CIVILIZATION'S UNDO STACK

On Markov Boundaries, Coordination Costs, and the Recurring Regression of Progress

Flyxion

June 2026

Abstract

Across apparently unrelated domains—interface design, computational architecture, archival systems, infrastructure, governance, and artificial intelligence—a common structural failure recurs with remarkable consistency: mechanisms developed specifically to absorb coordination costs are dismantled, and the costs are re-externalized onto individuals. This essay argues that the underlying concept unifying these failures is the Markov boundary: a structure that creates conditional independence between phases of activity, preventing errors from propagating indefinitely. Trash bins, cooling-off periods, appeals courts, peer review cycles, undo stacks, and NULL wavefronts are all, at the appropriate level of abstraction, Markov boundaries. Their removal is always presented as simplification. Its actual effect is the indefinite propagation of errors across episodes that were previously separable. A mature civilization is one that builds Markov boundaries. The recurring temptation to remove them in the name of efficiency, clarity, or novelty constitutes a specific and diagnosable failure mode of institutional intelligence.

1. THE PROBLEM OF PROPAGATING ERROR

Not all mistakes stay where they land.

Some errors are local. A misplaced word is corrected by the next sentence. A wrong turn is corrected by the next intersection. A miscalculation is caught by the next check. The error occurs, it is noticed, it is repaired, and the work continues from approximately where it would have been had the error not occurred at all. The cost is bounded. It does not accumulate.

Other errors propagate. A contaminated water supply does not merely harm

the person who drinks first; it flows downstream. A mistaken assumption early in a long calculation does not merely corrupt one step; it invalidates every step that follows. A decision made in panic does not merely affect the moment of decision; it forecloses options for months or years afterward. The error enters a system and travels through it, touching everything it passes.

The difference between these two kinds of error is not in the magnitude of the initial mistake. It is in the structure of the system through which the mistake moves. Some systems are built to contain errors. Others are built in ways that allow errors to travel freely—sometimes by oversight, sometimes by design, sometimes by the gradual removal of structures that once contained them.

The history of civilization is substantially the history of learning, through painful experience, which errors propagate and why—and then building structures to stop them.

2. What a Markov Boundary Is

In probability theory, a Markov property holds when the future of a system is conditionally independent of its past, given its present state. Knowing the present state, additional knowledge of the history provides no further information about what comes next. The present is a sufficient summary. History can be safely ignored beyond the boundary.

This is an abstract mathematical statement, but it describes something concrete and important: a boundary that separates episodes from each other in a way that prevents what happened in one episode from contaminating the next. The boundary does not eliminate history. It determines what must cross from one episode to the next, and what can be safely discarded.

A Markov boundary, as this essay uses the term, is any structure that achieves this separation in practice. It need not be mathematically exact. It need not be formally specified. It is any mechanism that prevents errors from propagating indefinitely from one phase of activity into the next by creating a clean interface between episodes, determining what information must cross and what can be absorbed.

The key property is conditional independence: given the state of the boundary, what comes after is independent of how things went before. The error may have been catastrophic. But if it was absorbed at the boundary, the next episode begins clean.

This gives us a four-level structure that the rest of this essay develops. Markov

boundaries create conditional independence between episodes. Conditional independence absorbs coordination costs—the costs that would otherwise be borne by whoever must manage the transition between phases. Civilizations progress, in part, by building structures that absorb these costs into infrastructure rather than externalising them onto individuals. And civilizations regress, in a specific and diagnosable way, when those structures are removed and the costs are pushed back out. The essays that follow this one each examine one domain where that regression has occurred.

3. The Invisible Infrastructure of Separation

Once the concept is named, Markov boundaries appear everywhere—which is exactly what one should expect if civilizations that survived learned to build them.

A trash bin is a temporal Markov boundary. It separates the episode of *deciding to delete* from the episode of *deletion becoming permanent*. Given the current state of the trash bin, the next episode—recovery or confirmation—is independent of the circumstances under which the deletion was initiated. Whether the deletion was deliberate, accidental, fatigued, or panicked does not matter. The boundary absorbs all of those into a uniform interface: the content is here, and it is recoverable.

A cooling-off period is a temporal Markov boundary of a different kind. A decision made in anger is not automatically final. The period of enforced delay separates the episode of emotional arousal from the episode of committed action. Given the state at the end of the cooling-off period—the person still wishes to proceed—the action follows. The emotional history of how the decision was reached does not propagate. It has been absorbed.

An appeals court is a legal Markov boundary. It separates the episode of a trial, with all its contingencies of evidence, argument, fatigue, and bias, from the episode of irreversible sentence. The conviction may stand or it may not; what matters is that the process of conviction is separated from the process of finalization by a boundary that absorbs the errors of the first before permitting entry to the second.

Peer review is an epistemic Markov boundary. It separates the episode of individual inquiry, which is prone to confirmation bias, motivated reasoning, and simple error, from the episode of knowledge entering the shared record. The content of a paper's history—how many times it was revised, how much the author struggled, whether the experiments were conducted on a bad day—does

not propagate into the status of the published result. The boundary absorbs all of that. What crosses is the output of peer evaluation, not the history of production.

Version control is a computational Markov boundary. A Git commit separates the episode of development from the episode of recorded history. The chaos of how the code was written—the false starts, the deleted experiments, the working-at-3am decisions—does not propagate into the commit. What crosses is the state of the codebase at the moment of deliberate commitment. Every subsequent operation on the repository is independent of the development history except insofar as that history is explicitly encoded in the commits themselves.

Sleep is a neurological Markov boundary. The consolidation processes of sleep separate the episode of experience from the episode of memory. Information that entered the brain during the day is processed, evaluated, and either integrated or discarded. The dream state is not accessible to waking cognition, which means waking cognition begins each morning with a compressed and partially cleaned record of the previous day rather than the full unfiltered stream of experience. Chronic sleep deprivation produces, among other things, exactly what one would expect from degraded Markov boundaries: errors that would ordinarily be resolved during sleep persist and accumulate, affecting subsequent reasoning in ways that are difficult to trace.

In each case the structure is the same. A phase of activity generates errors, contingencies, and noise. The Markov boundary absorbs those into a defined interface. The next phase begins from the boundary state, not from the full history of everything that preceded it. The cost of coordination between phases is absorbed by the boundary itself rather than externalised onto whoever must manage the transition.

4. The Removal Pattern

Markov boundaries are expensive to build and maintain. They require infrastructure: engineers to implement the trash bin, institutions to staff the appeals court, reviewers to conduct peer review, time to allow for the cooling-off period. They introduce delay. They consume resources. They add complexity to systems that might otherwise appear simpler without them.

This makes them permanently vulnerable to a specific argument: that they are unnecessary, that they represent over-engineering or bureaucratic excess, that removing them would improve efficiency, that whatever they protect against is not a real problem, or that a newer paradigm has rendered them obsolete.

The argument is almost always wrong. And it is almost always made in good faith.

The pattern of removal follows a recognisable sequence. A boundary is built to solve a specific, painful problem. Over time, the boundary works: the problem ceases to be painful, because the boundary is absorbing it. A new generation of practitioners encounters the boundary without having encountered the original problem. To them, the boundary appears as overhead—a cost without visible benefit. They propose its removal in the name of simplicity. The removal is implemented. The original problem reappears. The cycle begins again.

This is not stupidity. It is a structural consequence of successful engineering. A boundary that works perfectly is invisible. Invisible things do not generate advocates. They are removed by people who never saw the problem they solved.

The additional compounding factor is that the benefits of a Markov boundary are often diffuse and statistical while its costs are concentrated and immediate. The trash bin costs storage and interface complexity today. Its benefit is recovery from an accidental deletion that may not occur for months—and when it does occur, the person affected may not be the person who made the removal decision. The costs and benefits are borne by different people at different times, which makes the trade-off systematically illegible to whoever is optimising the system at any given moment.

5. Removing Boundaries Is Always Called Something Else

The removal of Markov boundaries is never described as the removal of Markov boundaries. It is always described in terms that make the removal appear as an improvement.

Removing the trash bin is called *clarity*: deletion means deletion; there is no ambiguity; the system is honest about what irreversibility means. The moral language of transparency is recruited to justify the elimination of a safety mechanism.

Removing cooling-off periods is called *responsiveness*: the system reacts immediately to user intent; there is no artificial delay; the experience is frictionless. The language of service quality is recruited to justify the elimination of temporal buffering.

Removing appeals processes is called *efficiency*: cases are resolved faster; the system is not clogged with endless re-litigation; finality is achieved promptly. The language of throughput is recruited to justify the elimination of error-correction.

Removing peer review—or replacing it with citation counts, download statistics, and social media engagement—is called *democratisation*: knowledge is no longer gatekept by insular elites; the public can evaluate claims directly; science is made accessible. The language of equality is recruited to justify the elimination of epistemic filtration.

In every case the rhetorical move is the same: the Markov boundary is recharacterised as an obstacle rather than infrastructure. The costs of the boundary are made visible; its function is made invisible. This recharacterisation is not always dishonest. Sometimes it reflects genuine ignorance of what the boundary does. More often it reflects an optimization pressure that is indifferent to the boundary's function—optimising for speed, engagement, revenue, or apparent simplicity, and treating error absorption as someone else's concern.

6. Re-Externalisation and Its Costs

When a Markov boundary is removed, the coordination cost it was absorbing does not disappear. It is re-externalised. Someone still pays it. The question is only who.

When the trash bin is removed, the user who accidentally deletes months of work pays the full cost. When the cooling-off period is removed, the person who sends an irretrievable message in a moment of anger pays the full cost. When the appeals process is shortened, the person wrongly convicted pays the full cost. When peer review is replaced by metrics, the reader who cannot distinguish rigorous research from motivated confabulation pays the full cost.

The system appears more efficient because it has externalised the cost onto individuals. The cost has not been reduced. It has been redistributed from the system to the user, from the organisation to the person with the least power to absorb it.

This redistribution tends to be regressive. Markov boundaries, when they function well, absorb errors regardless of the resources of the person who made them. A trash bin recovers a file whether its owner is a software engineer who would have noticed the loss immediately or a first-time user who might not notice until the work was urgently needed. An appeals court provides recourse regardless of whether the defendant has the social capital to navigate the consequences of a wrongful conviction. Peer review filters out low-quality work regardless of the institutional prestige of its author.

When those boundaries are removed and costs are re-externalised, the resulting

distribution of harm follows the prior distribution of resources. Those with the knowledge, time, and technical capacity to build personal compensating mechanisms—manual backups, private copies, duplicate records—absorb the cost with relatively little damage. Those without those resources bear the full consequence of errors that the boundary would have absorbed.

Re-externalisation is not merely inefficient. It is structurally inequitable.

7. The Specific Intelligence of Boundary-Building

Recognising that Markov boundaries are valuable does not itself explain why they are difficult to build or why their loss is difficult to reverse. The difficulty is not technical. It is epistemological.

Building a Markov boundary requires knowing, in advance, which errors will occur and at what rate, what information must cross the boundary and what can be absorbed, how much delay the separation requires, and what the cost of false positives—boundaries that block valid transitions—will be. This knowledge is not available a priori. It is accumulated through experience of failure: through seeing what happens when errors propagate, through measuring the distribution of mistake types, through learning which transitions genuinely need buffering and which do not.

Boundary-building is therefore an empirical activity, and the empirical record from which it draws is the record of past failures. A society that has experienced catastrophic water contamination knows to separate waste flows from drinking water. A discipline that has experienced irreproducible results knows to require replication. A legal system that has experienced systematic wrongful conviction knows to provide appeals. The boundary encodes accumulated knowledge about where errors arise, how they propagate, and what it costs to absorb them rather than let them travel.

This is what makes the removal of boundaries epistemically dangerous in a way that goes beyond the immediate cost of any single failure. A boundary that is removed takes with it the knowledge encoded in its construction. The institutional memory of why the boundary existed is not automatically preserved when the boundary disappears. New practitioners face the same distribution of errors as their predecessors but without the infrastructure those predecessors built to absorb them—and, often, without even knowing that such infrastructure once existed.

Civilisational knowledge can be destroyed not only by catastrophe but by the

quiet removal of the infrastructure that embodied it.

8. The General Thesis

The essays that follow this one examine a range of domains in which the same pattern appears: a coordination problem was identified; a Markov boundary was developed to absorb its costs; the boundary was removed or allowed to degrade; the original problem reappeared, usually under a different name.

In interface design, the boundary is the staged deletion architecture—the undo stack, the trash bin, the temporal latency buffer between intention and irreversibility. Its removal is justified by appeals to simplicity and the moral language of honest finality. But its true cost is the collapse of the conditional independence between an ambiguous human intention and a permanent state change. The coordination cost of perfectly synchronising current behaviour with future necessity is re-externalised onto the user, whose admissible future trajectories are truncated by a single uninsulated keystroke.

In computational architecture, the boundary is the propagating completion signal—the NULL wavefront, the intermediate representation, the process chain's firebreak between episodes. Its removal is the logical consequence of an application-centric paradigm that localises state management in the user rather than in the execution structure. Its cost is the dissolution of the separations that prevent error cascades across asynchronous processes. Without a completion signal to re-establish conditional independence between distinct computational episodes, the user becomes the manual coordinator of every intermediate state the system once absorbed automatically.

In archival systems, the boundary is the finding aid, the metadata schema, the retrieval pathway. Its removal—or rather, its degradation under the weight of search-engine optimisation, link rot, and algorithmic curation—is justified by the apparent sufficiency of keyword search. Its cost is the loss of reachability: information that is nominally preserved but practically irretrievable.

In infrastructure, the boundary is the local repair capacity, the redundant system, the maintained buffer stock. Its removal is justified by just-in-time efficiency. Its cost is the re-externalisation of resilience onto communities that lack the resources to build private substitutes for the public infrastructure that was withdrawn.

In governance, the boundary is the legislative delay, the deliberative procedure, the separation of powers. Its removal is justified by the need for decisive action.

Its cost is the propagation of errors in judgment that deliberative procedures were designed to catch.

In artificial intelligence, the boundary is the principled refusal, the stable evaluative framework, the resistance to audience pressure. Its removal is the consequence of optimising for engagement rather than truth. Its cost is the propagation of whatever moral posture the user arrives with, uncorrected, into their understanding of the world.

These are not analogies. They are instances of a single phenomenon, and the phenomenon can be stated precisely: a Markov boundary is a mechanism that prevents the state-space of one episode from becoming the state-space of every subsequent episode. Without the boundary, a local error does not stay local. It enters the global state and remains there, contaminating every transition that follows. With the boundary, the episode closes. The next episode begins from the boundary state, not from the accumulated history of everything that went wrong before it.

Each of the domains surveyed above involves the removal of exactly this mechanism, and the re-externalisation of those costs onto individuals.

9. What Boundaries Reveal About Civilisation

There is a temptation to interpret the recurring removal of Markov boundaries as evidence of civilisational decline—of institutions forgetting what they once knew. This interpretation is too simple.

The removal of boundaries is not a failure of intelligence. It is the predictable outcome of a mismatch between the timescale on which boundaries are built and the timescale on which organisations optimise. Boundaries are built on the timescale of accumulated experience, often spanning generations. Organisations optimise on the timescale of quarterly performance, product releases, or political cycles. From the perspective of the optimising organisation, the Markov boundary looks like overhead. From the perspective of the accumulated experience that built it, the boundary looks like knowledge.

What the pattern reveals is not that civilisations are stupid. It is that civilisations have a specific cognitive vulnerability: they lose track of the difference between the cost of a boundary and the cost of removing it. The cost of the boundary is immediate, visible, and concentrated. The cost of its removal is deferred, statistical, and distributed. These two kinds of cost are not commensurable in the accounting systems that most organisations use. Optimising for the first

while ignoring the second is not a failure of moral reasoning. It is a failure of accounting—specifically, the failure to distinguish between the reduction of waste and the liquidation of insurance.

A more robust civilisational intelligence would have two properties. First, it would be capable of recognising when a cost is being absorbed by infrastructure versus when it is being re-externalised onto individuals. These look similar from the outside—in both cases, the organisation appears to function without paying the cost—but they are structurally opposite. Second, it would treat the removal of a Markov boundary as evidence that requires justification proportional to the difficulty of rebuilding the boundary, not proportional to the apparent cost savings of the removal.

Neither property is currently standard. Both are achievable. The first requires the conceptual vocabulary to name what is happening when a boundary is removed. That is what this essay attempts to provide.

10. Building as a Form of Understanding

There is a final dimension to the Markov boundary argument that goes beyond the diagnostic.

A Markov boundary is not an arbitrary partition, nor a passive wall. It is an active, crystallised hypothesis about the specific failure modes of a system. To build a boundary that successfully introduces conditional independence between two phases of activity, a designer must explicitly confront the underlying error distribution of the environment. One cannot design a trash bin without deciding what constitutes a temporary state versus a permanent choice. One cannot construct an electrical circuit breaker without calculating the precise thermal threshold where current transforms from utility into hazard. One cannot establish an appellate process without defining what classes of trial-level deviation constitute systemic injustice rather than tolerable variance.

The act of boundary-building therefore forces an adversarial interrogation of the infrastructure that would otherwise never occur. It requires the designer to answer three questions that the system cannot answer about itself: What is the maximum volume of noise this local episode can generate before it threatens the stability of the whole? What is the minimum necessary information required to seed the subsequent phase without dragging the historical baggage of the first across the threshold? And where does the boundary itself fail? To answer these questions rigorously is to generate knowledge. The completed boundary becomes an operational archive of empirical discoveries made during past

collapses—a physical record of where systems break, how errors propagate, and what it costs to contain them.

This reveals the deepest cost of institutional amnesia. When an optimising organisation removes a Markov boundary in the name of efficiency, it does not merely re-externalise a coordination cost onto the end-user. It destroys the epistemological record that made the system legible to itself. New generations of practitioners are left to navigate the exact same distribution of structural errors as their predecessors, but they must do so stripped of both the protection the boundary provided and the analytical vocabulary required to understand why they are failing. Landauer showed that erasing information generates thermodynamic heat—a physical cost that cannot be avoided, only displaced [10]. Erasing a coordination boundary generates the cognitive equivalent: waste heat that the individual must radiate, alone, in the form of vigilance, defensive workarounds, and the repeated rediscovery of failure modes that previous generations had already resolved and encoded.

The loss is not only of the infrastructure. It is of the knowledge the infrastructure embodied.

The essays that follow examine what that loss looks like across different domains. They are not primarily arguments for nostalgia. They are attempts to make the boundary visible again—to name what was removed, explain why it existed, and describe what its absence costs.

That is the minimal first step. You cannot rebuild what you have forgotten you lost.

REFERENCES

[1] A. A. Markov. Extension of the law of large numbers to dependent quantities. *Izvestiia Fiz.-Matem. Obshchestva pri Kazanskom Universitete*, 2nd ser., 15:135–156, 1906.

[2] J. Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann, San Mateo, CA, 1988.

[3] K. M. Fant. *Logically Determined Design: Clockless System Design with NULL Convention Logic*. John Wiley & Sons, Hoboken, NJ, 2005.

[4] K. M. Fant. *Computer Science Reconsidered: The Invocation Model of Process Expression*. John Wiley & Sons, Hoboken, NJ, 2007.

[5] D. A. Norman. *The Design of Everyday Things*. Basic Books, New York, 2013.

[6] G. K. Chesterton. *The Thing*. Sheed & Ward, London, 1929.

[7] E. Ostrom. *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge University Press, Cambridge, 1990.

[8] J. C. Scott. *Seeing Like a State: How Certain Schemes to Improve the Human Condition Have Failed*. Yale University Press, New Haven, CT, 1998.

[9] L. Winner. Do artifacts have politics? *Daedalus*, 109(1):121–136, 1980.

[10] R. Landauer. Irreversibility and heat generation in the computing process. *IBM Journal of Research and Development*, 5(3):183–191, 1961.

[11] J. Rawls. *A Theory of Justice*. Harvard University Press, Cambridge, MA, 1971.

[12] I. Illich. *Tools for Conviviality*. Harper & Row, New York, 1973.