowledge becomes transmissible precisely because it ceases to depend upon the particular circumstances from which it emerged.

The scientist, the teacher, the comedian, the public intellectual, and the software engineer all confront different manifestations of the same challenge. They must construct representations capable of surviving passage through minds other than their own. The techniques differ, but the objective remains constant: the preservation of identity under transformation.

Seen in this light, scientific writing does not merely communicate knowledge. It participates in a broader search for forms that can outlive the interpretive instability of the worlds through which they travel. The disappearance of the personal is not the rejection of human

experience. It is the consequence of selecting those structures that remain when repeated projection has stripped away everything else.

16 The Observer Cannot Be Removed

Consider the simplest possible scientific report: *I saw a lunar eclipse*. At first glance the sentence appears autobiographical, since it reports an event in the life of a particular observer at a particular place and time. Yet it also functions as an observational claim about an external phenomenon, of exactly the kind astronomers have recorded for centuries. The sentence occupies an intermediate category, containing at least three distinct layers: an autobiographical layer (*I was present at a particular place and time*), an observational layer (*a lunar eclipse occurred*), and an evidential layer (*I am reporting direct observation rather than hearsay*).

Scientific writing characteristically removes only the first layer. Instead of *I saw a lunar eclipse*, one writes *a lunar eclipse was observed*, or simply *the eclipse occurred at a recorded time*. The external event survives while the observer disappears. But the observation itself never existed independently of an observer. At some point a person looked through a telescope, read a meter, saw a spectrum, heard a signal, or noticed an anomaly. Scientific writing performs a transformation that strips the observer-context away while preserving the observed event, but it cannot retroactively prevent the event from having been someone's experience in the first place.

This creates a continuum rather than a sharp boundary. At one end: *I felt moved watching the eclipse* – almost entirely autobiographical. At the other: *the eclipse reached totality at a given time, in a given coordinate frame* – almost entirely impersonal. Between them lie countless hybrid statements: *I observed a lunar eclipse; we measured the eclipse duration; astronomers recorded the eclipse; historical observers described the eclipse*. The closer a statement approaches direct observation, the harder it becomes to separate the observer from the observed. This is why laboratory notebooks, field notes, travel journals, and naturalist records remain so revealing: they sit precisely at the boundary between autobiography and science. When Darwin notes in a field journal that he has noticed something unusual about the finches of a particular island, it is simultaneously autobiography and biology. When Galileo records that he has observed moons orbiting Jupiter through a particular instrument on a particular night, it is simultaneously autobiography and astronomy.

The entanglement runs deeper still once the observation is recorded as evidence rather than merely reported as fact. A photograph of a lunar eclipse, for instance, typically carries a timestamp, and the geometry of the eclipse itself constrains the observer to a narrow band of the Earth's surface from which totality was visible at that time. Any landmarks captured incidentally in the frame – a distinctive skyline, a recognizable structure, a particular angle of horizon – further narrow that region, potentially to a single vantage point. The photograph is composed of overlapping observations simultaneously: the eclipse, the landscape, the time, and the presence and location of the observer. The astronomical content – the eclipse itself – may end up being the least identifying component of the image; the incidental context recorded for free alongside it can reveal far more about the observer than the

celestial event that motivated the photograph in the first place. The eclipse is public. The deck from which it was photographed is not.

This is not a flaw unique to amateur observation. It is a structural feature of observation as such. Every scientific fact begins as an autobiographical event in someone's life. Science can erase the observer from the final document, but it cannot erase the fact that an observation originally occurred to a particular person, at a particular time, in a particular place – the process Latour traces in his account of scientific inscription, in which a situated, instrument-mediated act of looking is progressively rewritten into an increasingly decontextualized and citable claim as it moves through a laboratory and into print [7]. The suppression of the observer in scientific prose is therefore not the elimination of autobiography but its compression into invisibility – a transformation

$$(\text{Observer, Observation}) \longrightarrow (\text{Observation}),$$

in which the right-hand side is all that survives in the text, even though the left-hand side is all that ever actually occurred.

17 A Brief Failure of Autobiography

At this point, a reasonable reader might object that the discussion has become excessively abstract. If autobiographical erasure is genuinely important, why not simply include more autobiography?

The suggestion appears straightforward. One could discuss places visited, people encountered, conversations remembered, mistakes made, or unusual experiences accumulated over time. Such material would seem to provide precisely the sort of personal trajectory whose absence has been repeatedly noted throughout this essay.

The difficulty is that autobiography immediately encounters the very communicative constraints under discussion.

Suppose, for example, that one wished to write in detail about a particular trip abroad. The resulting account would not concern a single individual. It would involve family members, friends, colleagues, strangers, guides, employers, fellow travelers, and countless incidental participants. Every story would implicitly contain representations of other people. Every anecdote would establish associations that those individuals neither requested nor necessarily desired.

The problem quickly becomes recursive. Before describing an experience, one must consider not only whether the description is accurate but whether others would consent to the particular representation being constructed. One must consider alternative interpretations, omitted context, incomplete memories, and unintended implications. The act of autobiography becomes increasingly indistinguishable from the management of a complex network of social projections.

As a result, the personal narrative begins to collapse under its own communicative requirements.

The irony is difficult to ignore. An essay concerned with autobiographical erasure discovers that even the attempt to restore autobiography immediately generates pressures toward erasure. The more people involved in a story, the greater the number of possible interpretations. The richer the context, the larger the burden of explanation. The more complete the narrative, the more other narratives become entangled with it.

Consequently, one is tempted to retreat once again toward abstraction. Instead of recounting the particular conversation, one extracts the general principle. Instead of describing the individual, one describes the pattern. Instead of recording the event, one studies the structure that multiple events appear to share.

The result is perhaps disappointing from a literary perspective but highly revealing from an analytical one. The disappearance of autobiography does not always occur because the personal lacks importance. Sometimes it occurs because the personal contains too much information. The trajectory is too entangled, too contextual, too dependent upon the lives of others to travel easily through public communication.

In this sense, abstraction may represent not a rejection of experience but a compromise with it. The theorist does not necessarily abandon the autobiographical. Rather, the autobiographical becomes transformed into a search for invariants capable of surviving once the particulars can no longer be transmitted.

The reader therefore encounters the theory instead of the trip, the pattern instead of the anecdote, the geometry instead of the itinerary – not because the journey never occurred, but because communicating the journey proved substantially more difficult than communicating what appeared to remain constant throughout it. This is, in miniature, the entire phenomenon under investigation: an attempt to talk about a particular place ends up talking about everywhere except that place.

18 Narrative Engineering and the Protection of Others

The difficulties surrounding autobiography are not merely problems of memory. They are also problems of representation. Human experiences rarely occur in isolation. Every personal narrative intersects with the lives of other individuals, each possessing their own perspectives, interests, vulnerabilities, and rights to privacy. Consequently, the act of telling one's own story frequently becomes inseparable from the act of representing others.

This challenge has produced a wide range of narrative techniques across literature, journalism, medicine, psychotherapy, law, and autobiography. These techniques are often understood as ethical or legal safeguards. Yet they can also be interpreted as mechanisms for managing interpretive complexity.

One common strategy involves changing names and identifying details. A memoirist may preserve the structure of an event while altering characteristics that would permit the identification of particular individuals. The resulting account remains recognizable to the author while reducing the risks associated with public representation.

A second strategy involves compositing multiple individuals into a single character. Novelists frequently employ this technique when adapting real experiences into fiction.

Rather than describing several people separately, characteristics are merged into a single representative figure. The resulting character does not correspond precisely to any individual who existed, yet preserves important relational or thematic structures present across multiple encounters.

A third strategy involves temporal compression. Events separated by months or years may be presented as though they occurred within a shorter interval. The objective is not necessarily deception but communicative tractability. The narrative seeks to preserve structural relationships while reducing complexity.

Similar practices appear in professional contexts. Physicians, psychologists, social workers, lawyers, and counselors routinely modify or anonymize case descriptions. Confidentiality obligations prohibit unrestricted disclosure. Case studies therefore become carefully engineered representations that preserve educational value while protecting participants. The resulting narrative is neither wholly factual nor wholly fictional. It occupies an intermediate space designed to balance communicative utility against ethical responsibility.

Power asymmetries further complicate the situation. Individuals occupying positions of authority often possess greater capacity to shape public narratives than those about whom they write. Teachers may describe students. Employers may describe employees. Researchers may describe participants. Journalists may describe subjects. In each case the narrator possesses disproportionate control over representation. Ethical communication therefore requires additional constraints intended to prevent narrative power from becoming representational domination.

These practices reveal an important feature of autobiography. The autobiographical narrative is rarely a direct reconstruction of events. Instead, it is a carefully managed artifact produced through a series of transformations. Names are changed. Identities are blurred. Multiple people become one person. Multiple events become one event. Particular details are omitted while others are emphasized.

The purpose of these transformations is not necessarily concealment. Rather, they allow the narrative to survive public communication without imposing unacceptable costs upon the individuals entangled within it. The resulting text preserves certain structures while sacrificing others.

From the perspective developed throughout this essay, such practices represent another instance of semantic compression. The complete historical trajectory is too complex, too contextual, and too socially interconnected to transmit directly. Narrative engineering therefore seeks a compromise between fidelity and communicability. The story becomes less literal in order to become more transportable.

This observation further complicates the notion of autobiography. Even when an author attempts to write personally, the resulting narrative is often shaped by obligations to privacy, confidentiality, fairness, and social responsibility. The autobiographical self that ultimately appears in print is therefore not simply remembered. It is constructed through a process of selective preservation.

The irony is that the more interconnected a life becomes, the more difficult pure autobiography becomes as well. Every meaningful experience accumulates additional participants,

additional perspectives, and additional ethical constraints. The resulting pressure does not merely encourage anonymity. It encourages abstraction. The author increasingly abandons the attempt to communicate every detail and instead searches for patterns, themes, and structures capable of surviving once the particulars have been removed.

The journey from memoir to theory is therefore not always a movement away from experience. Sometimes it is the consequence of experience becoming too entangled with the lives of others to be transmitted directly. Abstraction emerges not because the world lacks stories but because the stories contain more people than any single narrator has the right to fully explain. A natural extension of this observation, beyond this essay's present scope, is that scientific theories, legal precedents, myths, archetypes, and fictional character types may all be understood as increasingly anonymized compressions of vast numbers of individual experiences. At some point "one person" becomes "a type of person," and eventually "a pattern of interaction." That is nearly the exact trajectory from autobiography to theory.

19 Inference, Metadata, and the Accidental Autobiography

A common assumption is that communication consists of intentionally transmitted information. Closer examination suggests that a substantial portion of communication consists of information that is never explicitly intended at all.

Consider a photograph of a lunar eclipse taken from a residential deck. The apparent subject of the image is astronomical. The photographer wishes to share an observation concerning the Moon, the eclipse geometry, or the visual appearance of a celestial event. At first glance, the image appears largely scientific in character.

Photographs occupy this ambiguous position quite generally. Recent work on inter-medial autobiography has argued that the photograph is never a transparent record but simultaneously a trace, a piece of evidence, and an act of self-representation, accumulating further mediation each time it is reproduced, captioned, or recontextualized [11]. The claim extends a line running back through Barthes, for whom the photograph functions less as a likeness than as a certificate of presence – proof that something stood before the lens, without ever quite delivering the thing itself [2]. The eclipse photograph illustrates the point with unusual clarity: the same image that documents an external celestial event routinely carries unintended information about the circumstances, instruments, and presence of the person who made it.

Yet the photograph contains much more than astronomy. The timestamp associates the observation with a specific moment. The eclipse itself constrains the observation to a particular region of the Earth's surface. The orientation of the Moon constrains viewing direction. Visible landmarks constrain geography. Structures appearing within the image constrain architecture. Environmental features constrain local conditions. Together, these elements may reveal substantially more about the observer than about the eclipse.

What begins as an astronomical observation gradually becomes a form of inadvertent autobiography.

The situation is particularly interesting because none of this information is communicated directly. The image does not state an address. It does not provide coordinates. It does not announce the identity of the photographer. Nevertheless, a sufficiently motivated observer may infer all of these things from information embedded within the observation itself.

The distinction between observation and autobiography therefore begins to dissolve. The eclipse is public information. The location from which the eclipse was photographed may not be. Yet the two become inseparable within the image. The scientific observation acts as a carrier for personal information that was never intended to be the subject of communication.

This phenomenon is not limited to photography. Travel narratives reveal social networks. Research papers reveal institutional affiliations. Case studies reveal professional relationships. Social media posts reveal routines, schedules, habits, and geographic patterns. Even seemingly impersonal observations often contain hidden metadata regarding the circumstances under which those observations were produced.

Under such conditions, publication begins to resemble disclosure. The issue is not necessarily danger, nor even privacy in the conventional sense. Rather, the communicator becomes aware that every observation creates opportunities for inference. Information that appears trivial in isolation may become highly informative when combined with other publicly available information. A photograph, a timestamp, a landmark, and a previously known location may collectively reveal far more than any individual component would suggest.

This creates an unusual asymmetry between communication and interpretation. The communicator experiences the photograph primarily as an image of the eclipse. The observer may experience the same photograph as evidence regarding geography, property, routine, identity, or location. The two parties are therefore not necessarily communicating about the same object at all.

One response to this condition is delay. Rather than suppressing communication entirely, publication may be postponed until the informational value of the associated metadata has decayed. A photograph posted years after it was taken reveals less about current routines than one posted immediately. A travel narrative published long after the journey no longer functions as a real-time record of location. Temporal distance acts as a form of anonymization.

This strategy may be understood as a kind of self-imposed informational shadowing. The observation is preserved while its most sensitive contextual inferences are allowed to dissipate. The objective is not concealment but the reduction of unintended informational coupling between the communicated object and the circumstances of communication.

The resulting behavior may appear peculiar. An individual who is perfectly willing to discuss abstract theories, scientific observations, or public events may hesitate to share a seemingly innocuous photograph. Yet the hesitation becomes understandable once communication is viewed as a process of inference rather than transmission. The photograph communicates not only what the photographer intended to show but also what others may be able to deduce.

The paradox is that the eclipse itself may be among the least personal elements within the

image. The celestial event belongs equally to everyone. The deck, the landmarks, the timing, and the circumstances of observation belong to a particular life. The observer therefore discovers that sharing a picture of the Moon can feel strangely similar to publishing a partial map of one’s own existence – which is itself a perfectly compact illustration of the essay’s central theme: the author may be entirely comfortable publishing extensive theoretical work on cosmology, admissibility, or formal ontology, yet hesitate over a single photograph of the Moon, because the railing in the background carries more identifying information than the eclipse that motivated the photograph in the first place.

20 Erasure and Exposure: Autobiography as Negotiated Projection

Recent work in autobiography studies arrives at a strikingly similar conclusion from an entirely different direction. In her dissertation *Between Erasure and Exposure: Intermedial Autobiography Since Roland Barthes*, Van Arsdall argues, drawing on Barthes’s semiotics of the image and text [1, 2] and subsequent French photographic theory, that autobiographical media do not simply reveal a self, nor do they simply conceal one. They produce a self through partial disclosure, displacement, substitution, and mediation – a continual negotiation between what is shown and what is withheld [11]. The present essay has been developing a parallel argument through information theory, attribution, metadata, and pseudonyms rather than through photographic theory and the criticism of self-representation, but the destination is recognizably the same.

The vocabulary of erasure and exposure gives a name to a structure that has been implicit throughout the preceding sections. Let S denote the latent space of the self – memories, bodily states, relations, locations, histories, intentions, and unarticulated constraints – and let A denote an autobiographical artifact produced from it by some projection $\pi : S \rightarrow A$. *Exposure* is the information about S preserved by π ; *erasure* is the information about S discarded by π . Neither process can occur in isolation. If nothing is erased, communication is impossible, since the totality of a life cannot be transmitted through any finite artifact. If nothing is exposed, communication fails for the opposite reason: no observable representation survives at all. Autobiography therefore does not choose between revealing and concealing the self. It occupies the regime strictly between these two degenerate limits,

$$0 < E_{\pi}(S) < H(S),$$

in which some dimensions of the self are exposed precisely because others are simultaneously erased. (A formal treatment of $E_{\pi}(S)$, its complement, and the consequences of this inequality is given in the appendix.)

This reframes a recurring theme of this essay. Writing, photography, pseudonyms, metadata, and digital signatures are not simply different media for the same content. They are different projections of the same latent self, each preserving a different invariant and erasing a different remainder. A photograph preserves visual trace while erasing intention. A caption preserves interpretation while erasing bodily presence. A pseudonym preserves continuity of authorship while erasing legal identity. A cryptographic hash preserves document identity while erasing semantic content. None of these artifacts is the self. Each

is one projection among many, and the autobiographical object that a reader reconstructs is not any single artifact but the overlap among several: a self inferred from the intersection of the possible selves compatible with each available trace.

Van Arsdall's account is particularly suggestive on this last point. She repeatedly describes photographs not as direct records of a self but as "traces of traces" – objects that point toward a presence without ever fully delivering it, and that themselves accumulate further mediation each time they are reproduced, captioned, or recontextualized. This is close to the structure this essay has already proposed for pseudonyms, hashes, and writing style, and will propose again below for filter-order signatures: a pseudonym is a trace; a cryptographic hash is a trace; a profile image built from publicly available filters is a trace; a hierarchical filter application order is a trace; a writing style is a trace; and authorship itself, as argued throughout this essay, is never observed directly but reconstructed from the accumulation of such traces. Erasure and exposure are therefore not opposites in tension but coupled operations of a single process: every exposure narrows the space of selves compatible with the artifact, and every erasure preserves the ambiguity within which a person remains, to some degree, irreducible to any one of their traces.

21 Names, Pseudonyms, and the Limits of Attribution

The discussion of autobiographical erasure naturally raises a related question concerning authorship itself. If communication involves projection, compression, and inference, then how does one determine whether multiple artifacts originate from the same source? The answer appears straightforward until one examines the limitations of names.

Names are commonly treated as identifiers, yet they function imperfectly in this role. Multiple individuals may share the same legal name. A single individual may publish under different names. Names may change through marriage, migration, transliteration, cultural adaptation, or personal preference. Consequently, the relationship between a name and a person is neither unique nor permanent.

Pseudonyms introduce an additional layer of complexity. They are often interpreted as mechanisms of concealment, yet they may simultaneously improve attribution. A distinctive pseudonym consistently attached to a body of work can function as a more reliable reference than a common legal name. The pseudonym acts as a namespace linking documents, publications, repositories, correspondence, and ideas across time. In information-theoretic terms, the legal name says, in effect, *this work came from this person*, while the pseudonym says *these works belong to the same trajectory* – and the second is often closer to what readers actually care about. Most readers are trying to follow a chain of ideas through time, not reconstruct a birth certificate.

However, this apparent advantage should not be overstated. A pseudonym is not an identity. It is merely another signal from which identity may be inferred.

Readers frequently assume that different names imply different individuals. A legal name and a pseudonym may therefore be interpreted as distinct authors despite originating from the same source. The reverse problem also occurs. Different individuals may be

incorrectly assumed to be the same person because they share a name, a writing style, a field of study, or a set of interests. Attribution therefore remains an inferential activity rather than a direct observation.

The situation becomes even more complicated once incentives enter the system. A pseudonym functions effectively as an identifier only insofar as it remains difficult or unprofitable to imitate. If a name acquires sufficient reputational value, economic value, or social significance, incentives emerge for spoofing, imitation, impersonation, and strategic association.

This phenomenon resembles the dynamics described by Goodhart's Law. A measure may function effectively as a proxy for an underlying property until it becomes the target of optimization. Once the proxy itself becomes valuable, behavior shifts toward manipulating the proxy rather than preserving the underlying property it was intended to represent. A pseudonym may initially function as a useful indicator of authorship; as recognition increases, the identifier itself becomes a target. The signal that once served as evidence begins attracting attempts at replication, appropriation, or manipulation. The marker remains useful, but its reliability gradually declines when considered in isolation. A pseudonym functions well as a signature only so long as the cost of imitation exceeds the benefit of imitation; once that relationship changes, the pseudonym becomes vulnerable to the same dynamics affecting every other reputation system – signatures, seals, trademarks, diplomas, academic credentials, and verification badges have all followed the same arc, from reliable indicator, to valuable credential, to target of imitation, to diminished reliability.

For this reason, attribution rarely depends upon names alone. Readers implicitly rely upon a much richer collection of signals. Writing style, recurring themes, preferred terminology, publication history, technical interests, citation patterns, conceptual frameworks, and temporal continuity all contribute to judgments of authorship. Identity emerges not from a single identifier but from a pattern distributed across many observations.

This observation aligns closely with the broader argument of the present essay. Just as communication involves reconstructing meanings from incomplete evidence, attribution involves reconstructing authorship from incomplete evidence. In neither case is the underlying source observed directly. Instead, observers infer a hidden structure from a collection of traces.

The legal name, the pseudonym, the autobiography, the scientific paper, and the public persona all participate in this process. Each functions as a projection of a more complex underlying trajectory. None is identical to the person. Each preserves certain features while discarding others.

The resulting picture is neither one of complete anonymity nor complete transparency. Identity appears instead as a problem of continuity. Observers seek evidence that multiple artifacts, statements, and actions belong to the same underlying source. Names assist in this task, but they do not solve it. The continuity itself resides not in the identifier but in the trajectory that the identifier imperfectly represents.

From this perspective, a pseudonym should not be understood primarily as a disguise. It is better understood as a communicative artifact designed to preserve continuity across contexts. Like every other artifact discussed in this essay, it succeeds by compression. It

preserves enough structure to permit recognition while discarding information judged irrelevant to the task at hand. The resulting representation is necessarily incomplete, yet its incompleteness is precisely what allows it to function.

22 Pseudo-Identity and Collision Tolerance

A common assumption underlying discussions of identity is that identifiers uniquely determine the entities to which they refer. Closer examination reveals that many practical systems operate according to a weaker principle. Rather than eliminating ambiguity entirely, they merely reduce ambiguity to a sufficiently low level that communication can proceed effectively.

Modern computing provides numerous examples. Unix timestamps are routinely used to label events, files, transactions, and observations. A timestamp records the number of elapsed seconds relative to a reference epoch and therefore provides a compact reference for locating an event within a temporal sequence. Yet timestamps are not unique identifiers. Multiple events may occur during the same second. The timestamp functions successfully not because collisions are impossible but because collisions are often rare enough to be ignored.

Cryptographic hashes exhibit a similar structure. Given a document d , a hash function produces a digest $h = H(d)$. In practical systems the digest is frequently treated as though it were identical to the document itself. Software releases are verified through hashes. Files are deduplicated through hashes. Version-control systems identify commits through hashes. Entire histories become associated with compact hexadecimal strings.

Nevertheless, a hash is not an identity. The hash is a compressed representation whose utility derives from the improbability of collision rather than the impossibility of collision. Distinct documents may satisfy $H(d_1) = H(d_2)$ while remaining entirely different objects. The identifier therefore functions as a statistical approximation to identity rather than a perfect realization of it.

This distinction proves surprisingly relevant to human communication. Names, signatures, pseudonyms, professional titles, institutional affiliations, and writing styles all function in a similar manner. None uniquely determines a person. Instead, each reduces uncertainty regarding the underlying source. Communication proceeds because the probability of mistaken attribution is sufficiently small, not because ambiguity has been completely eliminated.

The resulting picture differs substantially from traditional notions of identity. Rather than being a singular marker attached to an individual, identity becomes a probabilistic reconstruction assembled from multiple traces. A name contributes evidence. A writing style contributes evidence. A publication history contributes evidence. A pseudonym contributes evidence. None is decisive in isolation. Together they reduce the probability of collision.

This observation suggests that many forms of identity should be understood as pseudo-identities. They are practical solutions to attribution problems rather than metaphysical

guarantees of uniqueness. The objective is not to establish perfect correspondence between marker and source. The objective is to reduce the set of plausible alternatives to a manageable size.

The same principle appears throughout scientific communication. Theories are identified by names despite evolving over time. Species are identified by classifications despite biological variation. Historical figures are identified by records despite incomplete documentation. In each case observers operate under conditions of uncertainty and rely upon markers whose reliability is statistical rather than absolute.

From this perspective, identity resembles an engineering problem more than a philosophical essence. The question is not whether collisions are possible – collisions are almost always possible. The relevant question is whether collisions occur frequently enough to interfere with the intended task.

A pseudonym therefore occupies a position analogous to a cryptographic digest. It does not uniquely determine its source in any absolute sense. Rather, it acts as a compact reference to a larger trajectory of artifacts, ideas, and behaviors. Recognition occurs because the probability that an unrelated trajectory would generate the same collection of signals remains sufficiently small. Identity emerges not from perfect uniqueness but from the practical management of collision risk.

23 Identity Through Transformation Order

An instructive example of pseudo-identity arises from the construction of ordinary digital artifacts. Consider a profile image produced from publicly available components: a source photograph, an edge-detection filter, a color-gradient mapping, and a steganographic embedding technique. Each operation, in isolation, is commonplace. A photograph of an everyday object is not unique. A Sobel edge detector is not unique. A red–blue gradient is not unique. Least-significant-bit steganography is not unique. None of these components, considered alone, provides a reliable identifier.

Suppose the image is constructed through the composed sequence

$$I = F_4 \circ F_3 \circ F_2 \circ F_1(X),$$

where F_1 extracts a photographic subject, F_2 applies an edge-detection filter, F_3 applies a color transformation, and F_4 embeds additional information through a steganographic technique applied to the result. Each operation is publicly known and freely available. Yet the resulting image may nevertheless become highly distinctive.

The reason is that these transformations generally fail to commute: for two operations F_A and F_B in general

$$F_A \circ F_B \neq F_B \circ F_A.$$

Applying the edge detector before the color mapping produces a different result than applying the color mapping before the edge detector. Embedding information before edge extraction produces a different result than embedding it afterward. Consequently, the order in which operations are applied becomes informative in its own right. The signature

of the artifact is not contained in any individual operation but in the trajectory through transformation space – what might be called a *hierarchical filter application order*.

This observation generalizes well beyond image processing. Writing, scientific reasoning, artistic production, and personal identity all exhibit similar properties. Individual components are frequently shared across an entire community: vocabulary, concepts, techniques, and conventions belong to populations rather than individuals. What distinguishes one trajectory from another is often not any particular ingredient but the characteristic ordering through which these elements are selected, transformed, combined, and refined.

A handwriting expert does not search for perfect repetition – perfect repetition is, if anything, a sign of forgery. The expert searches for consistent deviation: the systematic, slightly idiosyncratic way an individual departs from the population norm in stroke pressure, letter spacing, or rhythm. A painter’s signature is often not the subject matter at all but small, repeated irregularities in brushwork, in the treatment of edges, in the transition between colors. If P denotes the population distribution of some feature and P_i the distribution characteristic of a particular author, the relevant signal is the divergence $D(P_i \parallel P)$, not the overlap. Identity, in this sense, lives in the residue.

Identity therefore need not depend upon unique ingredients. It may emerge from unique compositions of non-unique ingredients. The resulting structure resembles a cryptographic hash generated from common symbols: the distinctiveness arises not from any particular component but from the improbability that another trajectory would reproduce the same sequence of transformations. The deepest identifier of a system may not be what it contains but how it transforms what it contains. In this sense, authorship, artistic style, and even personal identity can be understood as histories of transformation rather than collections of properties – a trajectory, once again, carrying more information than any single state drawn from it.

24 Authorship as an Inference Problem

The question of authorship is often treated as though it were fundamentally different from the question of interpretation. In practice, the two problems are remarkably similar. Just as readers reconstruct meanings from texts, they also reconstruct authors from texts. In neither case is the underlying object directly observed. Instead, observers infer a hidden source from a collection of visible traces.

This process becomes particularly evident when autobiographical markers are sparse. Traditional autobiographical writing continually reinforces the existence of a particular narrator. The text contains names, places, dates, relationships, memories, and experiences. The reader is repeatedly reminded that the work emerges from a single life. The identity of the author therefore becomes part of the content itself.

Scientific and theoretical writing operate differently. Personal details are minimized. Narrative chronology disappears. Attention shifts toward arguments, models, observations, and abstractions. As autobiographical information is progressively removed, the text increasingly resembles an autonomous object. The reader encounters the structure while

possessing relatively little information regarding the circumstances of its production.

Under these conditions, authorship itself becomes an inference problem – the same problem Foucault poses when he asks not who an author really was but what the “author function” does: a principle by which a body of texts is grouped, delimited, and granted a certain mode of existence, independent of any direct access to the writing subject [5].

Readers begin examining indirect evidence. They infer authorship from writing style, recurring terminology, conceptual preferences, citation patterns, technical sophistication, thematic continuity, and the internal coherence of the work. The resulting judgments are necessarily probabilistic. The text may suggest a particular type of author without uniquely identifying one.

This uncertainty becomes especially pronounced when a work spans multiple domains. A reader encountering discussions of mathematics, philosophy, linguistics, computer science, psychology, cosmology, and engineering within a single body of work may naturally wonder whether the material originated from one individual or from a group. The breadth of the artifact exceeds ordinary expectations regarding specialization, encouraging alternative hypotheses regarding its source.

Interestingly, such uncertainty is often produced by the very processes discussed throughout this essay. The suppression of autobiographical detail removes many of the cues readers ordinarily use to construct a model of the author. The work becomes easier to evaluate independently of its creator, but the creator simultaneously becomes more difficult to reconstruct.

The irony is that successful autobiographical erasure may produce uncertainty not merely regarding the author’s life but regarding the existence of the author as a singular entity. Readers begin asking whether they are encountering a person, a collaboration, an institution, a collective pseudonym, or an entire research group. The more the work functions as a self-contained structure, the less obvious its origin becomes. A reviewer who remarks that a body of work appears to have been produced either by a single individual or by a collective, adding that the distinction is difficult to determine from the text alone, is in this sense performing the precise operation that this essay describes: not reading a biography, but reconstructing a hidden trajectory from observable traces, and discovering that the trajectory admits multiple explanations.

This phenomenon reveals an important asymmetry. Scientific writing is often assumed to remove personal identity in order to foreground ideas. Yet removing personal identity does not eliminate the author’s presence. Instead, it transforms the author’s presence into an object of inference. The reader ceases to encounter the individual directly and begins reconstructing the individual from the properties of the artifact.

In this respect, authorship resembles many other hidden structures discussed in the present essay. A scientific theory is inferred from observations. A personality is inferred from behavior. A location is inferred from metadata. An intention is inferred from actions. An author is inferred from texts. In every case, observers reconstruct an underlying source from partial evidence.

The resulting situation produces a curious reversal. The less autobiographical a work

becomes, the more readers may find themselves wondering about the person who produced it. The disappearance of explicit autobiography does not remove the author from consideration. It simply changes the method by which the author is encountered. One might therefore regard authorship itself as another invariant-seeking process. Readers attempt to identify a stable source responsible for a collection of artifacts distributed across time. Names, pseudonyms, writing styles, conceptual habits, and publication histories all function as evidence in this reconstruction. None is sufficient on its own. Together they form a trajectory from which a hidden source may be inferred.

The question “Who wrote this?” therefore turns out to be structurally similar to the question “What does this mean?” Both involve the reconstruction of an unseen object from visible traces. Both remain vulnerable to ambiguity. Both depend upon incomplete information. And both reveal that communication is never merely the transmission of content. It is also the continuous generation of hypotheses regarding the sources from which that content emerged.

25 Authorship Attribution as Collision Detection

The preceding discussion suggests that identity should not be understood as a binary property but as a problem of discrimination among competing sources. This perspective permits a useful reformulation of authorship itself. Rather than asking whether two documents originate from the same author, one may ask whether the observable features of the documents are sufficiently unusual that alternative explanations become improbable.

Let D denote the space of documents and let $\phi : D \rightarrow \mathbb{R}^n$ be a feature extraction map. Components of $\phi(d)$ may include word frequencies, sentence lengths, punctuation habits, citation patterns, preferred examples, recurring metaphors, topic transitions, conceptual structures, and other measurable characteristics. Each author then generates documents occupying a region of feature space; let μ_i represent the characteristic feature distribution associated with author i . Authorship attribution becomes the problem of determining whether two documents d_1, d_2 could plausibly have originated from the same underlying distribution.

The conventional intuition is that authorship is established through similarities. However, similarities alone are rarely informative. Most writers use common grammatical constructions, common vocabulary, and common rhetorical devices. Such features contribute little evidence because they are shared by large populations. Instead, attribution derives primarily from persistent deviations from population norms. Let P denote the background distribution of features in a reference population and let P_i denote the feature distribution associated with a particular author. The distinctiveness of the author may be measured by a divergence $D(P_i || P)$. As this divergence increases, accidental collisions become less likely.

In this formulation, authorship attribution resembles collision detection in a high-dimensional space. Two documents may appear superficially similar while originating from entirely different sources. Conversely, documents that differ substantially in subject matter may still reveal a common origin through persistent structural regularities.

This observation becomes particularly important when considering pseudonyms and collective authorship. A pseudonym does not establish identity directly. Instead, it serves as a label attached to a region of feature space. Readers gradually learn to associate certain patterns with that label. Attribution succeeds when newly encountered documents occupy approximately the same region.

The possibility of collision can never be eliminated entirely. Distinct authors may occasionally produce similar outputs. A sufficiently skilled imitator may deliberately attempt to reproduce another author's style. Yet the probability of successful imitation decreases as additional dimensions of evidence are incorporated. The resulting process closely resembles modern cryptographic practice. A short identifier may collide relatively easily; a long identifier incorporating many independent features becomes progressively more reliable. Likewise, a single word choice provides little evidence regarding authorship, whereas thousands of correlated stylistic decisions accumulated across many documents become increasingly difficult to reproduce accidentally.

This perspective suggests that authorship should be understood not as a property of names but as a property of trajectories. The underlying object being identified is not a legal individual, a pseudonym, or a signature. Rather, it is a persistent generative process producing a sequence of artifacts over time.

This framework also clarifies why certain forms of imitation remain difficult even when superficial features are copied successfully. An imitator may reproduce vocabulary, formatting, notation, or subject matter. Yet deeper regularities often persist beneath these visible layers: the order in which distinctions are introduced, the progression of abstractions, the choice of examples, the preferred modes of explanation, and the sequence of conceptual transformations. The most distinctive signatures, in other words, may not be found in particular words or symbols at all – they reside in the hierarchical filter application order discussed in the preceding section, the characteristic sequence through which a cognitive system transforms observations into communicable artifacts. Such structures are difficult to imitate precisely because they emerge from long developmental trajectories rather than isolated stylistic choices.

Authorship therefore appears not as a static identity but as a dynamical process. The signature of a thinker is not a mark attached to a document after its completion. It is the persistent pattern of transformations that repeatedly generates documents possessing recognizable structure. The author's identity is encountered not through direct observation but through the statistical improbability that another trajectory would produce the same constellation of deviations across so many independent dimensions.

26 The Autobiography of Constraints

If the argument of this essay is correct, then a natural objection presents itself. Stripped of chronology, family history, emotional confession, and the ordinary furniture of memoir, does the present text actually tell the reader anything about its author at all?

It does – but not in the form expected. What it withholds in events, it supplies in con-

straints. Throughout this essay, the reader has not been told where the author was born, what was studied, whom the author has known, or what has been felt at any particular moment. Yet the reader has learned, indirectly and cumulatively, a great deal: which examples were chosen and which were avoided; which anxieties recur (interpretive instability, accidental disclosure, the entanglement of other people's lives with one's own); which communication strategies are practiced (delayed publication, narrative compositing, the preference for pseudonym over legal name, the retreat into formalism when narrative becomes too risky to transmit); and which intellectual commitments persist across the entire discussion (a fascination with projection, compression, invariance, and the geometry of what survives transformation).

None of this is autobiography in the conventional sense of *here is what happened to me*. It is autobiography in a different and arguably more revealing sense: *here is how I decide what counts as communicable*. A memoir exposes a sequence of events. This essay exposes the filter through which events would have to pass before they could appear in print at all. The two are not the same kind of disclosure, but they are not unrelated either – the filter itself is the residue of a life, compressed not into a chronology but into a set of standing constraints.

This is, in fact, exactly what the essay's own central claim would predict. If autobiography is repeatedly transformed under pressure into structure, then a sufficiently long meditation on *why* autobiography becomes structure should itself contain a recoverable trace of the person doing the meditating – not as content, but as the shape of the selection process. The autobiography was never simply omitted. It moved from the space of events into the space of constraints governing which events, and which kinds of events, are permitted to become visible.

27 Why Theorists Drift Toward Geometry

The preceding sections suggest a final, more general claim, toward which the essay has been moving from its opening pages. Different classes of representation differ systematically in how much of their content survives transformation, and this difference is not incidental to the choice between them – it is frequently the reason for the choice.

Narratives are highly observer-dependent. Their meaning relies on shared background, emotional resonance, cultural assumption, and tacit context that cannot be fully specified in advance; remove the context and the narrative degrades, sometimes badly. Categories and types are somewhat less observer-dependent: a narrative compressed into an archetype, a case study, or a recurring pattern travels further because much of its local detail has already been discarded. Formal structures – definitions, logical relations, models – are less observer-dependent still, since their validity is meant to be checkable from the structure itself rather than from shared history with the author. Geometric and mathematical invariants occupy the far end of this spectrum: a proof, a symmetry, an invariant quantity remains recognizable under translation, reformulation, change of notation, and reinterpretation by readers separated from the author by centuries and by entirely different conceptual vocabularies.

A theorist who has repeatedly encountered the earlier stages of this essay's argument in

their own working life – interpretive drift even between two people, exponential ambiguity under scale, statements detached from context and reattached to unintended narratives, autobiography that implicates other people before it can be told – has, in effect, repeatedly observed the fragility of the observer-dependent end of this spectrum and the comparative robustness of the other end. The drift toward mathematics, geometry, logic, and formal systems is therefore not best explained as a personality trait, an aesthetic preference, or a rejection of human experience. It is better explained as a rational response to a measured asymmetry: formal structures are simply more likely to remain recognizable after the transformations that communication inevitably imposes.

This is also, not incidentally, a fair description of why a constraint-first approach to ontology – one organized from the outset around what distinctions, projections, and reachability relations remain admissible under transformation, rather than around any particular observer’s narrative of how those structures came to be noticed – holds a natural appeal for someone who has lived through the earlier arguments of this essay. A framework built from history, distinction, and admissibility before object, state, and prediction is, among other things, a framework built to survive exactly the kind of projection this essay has been describing. The theorist does not migrate toward geometry because the personal has become uninteresting. The theorist migrates toward geometry because geometry is what remains visible after the personal has been forced, again and again, through a channel too narrow to carry it intact.

28 Admissible Self-Projection

The three formal results introduced across this essay and its appendix are not three separate achievements but a single progression, and it is worth stating that progression explicitly before drawing the essay to a close. The Projection Complexity Theorem establishes the *cost* of communication: why explaining a trajectory to an audience becomes disproportionately expensive once a shared interpretive frame can no longer be assumed, with the cost passing from linear to quadratic in audience size at exactly the point where interpretations must be reconciled with one another rather than merely received. The Narrative Chain establishes the *consequence* of that cost: faced with a channel that cannot carry the full burden, successive acts of communication compress the self through a sequence of projections whose information content is provably monotone decreasing, terminating in the symbolic, low-dimensional public figure each of us eventually becomes to people who do not know us directly. The Admissibility result, developed in this section and formalized in the appendix, establishes what is *optimal* under those two facts taken together: given that communication is costly and that compression is consequently unavoidable, what is the best compression available, subject to surviving reinterpretation and remaining fair to the people entangled in the story? In short,

$$\text{Complexity} \longrightarrow \text{Compression} \longrightarrow \text{Admissibility,}$$

where each arrow names a theorem rather than a metaphor. The first explains why communication is hard. The second explains why hard communication compresses. The third

explains what a good compression looks like once the first two facts are accepted as unavoidable. Read in this order, the three results stop looking like an appendix of loosely related propositions and start looking like a single argument: that autobiography, public identity, and scientific writing are successive points on one constrained optimization, not three unrelated phenomena that happen to share some vocabulary.

Taken together, the preceding sections describe an implicit optimization problem that every autobiography, public talk, memoir, social-media post, scientific paper, and interview is solving, whether or not its author would ever put it this way. Informally: maximize how much of the self is exposed, subject to the constraint that the resulting artifact survive reinterpretation, and subject to the further constraint that it remain socially admissible – that it not impose unacceptable disclosure costs on the people entangled within it, nor cross whatever privacy thresholds the author has set. In the notation developed throughout this essay, the problem is

$$\max_{\pi} E_{\pi}(S) \quad \text{subject to} \quad R_G(\pi(S)) > \tau \quad \text{and} \quad D_p(\pi(S)) < \varepsilon,$$

where $E_{\pi}(S)$ is autobiographical exposure, $R_G(\pi(S))$ is robustness under the transformations G that reinterpretation, audience scale, and time will impose, and $D_p(\pi(S))$ is the disclosure cost to any other person p entangled in the representation. (The appendix gives each term a precise definition and proves the proposition this section only states.)

The solution to this problem is never complete disclosure. An uncompressed self – π equal to the identity, full exposure, no compression at all – violates the robustness constraint almost immediately, since it is precisely the representation most vulnerable to misreading, recontextualization, and the entanglement of other people’s lives. The solution is instead a point on what might be called the admissibility manifold: the surface of projections that achieve the best exposure available without crossing either constraint. Memoir, scientific paper, pseudonymous essay, and social-media caption are not different kinds of self-representation so much as different solutions to the same constrained optimization, evaluated at different settings of τ and ε . A curriculum constrains G in advance. A disclaimer relaxes the effective threshold τ by reframing the audience’s interpretive stance before the content arrives. Narrative compositing lowers D_p at some cost to E_{π} . Every device discussed in this essay turns out to be a way of moving along, or reshaping, this same manifold.

Viewed this way, the relationship between autobiography and scientific writing proposed at the start of this essay can finally be stated precisely. Scientific writing is not a separate genre standing apart from autobiography. It is the extreme point of the admissibility manifold at which robustness under reinterpretation is prioritized almost completely over autobiographical exposure – the corner at which τ is set as high as the subject matter allows, and $E_{\pi}(S)$ is correspondingly driven down to whatever residue remains compatible with making a claim at all. Memoir sits near the opposite corner, where exposure is prioritized and robustness is risked. Most communication occupies the interior of the manifold, trading a little of each constraint against the other as audience, stakes, and the people involved

demand. The claim that closes this essay can therefore be put in a single line:

Autobiography = Admissible Self-Projection.

The self is never transmitted directly. Only admissible projections of the self are transmitted – and what has looked, across the sections of this essay, like a long series of unrelated case studies (eclipse photographs, pseudonyms, disclaimers, public personas, scientific papers) turns out in every instance to be the same constrained optimization, solved under different pressures, in a different corner of the same manifold.

29 The Autobiography Was Never Missing

This essay began with a puzzle. Scientific writing often appears strangely detached from the lives of the people who produce it. The conventions of academic prose suppress personal narrative, minimize subjective experience, and direct attention toward structures that can be evaluated independently of their authors. The resulting literature can seem almost anonymous. Ideas appear to emerge from nowhere. Observations appear without observers. Theories appear without theorists.

The preceding chapters have argued that this phenomenon cannot be explained solely through appeals to objectivity. Autobiographical erasure emerges from a convergence of deeper pressures. Communication occurs between distinct interpretive systems. Meaning drifts under projection. Audiences vary in scale and composition. Statements acquire unintended associations. Personal narratives become entangled with the lives of others. Observations reveal information beyond their intended subjects. Public communication compresses complexity into portable representations. Scientific discourse therefore participates in a broader search for forms capable of surviving transformation.

Yet the essay has also revealed a paradox. While repeatedly discussing the disappearance of autobiography, it has continuously introduced autobiographical material. References to travel, photography, publication practices, pseudonyms, audience management, privacy concerns, public communication, scientific observation, and interpretive uncertainty all reveal aspects of the individual producing the text. The reader encounters no conventional chronology, yet gradually acquires a sense of the concerns, habits, and intellectual trajectory from which the essay emerged.

The apparent contradiction dissolves once autobiography is understood as a problem of representation rather than disclosure.

A traditional autobiography attempts to preserve events. The present essay preserves constraints. A memoir records what happened. The present essay records the conditions under which events become communicable. One approach organizes experience chronologically. The other organizes experience structurally. Both remain autobiographical, but they operate through different forms of compression.

This distinction suggests that autobiography is not eliminated by abstraction. Instead, abstraction transforms the medium through which autobiography appears. The personal ceases to be expressed through narrative sequence and begins to appear through recurring

patterns of attention. The reader learns less about individual episodes and more about the criteria governing which episodes become visible. The resulting portrait is indirect, yet often surprisingly revealing.

The same observation applies to scientific writing more generally. A scientific paper may omit family history, travel experiences, political opinions, personal relationships, and emotional responses. Nevertheless, it often reveals a great deal about the intellectual landscape from which it emerged. The problems chosen, the assumptions adopted, the analogies employed, the methods preferred, and the questions considered important all provide traces of the observer who produced the work. The scientist disappears from the text only in a limited sense. The disappearance is never complete.

Indeed, complete erasure may be impossible. Every observation originates somewhere. Every theory addresses a particular class of problems. Every abstraction emerges from a history of encounters with the world. Even the most formal mathematical structure reflects decisions regarding what should be preserved and what should be ignored. These decisions are not arbitrary. They emerge from trajectories of attention accumulated across time.

The search for invariants therefore contains an unexpected autobiographical dimension. The structures that an individual chooses to preserve reveal something about the transformations they have experienced. The theorist who becomes fascinated by geometry, the scientist who seeks reproducibility, the teacher who emphasizes definitions, and the writer who worries about interpretation are all responding to particular encounters with instability. Their abstractions function simultaneously as descriptions of the world and records of the problems they have attempted to solve within it.

The argument of this essay can therefore be summarized in a simple inversion. Scientific writing does not eliminate autobiography. It relocates it. The autobiography moves from events to constraints. It moves from chronology to structure. It moves from narrative to invariance. It moves from the story itself to the principles governing what can be said about the story.

The reader who searches for personal details may therefore conclude that the autobiography has disappeared. The reader who searches for recurring concerns, persistent constraints, preferred abstractions, and chosen invariants discovers something quite different. The autobiography was never absent. It was transformed.

Scientific writing, mathematical reasoning, pseudonyms, public communication, disclaimers, educational curricula, and formal theories all participate in this transformation. Each reduces certain dimensions of experience in order to preserve others. The resulting artifacts may appear impersonal, yet they remain connected to the trajectories that produced them. What survives is not the entire history of a life but those structures that proved sufficiently stable to persist after repeated projection through time, interpretation, and communication.

The final irony is that this essay itself has served as an example of the phenomenon it set out to explain. An inquiry into autobiographical erasure has gradually become autobiographical – not because it abandoned abstraction, but because abstraction revealed the constraints through which a life becomes communicable. The narrative was never entirely removed. It was compressed into the geometry of its own preservation.

A Mathematical Formalism of Autobiographical Erasure

The propositions below give a formal – and deliberately schematic – expression to the arguments of the preceding sections. They are not offered as empirical claims but as a compact notation for relationships that have so far been stated in prose: that communication is lossy by construction, that ambiguity compounds with audience size, that scientific writing is a specific projection that discards an observer-context while preserving an event, and so on.

Let X denote the space of lived events, observations, memories, relations, and contextual dependencies available to an individual. A communicative act is modeled as a projection

$$\pi_A : X \longrightarrow Y_A,$$

where Y_A is the representational space available to an audience A . The projection π_A is never injective in realistic cases, since no audience receives the totality of the speaker’s life, context, memories, and intentions. The autobiographical loss associated with the projection is therefore

$$L_A(x) = \log \left| \pi_A^{-1}(\pi_A(x)) \right|,$$

where $\pi_A^{-1}(\pi_A(x))$ denotes the set of possible underlying events compatible with the communicated artifact.

Definition A.1 (Interpretive Variance). *Given an utterance $y \in Y_A$ and an audience $A = \{a_1, \dots, a_n\}$, let $I_i(y) \in Z_i$ denote the interpretation reconstructed by audience member a_i . The interpretive variance of y over A is*

$$\text{Var}_A(y) = \frac{1}{n} \sum_{i=1}^n d_Z(I_i(y), \bar{I}(y))^2,$$

where d_Z is a semantic distance and $\bar{I}(y)$ is an appropriate Fréchet mean of the interpretations.

Proposition A.1 (Communication Begins with Loss). *If π_A is non-injective, then there exist distinct lived contexts $x_1, x_2 \in X$ such that $\pi_A(x_1) = \pi_A(x_2)$. Thus the communicated artifact cannot uniquely determine the lived event that produced it.*

Proof. Non-injectivity means precisely that there exist $x_1 \neq x_2$ with identical image under π_A . Since the audience receives only $\pi_A(x)$, both x_1 and x_2 remain compatible with the same communicative object. Hence communication necessarily loses information about the originating context. \square

Proposition A.2 (Scaling of Interpretive Ambiguity). *Suppose each interpreter admits at least $k > 1$ plausible reconstructions of an utterance. For n independent interpreters, the number of joint interpretive configurations is bounded below by k^n . Interpretive ambiguity therefore grows exponentially with audience size under independence.*

Proof. If each of n interpreters admits at least k possible interpretations, then the product space of joint interpretations contains at least $\underbrace{k \cdot k \cdots k}_{n \text{ times}} = k^n$ configurations. Therefore the space of possible collective reconstructions grows exponentially. \square

The preceding proposition bounds the number of possible interpretations. A related but distinct question is how much the speaker must actually *say* – the explanatory cost of communicating a trajectory to an audience, rather than the size of the space of ways it might be misread. The following result gives that cost a precise growth rate, and identifies the exact condition under which it changes character.

Definition A.2 (Autobiographical Burden). *Let $T = (e_1, e_2, \dots, e_n)$ denote a life trajectory: a finite sequence of events, observations, and contextual dependencies. For an audience $A = \{a_1, \dots, a_m\}$, let $\pi_i : T \rightarrow R_i$ denote the projection of T required to produce a representation admissible to audience member a_i , and let $L(\cdot)$ denote the explanatory complexity of a representation – the length, in some fixed encoding, of the material the speaker must supply for a_i to reconstruct $\pi_i(T)$. The autobiographical burden of communicating T to A is*

$$B(T, A) = \sum_{i=1}^{|A|} L(\pi_i(T)).$$

Proposition A.3 (The Projection Complexity Theorem). *If every π_i coincides with a single shared projection π , so that no member of A requires an explanation tailored to any other member's interpretation, then $B(T, A) = O(|A|)$. If instead the audience reaches a scale at which interpretations diverge and the speaker must additionally reconcile each pair of divergent interpretations for the communication to remain coherent, then $B(T, A) = O(|A|^2)$.*

Proof. Linear regime. Under the shared-projection hypothesis, $L(\pi_i(T)) = L(\pi(T)) = c$ for a constant c independent of i . Hence $B(T, A) = \sum_{i=1}^{|A|} c = c |A| = O(|A|)$.

Quadratic regime. Suppose instead that distinct audience members produce distinguishable reconstructions, $\pi_i(T) \neq \pi_j(T)$ for $i \neq j$, and that coherent communication across A additionally requires the speaker to supply, for every pair $\{a_i, a_j\}$, an explanation of why their reconstructions differ. The per-member term still contributes $\sum_{i=1}^{|A|} L(\pi_i(T)) = O(|A|)$. The reconciling term contributes one additional explanation for each of the $\binom{|A|}{2} = \frac{|A|(|A|-1)}{2}$ unordered pairs; if each such explanation has complexity bounded by a constant, this term is $\Theta(|A|^2)$, which dominates the linear term as $|A|$ grows. Hence $B(T, A) = O(|A|^2)$. \square

This result gives a quantitative basis for a claim made qualitatively throughout the body of the essay: communication does not merely become harder as audiences grow, it becomes harder at a qualitatively different rate once a single shared interpretive frame can no longer be assumed and divergence between interpretations must itself be explained. The transition from $O(|A|)$ to $O(|A|^2)$ marks the point at which scientific writing's preference for explicit, audience-independent definitions becomes a strict computational advantage rather than a merely stylistic one: a representation invariant enough to keep $\pi_i = \pi$ for (nearly) every member of A collapses the quadratic reconciliation term back to linear.

Definition A.3 (Autobiographical Compression). *An autobiographical projection is a map $\alpha : X \rightarrow B$ from lived history X into a biographical representation B . Its compression ratio is*

$$C(\alpha) = \frac{\dim B}{\dim X}.$$

A conventional memoir has larger $C(\alpha)$ than a scientific paper, but both satisfy $C(\alpha) \ll 1$.

Proposition A.4 (Scientific Writing as Observer Suppression). *Let an observation be represented by a pair $(o, e) \in O \times E$, where o is the observer-context and e is the observed event. Scientific writing applies a projection $\sigma : O \times E \rightarrow E$ given by $\sigma(o, e) = e$. The observed event is preserved while observer-context is erased.*

Proof. For any $o_1, o_2 \in O$, $\sigma(o_1, e) = e = \sigma(o_2, e)$. Thus distinct observer-contexts become indistinguishable after projection. The event component survives, but the autobiographical context does not. \square

B Inference, Metadata, and Attribution

The problem of accidental autobiography can be modeled by distinguishing intended content from inferential content. Let $m = (c, h)$ be a message, where c is the intended communicated content and h is hidden or incidental metadata. An observer does not receive only c , but instead receives some transformation $r = \rho(c, h)$. The privacy risk of the message is then the mutual information

$$\mathcal{R}(m) = I(H; R),$$

where H is the hidden contextual variable and R is the received artifact.

Proposition B.1 (Accidental Disclosure). *If $I(H; R) > 0$, then the artifact R contains information about hidden context H , even when H is not part of the intended message.*

Proof. By definition, mutual information satisfies $I(H; R) = 0$ if and only if H and R are independent. Therefore if $I(H; R) > 0$, observing R changes the probability distribution over H . Hence the received artifact discloses information about the hidden context. \square

For example, a lunar eclipse photograph may intend to communicate an astronomical event c , while also carrying incidental variables h , including timestamp, viewing angle, local landmarks, atmospheric conditions, and architectural features. The photograph becomes risky precisely when $I(\text{location}; \text{image}) > 0$: it then functions not merely as astronomy but as partial geolocation evidence.

Definition B.1 (Temporal Anonymization). *Let H_t denote sensitive contextual information at time t , and let $R_{t+\Delta}$ be a delayed publication. Temporal anonymization occurs when*

$$I(H_t; R_{t+\Delta}) < I(H_t; R_t).$$

The delay reduces the inferential coupling between the artifact and the author's current situation.

Authorship can be formalized in the same way. Let S be a hidden source and let $D = \{d_1, \dots, d_n\}$ be a collection of documents. Attribution is the inference

$$P(S | D) \propto P(D | S) P(S).$$

A legal name, pseudonym, writing style, citation pattern, and conceptual vocabulary all contribute evidence to $P(D | S)$, but no single marker is identical to the source.

Proposition B.2 (Pseudonym Reliability and Spoofing). *Let p be a pseudonym and let S be the true source. If the benefit of spoofing p increases beyond the cost of imitation, then $P(S | p)$ decreases unless additional evidence is introduced.*

Proof. Initially, the pseudonym may be strongly associated with one source, so $P(p | S)$ is high and $P(p | S')$ is low for $S' \neq S$. If spoofing becomes rewarding, other sources increase $P(p | S')$. Since Bayesian attribution depends on the relative likelihoods $P(p | S)$ and $P(p | S')$, an increase in the latter reduces the posterior reliability of p alone. Additional signals are therefore required to preserve attribution. \square

C Erasure, Exposure, and the Trace Principle

Van Arsdall's account of intermedial autobiography turns on a productive tension between erasure and exposure: autobiographical media do not simply reveal the self, nor do they simply conceal it; they produce a self through partial disclosure, displacement, substitution, and mediation [11]. This section gives that distinction a formal counterpart within the framework developed above.

Let S denote the latent space of the self: memories, bodily states, relations, locations, histories, intentions, and unarticulated constraints. An autobiographical artifact A is produced by a projection $\pi : S \rightarrow A$.

Definition C.1 (Exposure and Erasure). *The exposed component of S under π is the mutual information*

$$E_\pi(S) = I(S; \pi(S)).$$

The erased component is the conditional entropy

$$R_\pi(S) = H(S | \pi(S)).$$

Proposition C.1 (Autobiography Occupies an Intermediate Regime). *For any non-trivial autobiographical projection π ,*

$$0 < E_\pi(S) < H(S).$$

Proof. If $E_\pi(S) = 0$, then S and $\pi(S)$ are independent, so the artifact conveys nothing about the self and communication fails by triviality. If $E_\pi(S) = H(S)$, then $R_\pi(S) = H(S) - E_\pi(S) = 0$, so $\pi(S)$ determines S exactly; the entire latent self would be recoverable from the artifact, which is both practically unattainable for any artifact of bounded dimension and, where attainable, indistinguishable from the self it claims merely to represent. Excluding both degenerate cases leaves $0 < E_\pi(S) < H(S)$. \square

This gives a formal reading of intermedial autobiography. Writing, photography, film, pseudonyms, metadata, and digital signatures are not merely different media; they are different projections $\pi_1, \pi_2, \dots, \pi_n$ of the same latent self-space S , each preserving a different invariant and erasing a different remainder. The autobiographical object available to a reader is therefore not a single representation but a bundle of partial projections,

$$A = \{\pi_1(S), \pi_2(S), \dots, \pi_n(S)\},$$

and the self that a reader reconstructs is approximated by the intersection of the preimages compatible with each trace:

$$S_{\text{inferred}} \approx \bigcap_{i=1}^n \pi_i^{-1}(\pi_i(S)).$$

Proposition C.2 (Additional Traces Narrow Inference). *For any projections π_1, \dots, π_n and an additional projection π_{n+1} ,*

$$\bigcap_{i=1}^{n+1} \pi_i^{-1}(\pi_i(S)) \subseteq \bigcap_{i=1}^n \pi_i^{-1}(\pi_i(S)).$$

Proof. Intersecting with an additional set $\pi_{n+1}^{-1}(\pi_{n+1}(S))$ can only remove elements from, never add elements to, the existing intersection. The inferred space of compatible selves is therefore non-increasing in the number of traces considered. \square

This proposition formalizes a point made throughout the body of the essay: authorship, identity, and self-representation are reconstructed not from any single marker but from the cumulative narrowing produced by many independent traces, each individually inconclusive.

Proposition C.3 (The Trace Principle). *As audience size, or more generally the heterogeneity and scale of the interpretive ecosystem through which a representation must travel, increases, the cost of direct autobiographical exposure increases correspondingly. Consequently, self-representation tends to migrate from direct disclosure toward traces, proxies, signatures, pseudonyms, artifacts, and transformation histories.*

Proof. Let A be an audience and let $\mathcal{H}(A)$ denote its heterogeneity. By the channel-capacity model developed elsewhere in this appendix, the effective transmissible bandwidth is

$$K_{\text{eff}}(A) = \frac{K_0}{1 + \mathcal{H}(A)}.$$

Thus if $\mathcal{H}(A_2) > \mathcal{H}(A_1)$, then $K_{\text{eff}}(A_2) < K_{\text{eff}}(A_1)$.

Let $C(\pi(S))$ denote the communicative complexity of an autobiographical projection. Reliable transmission requires

$$C(\pi(S)) \leq K_{\text{eff}}(A).$$

Direct autobiographical exposure corresponds to projections with comparatively high mutual information $E_\pi(S) = I(S; \pi(S))$. Since any representation carrying more information about S must, in general, require at least as much communicative complexity to transmit, $C(\pi(S))$ is monotone non-decreasing in $E_\pi(S)$. Hence increasing exposure increases or preserves the required channel capacity.

As $\mathcal{H}(A)$ increases, $K_{\text{eff}}(A)$ decreases. Therefore any projection π that was admissible for a more homogeneous audience may cease to satisfy $C(\pi(S)) \leq K_{\text{eff}}(A)$ for a larger or more

heterogeneous audience. To restore admissibility, one must choose a projection π' satisfying $C(\pi'(S)) \leq K_{\text{eff}}(A)$, which, by monotonicity, requires reducing or bounding $E_{\pi'}(S)$.

Thus, under increasing audience scale or heterogeneity, direct high-exposure autobiography becomes less transmissible. The communicative artifact must preserve only those lower-dimensional features capable of surviving the reduced effective bandwidth. These preserved features are precisely traces, proxies, signatures, pseudonyms, artifacts, and transformation histories. Self-representation therefore migrates from direct disclosure toward trace-based representation as audience scale and interpretive heterogeneity increase. \square

The Trace Principle ties the present formalism directly to the recurring examples of the essay. A pseudonym is a trace. A cryptographic hash is a trace. A profile image composed from publicly available filters, in a particular and idiosyncratic order, is a trace. A hierarchical filter application order is a trace. A writing style is a trace. Authorship itself, as argued throughout, is never observed directly but reconstructed from the accumulation of such traces. Erasure and exposure are accordingly not opposing tendencies but a single coupled process: every trace that is exposed narrows the space of selves compatible with the available evidence, while everything erased in producing that trace is precisely what keeps a person from being fully reducible to it.

D Invariants, Geometry, and the Drift Toward Formalism

Let G be a group or semigroup of transformations representing reinterpretation, translation, audience shift, temporal distance, political reframing, and contextual loss. A communicative structure $s \in S$ is invariant under G if $g \cdot s = s$ for all $g \in G$. More generally, s is approximately invariant if $d(g \cdot s, s) < \epsilon$ for all relevant $g \in G$.

Definition D.1 (Semantic Robustness). *The semantic robustness of a structure s under transformations G is*

$$R_G(s) = \exp(-\mathbb{E}_{g \sim G} d(g \cdot s, s)).$$

High robustness means the structure remains recognizable under reinterpretation.

Proposition D.1 (Formal Structures Maximize Transportability). *If a formal representation f has lower expected distortion than an autobiographical representation b , namely*

$$\mathbb{E}_{g \sim G} d(g \cdot f, f) < \mathbb{E}_{g \sim G} d(g \cdot b, b),$$

then $R_G(f) > R_G(b)$.

Proof. By definition $R_G(s) = \exp(-\mathbb{E}_{g \sim G} d(g \cdot s, s))$. Since the exponential function is strictly decreasing in its negative argument, lower expected distortion implies greater robustness. Hence $R_G(f) > R_G(b)$. \square

This gives a compact mathematical expression of the essay's central thesis. Scientific writing, mathematical notation, legal definitions, curricula, disclaimers, and pseudonyms are all strategies for increasing $R_G(s)$. They do so by discarding fragile contextual dimensions and preserving structures that survive transformation.

Definition D.2 (Autobiography of Constraints). *Let E be the set of life events and let K be the set of communicative constraints governing disclosure. A conventional autobiography projects $\alpha_E : X \rightarrow E$, while an autobiography of constraints projects $\alpha_K : X \rightarrow K$. The first tells what happened. The second tells what governs what can be said.*

Proposition D.2 (Indirect Autobiographical Recovery). *If the constraint projection α_K is stable across many examples, then a reader may infer properties of the author’s trajectory even without explicit chronology.*

Proof. Suppose the reader observes examples e_1, \dots, e_n and associated constraints k_1, \dots, k_n . If the constraints exhibit regularity, then there exists a latent structure T such that $P(k_1, \dots, k_n | T)$ is high. By Bayesian inference,

$$P(T | k_1, \dots, k_n) \propto P(k_1, \dots, k_n | T) P(T).$$

Thus repeated constraints allow reconstruction of a hidden trajectory T , even when ordinary biographical events are absent. \square

The conclusion follows naturally. Autobiography is not eliminated by abstraction. It is transformed into a different invariant. The life does not appear as a sequence of events, but as a pattern of constraints, projections, and preferred structures. In the notation of this appendix, the autobiographical content has moved from E to K , from event-space to constraint-space.

E Curricula, Disclaimers, and Frame Operators

Let U denote the space of possible utterances and let $I_A : U \rightarrow Z_A$ denote the interpretation map for an audience A . A frame operator is a contextual transformation

$$F : (U, A) \rightarrow (U, A_F),$$

which modifies the interpretive conditions under which an utterance is received. A disclaimer, curriculum, methodological preface, content advisory, or comedic setup may all be modeled as frame operators. Their purpose is to reduce the interpretive variance $\text{Var}_A(u)$ by restricting the admissible interpretive region available to the audience.

Definition E.1 (Admissible Interpretive Region). *For an utterance u , let $\mathcal{I}_A(u) = \{I_i(u) : a_i \in A\}$ be the set of interpretations produced by audience members. A frame operator F induces a constrained interpretive region $\mathcal{I}_{A,F}(u) \subseteq \mathcal{I}_A(u)$ when the presence of F rules out, discourages, or weakens certain reconstructions of u .*

Proposition E.1 (Disclaimers Reduce Expected Misframing). *Let $M \subseteq Z_A$ be the set of interpretations classified as misframings. Suppose a disclaimer F reduces the probability that an audience member enters M :*

$$P(I_i(u) \in M | F) < P(I_i(u) \in M).$$

Then the expected number of misframings in an audience of size n is reduced.

Proof. Let X_i be the indicator variable for the event $I_i(u) \in M$. Without the disclaimer, $\mathbb{E}[\sum_{i=1}^n X_i] = \sum_{i=1}^n P(I_i(u) \in M)$. With the disclaimer, $\mathbb{E}[\sum_{i=1}^n X_i \mid F] = \sum_{i=1}^n P(I_i(u) \in M \mid F)$. By hypothesis every term in the second sum is smaller than the corresponding term in the first. Therefore the expected number of misframings decreases. \square

This formalizes the role of warnings attached to television programs, jokes, software systems, and scientific documents. A content advisory preceding a violent drama does not alter the drama itself; it alters the probability distribution over interpretive states. Similarly, a statement that a generative system may hallucinate does not change the output; it changes the interpretive stance the user is expected to adopt toward the output.

Curricula function similarly but over a longer interval. Let $C = (F_1, F_2, \dots, F_m)$ be a sequence of frame operators applied over time, seeking convergence toward a target interpretation $z^*: I_A^{(m)}(u) \rightarrow z^*$.

Proposition E.2 (Curricular Convergence). *Suppose each curricular step F_j reduces the expected semantic distance from the target by a factor $0 < \lambda_j < 1$:*

$$\mathbb{E}[d(I_A^{(j)}(u), z^*)] \leq \lambda_j \mathbb{E}[d(I_A^{(j-1)}(u), z^*)].$$

Then after m steps,

$$\mathbb{E}[d(I_A^{(m)}(u), z^*)] \leq \left(\prod_{j=1}^m \lambda_j \right) \mathbb{E}[d(I_A^{(0)}(u), z^*)].$$

Proof. Apply the inequality recursively:

$$\mathbb{E}[d(I_A^{(m)}, z^*)] \leq \lambda_m \mathbb{E}[d(I_A^{(m-1)}, z^*)] \leq \lambda_m \lambda_{m-1} \mathbb{E}[d(I_A^{(m-2)}, z^*)] \leq \dots$$

Continuing yields the stated product bound. \square

Education, on this account, is not merely the presentation of content. It is the repeated application of interpretive constraints designed to make reconstruction converge.

F Narrative Engineering as Equivalence-Class Construction

Let P denote the set of actual persons involved in a life history, and let C denote the set of characters appearing in a public narrative. Fictionalization, anonymization, and compositing may be modeled as a map $\eta : P \rightarrow C$. When several persons are merged into one character, η is many-to-one. When identifying details are changed, η preserves certain relational properties while discarding individuating features.

Definition F.1 (Narrative Equivalence). *Two persons $p_1, p_2 \in P$ are narratively equivalent relative to a theme τ when $R_\tau(p_1) = R_\tau(p_2)$, where R_τ extracts the features relevant to the narrative theme. The equivalence class of p is $[p]_\tau = \{q \in P : R_\tau(q) = R_\tau(p)\}$. A composite character represents an equivalence class rather than a single person.*

This clarifies why fictionalization is not simply falsification. A composite character may fail to preserve individual identity while preserving relational structure. The relevant question is not whether the character corresponds one-to-one to a historical person, but which invariants of the historical relation have been preserved.

Definition F.2 (Ethical Disclosure Cost). *Let N be a narrative and let $p \in P$ be a real person represented by N . The disclosure cost to p is*

$$D_p(N) = \alpha I(p; N) + \beta H_p(N) + \gamma Q_p(N),$$

where $I(p; N)$ measures identifiability, $H_p(N)$ measures potential harm, and $Q_p(N)$ measures power asymmetry between narrator and represented person.

The ethical problem of autobiography is then not simply to maximize accuracy. It is to maximize narrative value subject to disclosure constraints:

$$N^* = \arg \max_N V(N) \quad \text{subject to} \quad D_p(N) \leq \varepsilon_p \quad \text{for every represented person } p.$$

Proposition F.1 (Compositing Can Reduce Identifiability). *Suppose a narrative character c is formed by merging k persons whose relevant features are distributed so that no single person contributes a uniquely identifying bundle. Then, under a uniform prior over candidates, the posterior probability of identifying any one source person from c is bounded above by $1/k$.*

Proof. If the composite character is compatible with each of k persons and no individuating feature uniquely selects one of them, then under a uniform prior the posterior mass is distributed across at least k candidates. Therefore the probability assigned to any one candidate is at most $1/k$. \square

This proposition captures the logic behind changing names, merging people, altering chronology, and suppressing details. These operations reduce identifiability while attempting to preserve the structure of the experience. The resulting narrative is less literal but may be more ethically transmissible.

Proposition F.2 (Power Asymmetry Increases Disclosure Constraint). *If $Q_p(N)$ increases while $I(p; N)$ and $H_p(N)$ remain fixed, then $D_p(N)$ increases.*

Proof. Since $D_p(N) = \alpha I(p; N) + \beta H_p(N) + \gamma Q_p(N)$ and $\gamma > 0$, increasing $Q_p(N)$ strictly increases $D_p(N)$ when all other variables remain fixed. \square

This formalizes why teachers, employers, clinicians, researchers, and public writers face stronger constraints when representing students, employees, patients, participants, or private individuals. The same anecdote may have different ethical costs depending on the relative positions of narrator and subject.

G The Narrative Chain and the Erosion of Information

The narrative self, introduced in the main text as an interface analogous to a name tag or uniform, and the public persona discussed in the sections on public communication, admit a direct joint formalization as a chain of successive projections.

Definition G.1 (The Narrative Chain). *Let S_{true} denote the full latent state of a person – the object denoted S elsewhere in this appendix. Define three further representations connected by projections*

$$S_{\text{true}} \xrightarrow{\pi_1} S_{\text{narrative}} \xrightarrow{\pi_2} S_{\text{public}} \xrightarrow{\pi_3} S_{\text{symbolic}},$$

where π_1 compresses the full self into the narrative self used in ordinary social communication, π_2 further compresses the narrative self into the public persona suitable for transmission to a heterogeneous audience, and π_3 compresses the public persona into the small number of recognizable, repeatable symbols by which a public figure is generally known. The fully composed map is

$$S_{\text{symbolic}} = (\pi_3 \circ \pi_2 \circ \pi_1)(S_{\text{true}}).$$

Proposition G.1 (Monotone Information Loss Along the Narrative Chain). *If each π_k is a deterministic, non-injective function of its argument, then*

$$H(S_{\text{symbolic}}) \leq H(S_{\text{public}}) \leq H(S_{\text{narrative}}) \leq H(S_{\text{true}}),$$

with strict inequality at each step whenever the corresponding π_k is properly non-injective on the support of its argument.

Proof. For any deterministic function f applied to a random variable Z , the data-processing inequality for entropy gives $H(f(Z)) \leq H(Z)$, with equality if and only if f is injective on the support of Z . Applying this to π_1 (with $Z = S_{\text{true}}$), to π_2 (with $Z = S_{\text{narrative}}$), and to π_3 (with $Z = S_{\text{public}}$) yields the stated chain of inequalities. Each projection in this essay’s argument is non-injective by construction – a narrative self never recovers the whole person, a public persona never recovers the whole narrative self, and a symbol never recovers the whole persona – so each inequality is strict. \square

This proposition gives a formal derivation of a phenomenon described qualitatively throughout the body of the essay: a public figure does not become a caricature by accident, nor solely through media distortion. Caricature is the predictable consequence of composing several individually reasonable compressions whose effects on the information content of the resulting representation are monotone and cumulative. No single step in the chain is responsible for the result. The chain itself is.

H Public Persona and Channel Capacity

Let M denote the full conceptual state of a speaker and let Y denote the message received by a public audience. Communication through a public channel is constrained by a capacity K , so that the reliably transmitted information satisfies $I(M; Y) \leq K$. As the heterogeneity of the audience increases, the effective capacity for technical nuance decreases.

Definition H.1 (Audience Heterogeneity). Let $A = \{a_1, \dots, a_n\}$ be an audience, and let B_i represent the background knowledge of a_i . Define audience heterogeneity by

$$\mathcal{H}(A) = \frac{1}{n} \sum_{i=1}^n d_B(B_i, \bar{B})^2,$$

where \bar{B} is the Fréchet mean of audience backgrounds.

We may model effective public bandwidth as

$$K_{\text{eff}} = \frac{K_0}{1 + \mathcal{H}(A)}.$$

This is not intended as an empirical law but as a formal expression of the intuition that increasing heterogeneity lowers the amount of specialized information that can be safely transmitted without severe distortion.

Proposition H.1 (Compression Pressure Under Audience Growth). *If K_{eff} decreases as $\mathcal{H}(A)$ increases, then public communication must either reduce message complexity or tolerate greater interpretive error.*

Proof. Let message complexity be $C(M)$. Reliable transmission requires $C(M) \leq K_{\text{eff}}$. If K_{eff} decreases while $C(M)$ remains fixed, this inequality eventually fails. Therefore the communicator must reduce $C(M)$ or accept unreliable transmission. \square

This formalizes why public intellectuals, popular scientists, teachers, and entertainers often develop simplified personas. The persona is a compressed interface $\psi : M \rightarrow P$ from the high-dimensional conceptual state M to a lower-dimensional public representation P . Its function is to preserve recognizability while reducing complexity.

Definition H.2 (Persona Loss). *The loss induced by a public persona ψ is $L_\psi(m) = \log |\psi^{-1}(\psi(m))|$. Large L_ψ means that many distinct conceptual states are collapsed into the same public representation.*

Proposition H.2 (Public Recognizability Requires Dimensional Reduction). *If a public representation must be recognized by an audience with effective capacity K_{eff} , then any persona $\psi : M \rightarrow P$ satisfying reliable recognition must obey $\dim P \leq K_{\text{eff}}$ in the relevant representational units.*

Proof. Reliable recognition requires that the audience can reconstruct the public representation P from the transmitted message. If $\dim P > K_{\text{eff}}$, the channel cannot reliably transmit enough information to distinguish the representation. Therefore $\dim P \leq K_{\text{eff}}$. \square

Thus the public persona is not merely an artificial mask. It is a bandwidth-constrained projection of a larger conceptual system. The caricature emerges when repeated public transmission selects only the most stable low-dimensional features and discards the rest.

I Admissibility and the Geometry of Self-Projection

The preceding appendix sections have introduced exposure $E_\pi(S)$, semantic robustness $R_G(s)$, and disclosure cost $D_p(N)$ as separate quantities, each capturing one constraint operating on a single autobiographical or scientific artifact. They are, however, stages of a single argument rather than independent results: the Projection Complexity Theorem (App. A) establishes that communication has a cost which grows nonlinearly once a shared interpretive frame fails; the Narrative Chain (above) establishes that this cost forces a provably monotone compression of the self; and the present section characterizes the best compression available once both facts are taken as given. This final section assembles exposure, robustness, and disclosure cost into the single constrained optimization problem that organizes the whole of this essay.

Definition I.1 (Admissible Self-Projection). *Let S be the latent self, G a group of reinterpretation transformations, and $P = \{p_1, \dots, p_k\}$ the set of other persons entangled in a candidate representation, with disclosure cost $D_p(\cdot)$ defined as for narratives above and applied here to an arbitrary projection of S . A projection $\pi : S \rightarrow A$ is admissible at thresholds (τ, ε) if*

$$R_G(\pi(S)) > \tau \quad \text{and} \quad \max_{p \in P} D_p(\pi(S)) < \varepsilon.$$

The admissibility manifold $\mathcal{M}(\tau, \varepsilon)$ is the set of all such admissible projections.

Definition I.2 (The Self-Projection Problem). *Given thresholds (τ, ε) , the self-projection problem is*

$$\pi^* = \arg \max_{\pi \in \mathcal{M}(\tau, \varepsilon)} E_\pi(S).$$

Proposition I.1 (Non-Existence of an Unconstrained Maximizer). *The identity projection $\pi = \text{id}_S$ maximizes $E_\pi(S)$ over all projections without constraint, but for any non-trivial group G of reinterpretation transformations acting non-trivially on S and any τ sufficiently close to 1, $\text{id}_S \notin \mathcal{M}(\tau, \varepsilon)$.*

Proof. Since $\pi(S) = S$ under the identity, $E_{\text{id}}(S) = I(S; S) = H(S)$, the maximum value attainable by mutual information between S and any function of S ; no other projection can exceed it. However, $R_G(\text{id}_S) = \exp(-\mathbb{E}_{g \sim G} d(g \cdot S, S))$, and because G acts non-trivially on S , $d(g \cdot S, S) > 0$ for a non-negligible measure of $g \in G$, so $R_G(\text{id}_S) < 1$. For τ sufficiently close to 1 this falls below threshold, and $\text{id}_S \notin \mathcal{M}(\tau, \varepsilon)$ regardless of ε . \square

Proposition I.2 (Scientific Writing as a Boundary Solution). *As $\tau \rightarrow \tau_{\max}$, the supremum of robustness attainable by any non-degenerate projection of S , the solution $\pi^*(\tau, \varepsilon)$ to the self-projection problem converges to the observer-suppressing projection $\sigma(o, e) = e$ established earlier in this appendix as characteristic of scientific writing.*

Proof. Let τ_{\max} denote the supremum of $R_G(\pi(S))$ over all non-degenerate projections π . By definition,

$$R_G(\pi(S)) = \exp(-\mathbb{E}_{g \sim G} d(g \cdot \pi(S), \pi(S))).$$

Maximizing R_G is therefore equivalent to minimizing the expected distortion $\mathbb{E}_{g \sim G} d(g \cdot \pi(S), \pi(S))$.

By the earlier proposition that formal structures maximize transportability, projections that discard observer-dependent and context-dependent information have lower expected distortion under G than projections that preserve autobiographical detail. In particular, an observation represented as $(o, e) \in O \times E$, with observer-context o and event e , is made more robust by the scientific projection $\sigma : O \times E \rightarrow E$, $\sigma(o, e) = e$. For any two observer contexts $o_1, o_2 \in O$,

$$\sigma(o_1, e) = e = \sigma(o_2, e),$$

so variation in observer-context is eliminated entirely. Any transformation $g \in G$ acting primarily on observer-context, rhetorical framing, biographical situation, or social interpretation therefore produces no change in $\sigma(o, e)$. Hence σ removes precisely those dimensions responsible for much of the expected distortion.

Now consider the admissible set

$$\mathcal{M}(\tau, \varepsilon) = \left\{ \pi : R_G(\pi(S)) > \tau, \max_{p \in P} D_p(\pi(S)) < \varepsilon \right\}.$$

As $\tau \rightarrow \tau_{\max}$, this feasible set contracts toward projections whose robustness is arbitrarily close to maximal. Since σ is the projection that removes observer-context while preserving the event, and since observer-context is the principal source of autobiographical variance in scientific communication, any sequence of maximizers $\pi^*(\tau, \varepsilon)$ with $\tau \rightarrow \tau_{\max}$ must converge, up to equivalence of event-preserving projections, toward σ .

Finally, because σ preserves only e and discards o , the exposure term $E_{\pi^*}(S)$ converges toward the residual exposure compatible with the observed event alone. This is exactly the scientific limit: the bare observation remains, while the observer is suppressed. Scientific writing is therefore the boundary solution of the self-projection problem as robustness under reinterpretation is driven toward its maximum. \square

This completes the formal counterpart to the essay's central claim. Memoir, pseudonymous essay, scientific paper, and symbolic public persona are not different kinds of object but different solutions $\pi^*(\tau, \varepsilon)$ to the same constrained optimization, evaluated at different thresholds of robustness and disclosure cost. A single relation holds across the entire manifold $\mathcal{M}(\tau, \varepsilon)$:

$$\text{Autobiography} = \text{Admissible Self-Projection.}$$

The self is never transmitted directly. Only admissible projections of the self are transmitted, and scientific writing is not a special case standing outside this account but its most extreme and most robust corner – the point on the manifold at which surviving reinterpretation has been prioritized almost completely over exposing the self that did the discovering.

References

- [1] Barthes, R. (1977). *Image, Music, Text*. Hill and Wang.

- [2] Barthes, R. (1981). *Camera Lucida: Reflections on Photography*. Hill and Wang.
- [3] Bateson, G. (1972). *Steps to an Ecology of Mind*. Chandler Publishing.
- [4] Eco, U. (1990). *The Limits of Interpretation*. Indiana University Press.
- [5] Foucault, M. (1977). What Is an Author? In *Language, Counter-Memory, Practice*. Cornell University Press.
- [6] Goffman, E. (1959). *The Presentation of Self in Everyday Life*. Doubleday.
- [7] Latour, B. (1987). *Science in Action: How to Follow Scientists and Engineers through Society*. Harvard University Press.
- [8] Ricoeur, P. (1992). *Oneself as Another*. University of Chicago Press.
- [9] Scott, J. C. (1998). *Seeing Like a State: How Certain Schemes to Improve the Human Condition Have Failed*. Yale University Press.
- [10] Shannon, C. E. (1948). A Mathematical Theory of Communication. *Bell System Technical Journal*, 27(3), 379–423.
- [11] Van Arsdall, L. E. (2015). *Between Erasure and Exposure: Intermedial Autobiography Since Roland Barthes*. Ph.D. dissertation, University of California, Los Angeles. eScholarship: <https://escholarship.org/uc/item/0v69m960>.