

The Ladder and the Artifact: Publication as Projection

Flyxion
Independent Researcher

July 2026

Abstract

A finished paper is ordinarily encountered as though it had always existed in its final form: a fixed sequence of definitions, propositions, and proofs, presented all at once. The process that produced it, a long and typically disordered sequence of drafts, is not part of what the reader receives. This essay argues that the relationship between a draft history and its published artifact is a specific instance of the operator taxonomy developed in two companion analyses of continuation and recoverability, and that naming the relationship precisely resolves an ambiguity in the usual metaphors for it. Publication is not well modeled as simple narrowing, since a finished paper is not merely a subset of everything that was drafted, nor as an override, since the final text is causally shaped by the drafts rather than independent of them. It is modeled here as a coordinate projection Π taking a drafting history to its terminal draft, which induces an equivalence relation identifying any two histories arriving at the same final text; a property of a drafting history is shown to survive publication exactly when it factors through Π , giving an exact criterion for what an artifact preserves in place of the earlier informal language of claims surviving while process is lost. Version control systems are examined as a considerably less lossy member of the same family of projections, one that preserves discarded material externally rather than eliminating it, while remaining subject to genuinely destructive operations of its own, squashing, reference reassignment, and garbage collection, that are shown not to reduce to the override case even though they erase as thoroughly. The essay closes by arguing that this erasure is not a defect of publication but constitutive of what publication is for, that the quotienting performed by Π is precisely the mechanism by which a historical construction becomes usable as a state description, and that the tradeoff this exacts, reusability purchased at the price of an unrecoverable process, is exact enough to state rather than merely to gesture at.

1 The ladder and the artifact

There is an old image, familiar from the closing remarks of Wittgenstein's early philosophy, of a ladder that must be climbed and then discarded once the climber has reached the vantage point the ladder existed to provide. The image is usually invoked to describe a certain kind of philosophical text, one whose propositions are tools for arriving somewhere rather than claims meant to be retained once the destination is reached. The same image applies, with less drama but more literalness, to almost anything written with revision. A finished paper is the vantage point. The drafts are the ladder. And the reader, arriving at the vantage point directly, is not given the ladder at all.

This essay is about what exactly gets left behind when the ladder is discarded, and about the fact that this can now be stated with some precision rather than left as a metaphor. Two companion analyses developed a vocabulary for exactly this kind of question, though neither was written with drafting in mind. One classified the ways a sequence of operators acting on a space of possibilities can depend on the order in which the operators are applied, distinguishing static narrowing, state-dependent narrowing, override, and persistence. The other asked a finer question of the same objects, not merely whether history matters but what specifically survives it, and showed that several of these mechanisms are best understood not by what they discard but by what they leave recoverable. This essay applies that vocabulary directly to the relationship between a sequence of drafts and the paper that results from them.

2 Drafting operations, classified

A sequence of drafts $D_0, D_1, \dots, D_n$, where D_0 is a blank page or an early outline and D_n is the version submitted or published, is a chain in exactly the sense the companion analyses treat formally, except that the space being acted on is not a set of admissible continuations but something closer to the space of all papers the author might yet write. Each transition from D_i to D_{i+1} is an editorial operation, and the operations that occur in ordinary drafting fall cleanly into the categories already established.

An instruction such as *remove all references to the abandoned example* behaves as static narrowing: it rules out a fixed class of content, identified independently of whatever else the draft currently contains, and every subsequent draft is simply the prior one with that class excluded. An instruction such as *make this proof clearer* or *rewrite the introduction so it matches the conclusion* is state-dependent narrowing in the strict sense of the companion analyses: it cannot even be carried out without reference to the current state of the proof or the current state of the conclusion, since what counts as clearer, or as matching, is defined relative to whatever is already on the page. An instruction such as *keep the mathematics, cut the exposition* is a projection in the technical sense developed in the second analysis, retaining recoverable content along one dimension of the draft while discarding it along another. An instruction such as *throw out the whole section and start over* is an override, in the strongest sense: whatever the discarded section actually said becomes entirely unrecoverable from the draft that follows it, not merely absent but undetectable as having ever been attempted. And a standing constraint such as *this paper must remain self-contained*, or, in the drafting of the companion essays themselves, *avoid cultural examples the reader must have independently*

encountered, behaves as persistence: not a single edit but a condition re-applied at every subsequent stage, regardless of what other editing occurs alongside it.

Observation 1. *The taxonomy developed for continuation chains was not designed with drafting in mind, and yet every ordinary editorial operation falls into one of its categories without residue. This is not because drafting was secretly always about continuation operators, but because the taxonomy was built from the general question of what a sequence of operations acting on a space of possibilities can do to that space, and drafting is one instance of a much more general phenomenon.*

3 Publication as projection

The relationship between the full sequence $D_0, \dots, D_n$ and the single artifact a reader receives can now be stated precisely rather than left as the informal observation that finished work hides its process. Publication is not itself one of the elementary operators; it is better understood as the composition of everything that happened across the whole drafting chain, evaluated for what it leaves recoverable in the final text alone.

The distinction between source and artifact developed elsewhere appears here in a particularly concrete form. The drafting chain is the source object: a history of transformations, revisions, refusals, recoveries, and replacements. The published paper is a rendering of that source under a projection that preserves only a restricted set of coordinates. The reader ordinarily encounters the rendering and treats it as primary, not because the source ceased to exist, but because the projection has become the socially authoritative representation of it.

Definition 1. *Let $D_0, \dots, D_n$ be a chain of drafts as above. The published artifact is D_n together with nothing else: no record of $D_0, \dots, D_{n-1}$, no record of which operators were applied at which stage, and no record of the order in which they occurred. Publication is the map taking the full chain to this artifact alone.*

This definition deliberately abstracts a single informational function from a much larger institution. Publication, as a social and scholarly practice, involves peer review threads, editorial correspondence, revision histories some venues do publish alongside a final paper, and supplementary material that may itself preserve part of the discarded process. None of that is being denied or argued away here. The claim under examination is narrower: whatever else a given venue's practices happen to preserve, the artifact a typical reader engages with, cites, and builds on is D_n considered in isolation, and it is that engagement, not the full institutional record, whose informational structure is being analyzed. A reader applying a cited theorem is almost never also consulting that paper's referee reports or rejected earlier drafts, even when those happen to survive somewhere; the operator $(D_0, \dots, D_n) \mapsto D_n$ describes what actually functions as the artifact in ordinary use, whatever a complete institutional archive might separately retain.

Remark 1. *Publication, as just defined, is a coordinate projection in the ordinary mathematical sense: writing $\pi_n(D_0, \dots, D_n) = D_n$, publication simply is π_n . Because π_n discards every coordinate but the last, it is highly many-to-one: two drafting histories $H_1 = (D_0, \dots, D_n)$ and $H_2 = (D'_0, \dots, D_n)$ that differ arbitrarily in every draft before the last, but happen to arrive at the same final text, satisfy $\pi_n(H_1) = \pi_n(H_2)$ without satisfying $H_1 = H_2$. Publication therefore*

induces an equivalence relation on drafting histories, $H_1 \sim H_2$ if and only if $\pi_n(H_1) = \pi_n(H_2)$. What a reader receives is not a representative history from the resulting equivalence class, but the common image shared by every history in it. The artifact identifies the fiber $\pi_n^{-1}(D_n)$ without indicating which member of that fiber was the actual history, how its stages were ordered, or how large the fiber is. Two authors who reached the identical final proof by entirely different routes, one who saw the argument immediately and one who spent months on false starts, are publication-equivalent in exactly this sense: indistinguishable after π_n , however distinguishable they were before it.

Remark 2. Strictly, π_n is a family of maps rather than a single operator, since its domain, tuples of length $n + 1$, changes with n , and a drafting history of five stages and one of five hundred stages are formally elements of different spaces. This is easily repaired: let $\mathcal{H} = \bigcup_{n \geq 0} \mathcal{D}^{n+1}$ be the set of all finite drafting histories of any length, and let $\Pi : \mathcal{H} \rightarrow \mathcal{D}$ be defined by $\Pi(D_0, \dots, D_n) = D_n$ for every history, regardless of its length. Publication is properly this single operator Π , and π_n is its restriction to histories of a fixed length. The essay continues to write π_n where a particular history's length is fixed by context, since doing so keeps the worked examples concrete, but the underlying operator being analyzed throughout is Π , defined uniformly across drafting histories of every length.

Ordinary publication, defined this way, is one point in a larger family of possible projections rather than the only one available. Full archival publication, $(D_0, \dots, D_n) \mapsto (D_0, \dots, D_n)$, is the identity map and therefore the least lossy member of the family, retaining every history perfectly and collapsing nothing. Publishing a revision window, $(D_0, \dots, D_n) \mapsto (D_k, \dots, D_n)$ for some k between 0 and n , retains a suffix of the drafting history and sits strictly between the two extremes, more informative than ordinary publication and less than full archiving. Ordinary publication, $\Pi(D_0, \dots, D_n) = D_n$, is not the most lossy projection available in any absolute sense; a projection to the paper's title alone, or a constant map recording only that some paper exists, would each discard more. It is instead the most lossy projection among those that still present the completed manuscript itself as the artifact of record: it retains the smallest amount of the history compatible with the artifact still being D_n , while remaining, unlike an override, genuinely a function of the whole of it. Characterizing ordinary publication this way makes clear that the scholarly practice actually in use has selected, whether deliberately or not, close to the maximally lossy end of a spectrum that had less lossy options available at every point along it, even holding fixed the requirement that the artifact of record be the manuscript itself rather than some further summary of it.

Under this definition, publication is a projection in exactly the technical sense of the companion analysis on recoverability, applied not to D_n in isolation but to the whole chain $(D_0, \dots, D_n)$: the companion analysis defines an element of the domain as retaining recoverable membership under an operator if it can be reconstructed from the operator's output together with a description of the operator, and here the operator in question is $\pi_n(D_0, \dots, D_n) = D_n$ itself. Certain functions of the chain remain exactly recoverable from the published artifact: the claims made in D_n are exactly the claims a reader encounters, and the proofs given are exactly the proofs a reader can check, since these are simply properties of D_n itself and D_n is precisely what survives. But a much larger set of properties of the chain are not recoverable at all. The order in which the central results were discovered relative to one another is not recoverable from D_n , even in principle, unless the paper happens to narrate its own history, which most do not. The definitions that were tried and

abandoned before the ones that appear in D_n were settled on are not recoverable at all; a reader cannot distinguish a paper whose final definition was the author's first attempt from one where it was the eleventh, since D_n looks identical either way. The local uncertainty that attended each step, which claims felt solid and which felt shaky at the time they were written, is not recoverable, since D_n presents every surviving claim with the same uniform declarative confidence regardless of how it was arrived at.

Observation 2. *A property f of a drafting history survives publication exactly when it factors through the publication map: there exists some function g on published artifacts such that $f = g \circ \pi_n$. The final claims, proof text, and argument structure survive in this sense because each can be determined from D_n alone. The order of discovery, the number of abandoned definitions, and the uncertainty attending earlier stages do not survive, because histories with the same image under π_n may differ arbitrarily with respect to each of them. This is exactly the earlier equivalence relation restated as a criterion: a property survives publication if and only if it is constant on publication-equivalence classes, since $f = g \circ \pi_n$ holds precisely when $\pi_n(H_1) = \pi_n(H_2)$ forces $f(H_1) = f(H_2)$.*

Observation 3. *The same distinction can be expressed in the language of procedural ontology. What survives publication is principally state information: the final claims, definitions, proofs, and argument structure. What is discarded is procedural information: the sequence of operations by which those states were reached. The artifact therefore presents ontology in its rendered form while suppressing ontology in its generative form. The reader encounters what exists at the end of the process rather than the process through which it came to exist. This is also the point at which the drafting chain's own state-dependent narrowing operations, each defined only relative to whatever the manuscript currently was, cease to be visible as operations at all: a state-dependent instruction like make this proof clearer leaves no trace in D_n of having been relative to anything, since D_n presents the clarified proof as though it had simply been written that way. Publication does not merely discard the earlier drafts; it converts what were once state-dependent transformations into an object that reads as though no transformation had been necessary to produce it.*

Remark 3. *This is a stronger claim than saying publication merely compresses the drafting process. Compression, in ordinary usage, suggests that the same information is present in smaller form and could in principle be decompressed. Projection, in the sense established here, makes no such promise. The region of the drafting history that publication discards is not present in compressed form inside D_n ; it is simply absent, in exactly the way a static narrowing operator with region A leaves no trace, recoverable or otherwise, of what happened outside A .*

4 What version control preserves, and what it does not

A version-controlled repository is the most obvious candidate for a system that resists this erasure, since a commit history preserves $D_0, \dots, D_{n-1}$ as accessible objects rather than discarding them the moment D_n is finalized. It is worth being precise about how much this actually changes.

Where an ordinary published artifact is the result of a projection retaining only D_n , a repository is closer to a projection retaining a much larger region, potentially the entire chain, together with metadata about which operator connected each consecutive pair of drafts. This is a real and substantial difference: a reader of the repository can, at least

in principle, recover the order of discovery, inspect abandoned definitions still sitting in deleted branches, and observe the local uncertainty visible in commit messages written in the moment rather than reconstructed after the fact. In the vocabulary of the companion analyses, the two differ in how much of the chain remains recoverable, not in the fundamental character of the operator each performs: ordinary publication is an extremely lossy projection, retaining almost no recoverable information about the chain beyond its terminal draft while remaining, unlike an override, causally dependent on the whole chain that produced it, whereas a repository is a far less lossy projection with a retained region that happens to include almost everything. Calling ordinary publication override-like would blur exactly the distinction this essay depends on: an override's output is independent of its input, but D_n is never independent of $D_0, \dots, D_{n-1}$, only unrecoverable from them. Publication forgets almost every coordinate of the chain. It does not, in the technical sense, discard its dependence on them.

But the exception is only partial, for two reasons. First, a repository is not immune to operations that erase the history it otherwise preserves. A squashed merge projects many commits into a single composite state; a force-push can replace the branch history treated as authoritative; and garbage collection can permanently destroy objects that have ceased to be reachable from any retained reference. These operations differ formally, and not all are overrides in the strict sense established by the companion analyses: a squash remains causally dependent on the commits it condenses, and garbage collection is conditioned by the repository's current reference structure, so neither is a constant map independent of its input in the way an override is defined to be. What they share is not input-independence but the destruction of distinctions among prior histories. The repository resists erasure only insofar as its references, objects, and metadata continue to preserve the information required for reconstruction, and there is nothing about version control as a technology that guarantees that they will. Preservation is a discipline maintained through repository practice, in the same sense that avoiding cultural examples was a persistent constraint maintained across the drafting of the companion essays rather than an automatic property of the drafting process itself. Second, even an untouched, fully preserved repository does not preserve everything: the reasoning behind a discarded definition, as opposed to the definition itself, is recoverable only if it was written down somewhere, and most of it, in most projects, was not. A repository preserves more of the ladder than a published paper does. It does not preserve the ladder entire.

Remark 4. *Reverting a commit, restoring a deleted branch, or cherry-picking an abandoned definition back into active use are all, in the vocabulary developed here, instances of an operator reintroducing content that a prior narrowing step had excluded. This is exactly the kind of recovery the foreclosure results in the companion analyses show to be impossible for a chain of pure narrowing operators acting on a single evolving space. What makes it possible in a repository is that the excluded material was never actually removed from the larger space the repository maintains; it was only removed from the branch currently regarded as authoritative. Foreclosure, in other words, is a property of chains that discard what they exclude, not a property of chains that merely stop pointing at it.*

This distinction parallels the difference between deletion and repair developed elsewhere. A deleted possibility that remains recoverable through an external trace has not

been annihilated; it has merely ceased to occupy a privileged position within the active structure. True foreclosure occurs only when the information required for reconstruction has itself been destroyed. The repository functions as a repair substrate precisely because it preserves traces that ordinary publication discards.

5 The erasure is not a defect

It would be easy to read the argument so far as a complaint: that publication ought to preserve more of its own history than it does, and that the loss described above is a failure of scholarly practice correctable by better habits or better tooling. That is not the claim being made here. The erasure that publication performs is close to constitutive of what publication is for, not an unfortunate side effect of it.

A drafting history contains many still-admissible futures. Publication presents one realized trajectory as an object. The alternatives cease to function as part of the artifact, even when they once participated in its construction.

A state description, in the vocabulary of the first companion analysis, is valuable precisely because it can be used without requiring whoever uses it to have witnessed the history that produced it. A reader who wants to apply a theorem does not need to know which of the author's earlier, abandoned formulations of it were wrong and why; the reader needs the theorem and its proof, which is to say the reader needs exactly the region that publication retains. If every published paper carried its full drafting history as a condition of being read, the artifact would cease to function as a state description at all, and every act of using a result would require re-traversing a historical construction that the whole point of publication was to make unnecessary to re-traverse. The projection is not incidental to this function; it is the mechanism by which the function is achieved. A ladder that could not be discarded would not have been much of a ladder.

This is the precise sense in which publication converts a historical construction into a state description. The first companion analysis distinguishes the two directly: a state description is recoverable from the unordered set of constraints in force, while a historical construction requires the specific sequence by which those constraints came to bear. A drafting history is a historical construction in exactly that sense, its meaning inseparable from the order in which decisions, retractions, and discoveries occurred. The equivalence relation $\sim$ is what performs the conversion: by identifying every history that arrives at the same D_n , publication discards precisely the sequencing information that made the drafting process a historical construction rather than a state description, and retains precisely the claims and structure that a state description consists of. Quotienting by $\sim$ is not merely analogous to the conversion from history to state; it is that conversion, stated as an equivalence relation on the object being converted.

Seen from this perspective, publication is one instance of a more general phenomenon that appears repeatedly across scientific, computational, and social systems: histories generate structures, but structures are easier to store, transmit, compare, and reuse than histories. The resulting pressure toward state descriptions is not unique to scholarship. It is a recurring consequence of any system that must trade reconstructive fidelity for portability. The artifact survives because it is cheaper to circulate than the process that produced it.

This asymmetry mirrors a recurring pattern in distinguishability-oriented work. Histories are often richer than the states they generate, but states are typically easier to compare, classify, and transmit. The projection from history to state therefore increases portability at the cost of distinguishability, replacing many historically distinct trajectories with a single artifact that no longer records which trajectory produced it.

What can be said precisely, rather than merely asserted, is the shape of the tradeoff. Reusability requires exactly the kind of lossy projection described above, retaining claims and structure while discarding process, and the cost of that reusability is that the process becomes genuinely unrecoverable, not temporarily hidden, from the artifact alone. This is not a claim that could have been stated this precisely without the vocabulary the companion analyses provide, and it is the reason this essay exists as a third piece rather than as a closing remark appended to either of the first two: the question it asks, why intellectual artifacts preferentially preserve some histories while erasing others, is not the same question as when a chain is history-dependent, and it is not the same question as what a given operator preserves. It is the question of why the artifact that functions as the field's working unit, whatever else a given venue's fuller institutional record separately retains, is built around one particular and highly lossy operator rather than another, and the answer offered here is that the artifact's usefulness as something citable and reusable depends on exactly the erasure that makes its history unrecoverable from the artifact alone.