

Time Is Not Authority

On the Structural Failure of Temporal Control Systems

January 30, 2026

Mats Heming Julner

Abstract

Modern digital systems increasingly treat temporal progression as a source of legitimacy: actions execute because time advanced. This paper argues that this is a fundamental category error. Time can order events, but it cannot authorize them. When systems collapse logging and execution and infer permission from sequence, they become structurally unstable at scale. This failure mode now underlies financial automation, blockchains, autonomous agents, and large-scale AI systems. We propose explicit authorization as a missing control primitive and argue that no time-driven system can remain stable without it.

1. The category error

Time is an ordering relation. Authority is a legitimacy relation. They are not the same. Time can answer questions about sequence: what came before, what came after, and what happened when. It can describe the structure of events and their temporal relationships. However, time cannot answer questions about legitimacy. It cannot tell us who is allowed to act, whether an action should occur, or whether execution is justified. Despite this, modern systems routinely assume that temporal order implies legitimacy. If an event occurs later in a sequence, it is treated as valid to execute. This assumption is rarely stated explicitly, but it governs nearly all automated infrastructure. Actions occur because a scheduler fired, a block was produced, a loop continued, or a timer expired. In each case, temporal progression itself becomes the implicit source of authority. We refer to systems in which execution is triggered solely by temporal progression, rather than by explicit permission, as temporal control systems. This substitution of sequence for legitimacy is the Temporal Authority Fallacy.

Time's apparent special status is an illusion. In any real system, time is merely one coordinate among many in the system's state: alongside balances, positions, identities, prices, and physical or informational location. Treating time as a source of legitimacy is therefore equivalent to treating any other coordinate as a source of legitimacy. A system that allows actions because a timestamp advanced is not meaningfully different from one that allows actions because a block height, account balance, or geographic position changed. The error is not that time was misused, but that a descriptive axis of state was promoted into a control variable. Temporal authority is thus not a special case; it is one instance of a more general category error: mistaking state for permission.

2. Logging, execution, authorization

In any control system, three layers are conceptually distinct. Logging records that something happened. Execution makes something happen. Authorization grants permission for something to happen. In well-designed systems, these layers are separated: authorization precedes execution, and execution precedes logging. Someone decides, something occurs, and then it is recorded. Modern

digital systems collapsed this structure. Logging became execution, and execution was inferred as authorization. The control flow was inverted. Instead of authorization determining what may happen, time determines what happens, and whatever happens is assumed to have been authorized. Execution is no longer grounded in explicit permission. It is inferred from the fact that the system is running, the event occurred, the state updated, or the sequence advanced. Permission is no longer checked. It is assumed.

3. Why time replaced authority

This substitution did not occur for philosophical reasons. It occurred for purely technical ones. Clocks are universal. Schedulers are simple. Time is machine-native. Humans are slow, ambiguous, and difficult to formalize. Faced with the need for reliable triggers, engineers did what was convenient: they used time. Once enough systems adopted this pattern, the implicit rule became that if a process continues, it may act. Time quietly became the de facto authority layer; not because it was correct, but because it was easy. Over time, this control logic became invisible. Entire infrastructures shifted from being governed by explicit decisions to being governed by continuity itself.

3.1 Relation to existing work

Prior work analyzes instability, incentives, and feedback failures across domains, but typically treats time-triggered execution as a neutral implementation detail. This paper argues that, at scale, temporal triggering becomes a de facto legitimacy rule, and that this is the missing category. This paper does not argue for improved governance, oversight, or monitoring of automated systems. It addresses a different and more fundamental layer: the source of legitimacy for execution. Existing work on access control, workflow engines, governance, accountability, and human-in-the-loop systems assumes that time-triggered execution is structurally valid and merely requires better supervision. In contrast, this paper argues that execution itself must be gated by explicit authorization, not by temporal progression.

Classical control theory demonstrates that systems without effective negative feedback cannot stabilize (Wiener, 1948; Ashby, 1956), but does not identify time-based triggering as a source of authority. Financial and AI safety literature documents instability, incentives, and error propagation (Minsky, 1986; Brunnermeier & Oehmke, 2013; Amodei et al., 2016), yet does not isolate temporal progression itself as the control primitive. The contribution here is to identify temporal authorization as the missing category.

4. An informal control claim

Consider a system in which execution is justified by continuity, and continuity never stops. Such a system has no semantic brake, no natural interrupt, and no stable equilibrium. It can only accelerate, propagate, and amplify. This is consistent with classical control theory: a system without a negative feedback channel cannot regulate itself (Wiener, 1948; Ashby, 1956). Time cannot say “no”. It cannot pause, revoke, or withdraw permission. It can only advance.

5. Finance as an early failure

Digital finance adopted a temporal model of authorization long before it was fully computerized. Interest accrues because time passes. Margin calls and liquidations trigger when time-based state transitions occur. Payments execute on dates. Leverage is rolled forward as cycles complete. In each case, continued operation is treated as implicit permission to proceed. Risk is not re-authorized; it is allowed to persist because nothing has yet stopped it. In DeFi this temporal triggering is literal (block-time), while in TradFi it appears as calendar time, settlement windows, margin cycles, and continuously recomputed risk limits; the mechanism differs, the control primitive is the same. The 2008 financial crisis was not primarily a failure of intelligence or incentives. It was an early instance of time being treated as authority. Risk executed because the system progressed, not because anyone explicitly renewed permission for that risk to exist. This aligns with Minsky’s financial instability hypothesis (Minsky, 1986), in which periods of stability encourage progressively riskier financial structures, and with later work on maturity mismatch and reflexive leverage (Brunnermeier & Oehmke, 2013). However, this paper reframes these dynamics not as pricing failures, but as failures of authorization: historical survival was mistaken for legitimacy, allowing exposure to grow without explicit, bounded consent.

6. Blockchains as the limit case

Blockchains are pure temporal authority machines. They impose a global ordering and treat that ordering as reality. In a blockchain, whatever is next in sequence must execute. There is no concept of permission, intent, context, or legitimacy beyond ordering. Once something is included in the canonical sequence, it produces consequences by definition (Nakamoto, 2008). All power therefore collapses into control over time slots. Empirical work on miner extractable value (Daian et al., 2020) and high-frequency market design (Budish, Cramton & Shim, 2015) demonstrates that temporal priority determines economic outcomes. Bitcoin does not merely reproduce the structural failure of digital finance. It isolates it. By stripping away institutions, discretion, interpretation, and external authority, Bitcoin reduces financial control to a single primitive: temporal ordering. Legitimacy becomes nothing more than being next in sequence.

Because execution is bound to temporal priority, blockchains systematically reward latency advantage and order-positioning. Prices, liquidations, and state transitions become highly sensitive to congestion, inclusion uncertainty, and order-flow shocks. Volatility, in this sense, is not only a property of speculative demand; it is also a structural consequence of legitimacy being assigned by sequence under throughput constraints. As participation scales, these systems tend to oscillate between periods of smooth inclusion and periods of contention, because the scarce resource is always priority in the next slot.

7. AI as the amplification

AI systems inherit the same control model. Agents act because loops continue. Models update because schedules fire. Decisions propagate because state advances. But AI systems differ in one crucial way: they increasingly control real-world execution. Not just money, but logistics, infrastructure, information, coordination, and decision pipelines. Concerns about specification gaming and unintended action propagation (Amodei et al., 2016; Russell & Norvig, 2021) reflect this structural risk. AI is not dangerous because it is intelligent. AI is dangerous because it is a time-based control system with global execution power.

8. The missing primitive

What all these systems lack is a first-class representation of authorization. Authorization must be explicit, separate from execution, inspectable, time-bounded, and revocable before action. Without this, systems default to a single rule: if execution is possible, it is allowed. This rule is survivable in slow human systems and catastrophic in fast automated ones. By authorization, we mean an explicit, inspectable permission state that can be verified prior to execution and revoked prior to effect. Authorization is distinct from capability, scheduling, and historical precedent.

9. Reversing the control flow

The correct architecture is:

$$\textit{Authorization} \rightarrow \textit{Execution} \rightarrow \textit{Logging} \rightarrow \textit{Time}$$

Not:

$$\textit{Time} \rightarrow \textit{Logging} \rightarrow \textit{Execution}$$

Time may describe history. It must never govern action. Sequence is memory. Authorization is control.

10. Social media as temporal authority

The same structural failure appears in contemporary social media systems, although its effects manifest psychologically rather than economically. Social media platforms operate as time-driven execution systems applied to human attention. Feeds refresh, timelines update, and recommendations are recomputed continuously as a function of temporal progression. As in financial and computational systems, action is justified by continuity: if time advances, the system acts. In this domain, execution corresponds to exposure. Content is delivered because the system is running, not because any contextual authorization or completion condition has been satisfied. This mirrors the engagement-optimization logic documented in large-scale recommender systems, which are trained to maximize watch time, dwell time, and interaction rather than experiential closure (Covington, Adams & Sargin, 2016; Zuboff, 2019). Human communication, however, relies on bounded interaction and situational permission. By replacing these with continuous time-based triggers, social media becomes structurally incapable of representing the state “this interaction should end”. The resulting restlessness and compulsive engagement are therefore not failures of willpower, but direct consequences of treating temporal sequence as legitimacy in a domain that requires the ability to stop.

11. Conclusion

Modern systems did not fail because they became intelligent, decentralized, or automated. They failed because they substituted temporal order for legitimacy. Time can record events. Time cannot justify them. Until explicit authorization becomes a first-class control primitive, large-scale digital systems will remain structurally unstable by design. This is not a moral claim. It is a control-theoretic one. Time is not authority.

References

- Nakamoto, S. (2008). *Bitcoin: A Peer-to-Peer Electronic Cash System*.
- Daian, P., Goldfeder, S., Kell, T., Li, Y., Zhao, X., Bentov, I., Breidenbach, L., & Juels, A. (2020). Flash Boys 2.0: Frontrunning in decentralized exchanges, miner extractable value, and consensus instability. In *2020 IEEE Symposium on Security and Privacy (SP)*.
- Budish, E., Cramton, P., & Shim, J. (2015). The high-frequency trading arms race: Frequent batch auctions as a market design response. *The Quarterly Journal of Economics*, *130*(4), 1547–1621.
- Minsky, H. (1986). *Stabilizing an Unstable Economy*.
- Brunnermeier, M. K., & Oehmke, M. (2013). The maturity rat race. *The Journal of Finance*, *68*(2), 483–521.
- Wiener, N. (1948). *Cybernetics: Or Control and Communication in the Animal and the Machine*. MIT Press.
- Ashby, W. R. (1956). *An Introduction to Cybernetics*. Chapman & Hall.
- Russell, S., & Norvig, P. (2021). *Artificial Intelligence: A Modern Approach* (4th ed.). Pearson.
- Amodei, D., Olah, C., Steinhardt, J., Christiano, P., Schulman, J., & Mané, D. (2016). Concrete problems in AI safety. *arXiv:1606.06565*.
- Zuboff, S. (2019). *The Age of Surveillance Capitalism*.
- Covington, P., Adams, J., & Sargin, E. (2016). *Deep Neural Networks for YouTube Recommendations*. In *Proceedings of the 10th ACM Conference on Recommender Systems (RecSys '16)*.