

Beyond State: Process, Time, and Legitimacy in Blockchain Systems

Preprint , Position Paper , Work in Progress

Version: v0.2 (v0.1: January 2026)

Date: July 2026

Author

Oualid Barhoumi

NOORCHAIN Research

Correspondence: contact@noorchain.io

This document is a preprint position paper exploring conceptual limits of state-based blockchain systems and proposing a process-oriented perspective on trust, time, and legitimacy.

NOORCHAIN is referenced as an experimental instantiation under active development, operating in controlled environments. This document does not constitute a technical specification, protocol implementation, roadmap, production-readiness statement, external audit or security attestation, legal or regulatory advice, tokenomics document, or investment proposal.

The ideas presented are intended to contribute to ongoing research and discussion at the intersection of distributed systems, cryptography, and institutional design.

Status

- This document may be revised.
- Feedback is welcome.
- Citations should reference the version number.
- Changes in v0.2: discussion of time-based primitives (VDFs and related constructions) added in section 3; one clarifying sentence on PoSS snapshot mechanics added in section 7; author and correspondence details completed. No other content changed.



1. Introduction , The State-Centric Assumption

Over the last decade, blockchains have demonstrated a robust capacity to secure shared digital states. Through well-understood cryptographic mechanisms and distributed consensus protocols, they enable parties with limited mutual trust to maintain a common ledger that is verifiable and resistant to tampering. This technical achievement has supported a wide range of applications, from digital assets and decentralized finance to public registries and globally accessible application infrastructures.

This success, however, rests on an assumption that is rarely examined directly: that truth, security, and legitimacy can be fully expressed and therefore secured through a final state. In the prevailing paradigm, a system is considered sound once its current state is valid, its state transitions obey formal rules, and the associated signatures or proofs are cryptographically correct. This state-based approach is coherent and highly effective for domains where the objects of interest are predominantly informational or financial.

When blockchains attempt to extend beyond the recording of assets and into human, social, or institutional systems, the limits of that assumption become evident. In such contexts, legitimacy does not arise solely from the final result; it depends on the process by which that result is reached. Procedure, temporal continuity, the absence of opportunistic shortcuts, and the persistence of commitments often matter as much as the outcome itself in determining what is considered valid, acceptable, or fair.

This tension exposes a structural blind spot in many current architectures. Blockchains are well suited to proving that a state is correct, but they are far less capable of demonstrating that a constrained path has been followed. They can certify that a decision was recorded, but not necessarily that it emerged from a legitimate process, one that was not manipulated and that complied with explicit temporal and procedural rules. In other words, they secure outcomes effectively, yet they secure trajectories only weakly.

This is not a critique of cryptography. Modern primitives, signatures, hashes, and zero-knowledge proofs, deliver the guarantees they were designed to provide: integrity, authenticity, and formal validity. The issue lies elsewhere: in the gap between what cryptography can prove efficiently and what human systems require to establish legitimacy.

Across institutional settings, trust is rarely anchored to a single instant. It accumulates, must be maintained, and can be lost when continuity is broken. It is not easily compressible; it is not trivially accelerable; and it is intrinsically linked to time. Yet most blockchain architectures treat time primarily as a technical parameter, something to minimize, optimize, or abstract away, rather than as a foundational constraint of security and legitimacy.

This paper advances the hypothesis that addressing these limits does not necessarily require a new cryptographic primitive or an alternative consensus design. Instead, it calls

for a shift in conceptual framing: from the validity of states to the verifiability of processes and trajectories over time. We refer to this as process-based trust, an approach that treats time, continuity, and non-compressibility as first-class elements of what a distributed system can meaningfully secure.

On this basis, we examine structural constraints of state-based systems, outline desirable properties for process-oriented mechanisms, and discuss implications for governance and institutional design. We then refer to NOORCHAIN as an early experimental instantiation of these ideas, while explicitly acknowledging the open questions and engineering challenges that remain.

2. State-Based Truth vs Process-Based Reality

Contemporary blockchain systems rely on an implicit but foundational model of truth: truth defined by a state. A state is considered valid when it satisfies formal rules and is accepted through a consensus mechanism. This has proven effective for systems in which value and coherence can be captured by discrete variables, balances, ownership, permissions, or computable outcomes.

Within this model, security is largely measured by the ability to ensure that each state transition is correct and that the resulting state cannot be altered without violating explicit cryptographic assumptions. The underlying premise is that if the state is valid, the system is valid. This equivalence holds in domains where the manner in which a result is obtained carries limited significance compared with the result itself.

Human and institutional systems operate under different premises. Legitimacy is not attached only to outcomes; it depends on how those outcomes are produced. The same end state may be viewed as legitimate or illegitimate depending on the path taken to reach it. A decision may be formally correct yet lose legitimacy if it was obtained by bypassing procedure, by artificially compressing time, or by exploiting informational asymmetries.

This distinction highlights a basic asymmetry: blockchains can represent and verify states with high precision, yet they provide few native mechanisms for representing trajectories. A trajectory is not merely a sequence of states. It may impose constraints on order, duration, continuity, and sometimes irreversibility. It is tied to time not as a simple index, but as a condition of validity.

In most protocols, time appears as a secondary technical parameter, timestamps, block numbers, and intervals. This abstraction is sufficient to order transactions and coordinate execution, but it is insufficient to express requirements such as sustained commitment, persistent behavior, or the inability to compress a long process into an instantaneous action.



The limitation becomes particularly salient in decentralized governance and collective coordination. Existing mechanisms often reduce complex processes to aggregated states: a final vote, a score, a quorum. In doing so, they lose crucial information about the dynamics preceding the result, stability of positions, continuity of participation, and the absence (or presence) of opportunistic manipulation over time.

This observation motivates a change in focus. Rather than indefinitely enriching the representation of state, it becomes necessary to ask how a distributed system might attest that a constrained process has occurred over time, even if only partially or indirectly. This transition, from state-defined truth to process-defined reality, provides the starting point for the framework developed in this paper.

Transition , On the Role and Limits of Smart Contracts

Smart contracts can already encode temporal constraints in blockchain systems. They can enforce delays, order sequences of actions, and reject state transitions unless predefined conditions are met. In that sense, they are essential tools for automating rules that depend on time.

However, these mechanisms remain limited to checking discrete, momentary conditions that are observable on-chain at execution time. A smart contract can establish that an action is valid at a given instant; it cannot, by itself, attest that the action resulted from a process that was actually followed over time, without interruption or artificial compression. Two trajectories that differ substantially in human or institutional terms may produce an identical and perfectly valid final state from the contract's perspective.

This distinction does not reflect a deficiency in smart contract engineering; it reflects a structural limitation of the evaluation model that smart contracts implement. Time is treated as a control parameter rather than as evidence of continuity. The gap between formal temporal constraints and lived trajectories is precisely where the problem addressed by this paper emerges.

3. Why Cryptography Alone Is Not the Missing Piece

When confronted with these limits, a common response is to invoke more advanced cryptographic mechanisms. It is tempting to treat the difficulties as a purely technical deficit that could be resolved by new primitives or more complex constructions. That interpretation is incomplete.

Cryptography excels at proving well-specified properties: possession of a secret, correctness of a computation, integrity of a message, or compliance of a state transition with formal rules. These capabilities are indispensable and form the security foundation



of modern distributed systems. Yet they operate within a framework in which the objects being proven are, by design, discrete and formalizable.

The issue discussed here is not that cryptographic tools fail at their intended tasks. Rather, it is that certain forms of legitimacy depend on properties that are not naturally expressible as static statements. Proving that a signature is valid does not establish that a commitment was maintained over time. Verifying that a computation is correct does not imply that intermediate steps were not bypassed through off-system optimization.

Even highly sophisticated cryptographic constructions struggle to capture properties such as non-accelerability, behavioral persistence, or adherence to implicit social procedures. These properties are tied to duration, repetition, and the absence of discontinuity, dimensions that cryptography, as a discipline, is not primarily designed to represent directly.

This includes the primitives explicitly designed around time. Verifiable Delay Functions (VDFs) and related constructions, timelock puzzles, proof-of-elapsed-time schemes, can prove that a sequential computation was performed and therefore that a minimum amount of computational time has elapsed. These are valuable tools, and they demonstrate that duration itself can be made verifiable. But they attest machine time, not institutional time: a VDF proves that a processor worked without shortcuts, not that a person or an organisation sustained a commitment, renewed a role, or remained exposed to the possibility of withdrawal over a period. The non-compressibility discussed in this paper concerns trajectories of human and institutional behavior, repeated, interruptible, and context-dependent, for which computational delay is at best a partial proxy. The same holds for adjacent notions in distributed systems: reputation systems aggregate historical outcomes without proving continuity, and accountability mechanisms in Byzantine fault-tolerant replication attribute misbehavior after the fact without qualifying the legitimacy of a compliant path. These tools are complementary to the perspective developed here, not substitutes for it.

It would therefore be mistaken to pursue an answer solely by adding cryptographic layers. Doing so risks obscuring the central point: some forms of trust cannot be established by verifying states or by producing point-in-time proofs alone. They require explicit attention to process, understood as a constrained sequence of actions unfolding across time.

Recognizing this boundary does not diminish cryptography. It clarifies its role: a necessary instrument within a broader framework that can incorporate temporal and procedural constraints. The research space explored in the remainder of this paper lies in that articulation, between cryptographic guarantees and the verifiability of trajectories.



4. Toward Process-Based Trust

If the limitations identified above are not resolved by finer-grained state validation, nor by stacking additional cryptographic primitives, then the reference frame itself must shift. Rather than encoding ever more meaning into instantaneous states, we need to treat the verifiability of processes as a first-class design objective. This paper uses the term process-based trust to describe that shift.

By process-based trust, we mean a form of trust that does not derive solely from the validity of a final result, but from the ability to establish that a constrained path was followed over time. That path is not reducible to a simple sequence of states. It includes requirements on the order of actions, their temporal spacing, their continuity, and, in some cases, the irreversibility of intermediate steps. Under this lens, trust is not anchored to a single moment; it is anchored to the coherence of a trajectory.

A process-oriented framework implies several distinctive properties.

First, non-compressibility. A process that carries institutional meaning should not be reproducible instantaneously through off-system preparation. If a trajectory can be reconstructed at negligible cost by staging work outside the system and publishing only the final act, then time has no structural role, and any trust attached to “participation over time” becomes fragile. In a process-based model, duration is not an optional parameter; it is part of the security surface.

Second, resistance to strategic shortcuts. Many mechanisms fail not because they are formally incorrect, but because they can be optimized in ways that satisfy the letter of the rules while violating their purpose. Process-based trust aims to limit this class of optimization, not necessarily by enumerating forbidden behaviors, but by making shortcuts non-equivalent to compliant trajectories, either by making them detectable ex post or by ensuring they do not produce the same standing in the system.

Third, ex post auditability. The goal is not to publish every detail of a process or to maximize transparency. Rather, it is to enable a third party to verify, after the fact that certain constraints were respected over time. This auditability can be partial, indirect, or probabilistic, but it must be sufficient to distinguish a compliant trajectory from an opportunistic one that happens to produce the same final state.

It is important to separate process-based trust from adjacent concepts that are often conflated with it. It is not a synonym for reputation, which aggregates historical outcomes without necessarily proving continuity or non-compressibility. It is not equivalent to randomness, which introduces unpredictability but does not attest to a path being followed. And it is not a replacement for consensus, which synchronizes a shared state but does not qualify the procedural legitimacy of how that state was achieved.



Adopting a process perspective does not imply that everything must be formalized or quantified. Human systems contain dimensions that cannot be reduced to strict rules without distortion. The claim is narrower and more practical: certain requirements, persistence, continuity, and resistance to acceleration can be partially captured through verifiable artifacts if those artifacts are designed explicitly to reflect process, rather than merely outcome.

This approach opens an intermediate space between two extremes: on one side, purely formal proofs over static states; on the other, opaque social processes that remain entirely off-chain and non-auditable. Process-based trust does not eliminate human judgment. It aims to constrain and stabilize some of the most manipulation-prone parts of institutional coordination by making time, continuity, and procedure observable in a way that can be checked and contested.

5. Time as a First-Class Constraint

In most blockchain architectures, time is treated as an auxiliary variable. It orders transactions, regulates block production, and enforces minimum delays before certain actions can occur. These functions are essential for coordination, yet they frame time primarily as an operational parameter something to be measured, indexed, and, where possible, minimized.

This framing reflects a broader engineering intuition: time is often understood as a cost. Latency, finality, throughput, and responsiveness are commonly optimized under the assumption that faster convergence toward a valid state is inherently desirable. Within this logic, a system is judged effective if it can produce correct outcomes quickly, regardless of the temporal structure of the process that led to those outcomes.

Institutional and social systems operate under a different logic. In many such systems, time is not merely a medium for coordination; it is a condition of legitimacy. Certain commitments derive their meaning from being sustained. Certain responsibilities can only be demonstrated through repetition and persistence. Trust, in these cases, weakens if it can be acquired instantaneously, without exposure to duration, uncertainty, or the possibility of withdrawal.

This divergence creates a structural tension. Blockchain systems tend to reward the ability to reach a valid state efficiently, while many human processes require precisely the opposite: the impossibility of compressing the path. When a system allows an actor to condense, into a single transaction, what is institutionally expected to unfold over time, it undermines the very significance of the commitment it claims to encode.

Treating time as a first-class constraint requires a change in perspective. The question is no longer whether an action occurs after a predefined delay, but whether duration itself



carries information that the system must preserve. A process that unfolds over time exposes participants to changing conditions, sustained effort, and the ongoing risk of disengagement. That exposure is not incidental; it is what gives weight to the process.

Within a process-oriented framework, time becomes structurally non-compressible. It cannot be substituted by an instantaneous proof, nor can it be fully simulated off-system without losing equivalence. Even when preparation occurs outside the system, the legitimacy of the process depends on the fact that certain steps were completed at distinct moments, under explicit constraints, and without the possibility of retroactive optimization.

This does not require perfect or continuous time measurement. What matters is not temporal precision but the existence of opposable temporal thresholds and irreversible transitions. Time functions as a form of institutional friction, limiting the capacity of actors to optimize locally without bearing long-term consequences. It introduces asymmetry between genuine participation and opportunistic replication.

Reintroducing time as a core constraint also clarifies several persistent ambiguities in decentralized systems. Concepts such as credibility, loyalty, or legitimacy cease to be purely subjective assessments and become partially grounded in observable trajectories. Time does not guarantee integrity, but it raises the cost of simulating integrity without incurring its associated obligations.

This perspective does not reject efficiency or scalability where they are appropriate. It suggests, rather, that not all functions should be optimized along the same temporal axis. For certain critical domains, governance, social validation, and recognition of sustained contribution relative slowness is not a deficiency but a security property.

By elevating time from a secondary parameter to a foundational constraint, distributed systems can begin to support forms of trust that cannot be reduced to instantaneous verification.

6. Governance, Legitimacy, and Anti-Arbitrariness

When blockchain systems address governance, the problem of security shifts toward a problem of legitimacy. The question is no longer limited to whether rules are enforced correctly, but extends to who defines those rules, how they evolve, and on what grounds their outcomes are accepted. In this domain, the limitations of purely technical approaches become especially visible.

A recurring observation in decentralized systems is that effective power does not reside solely in rule execution, but in rule specification. Thresholds, calendars, priority orders, and tie-breaking procedures are not neutral artifacts; they are governance decisions.

Even when these parameters are public and auditable, they retain a discretionary character, and with it the potential perception of bias, capture, or manipulation.

On-chain governance mechanisms often attempt to address this concern through voting and preference aggregation. While such tools are useful for expressing collective decisions, they do not eliminate arbitrariness at the structural level. Voting rules, electorate definitions, weighting schemes, and timing are themselves the result of prior human choices. The act of aggregation does not neutralize discretion; it merely relocates it.

A process-based perspective suggests a different point of intervention. Reducing arbitrariness does not necessarily require more participation or increasingly complex procedures. Instead, it involves identifying domains where human discretion adds limited value and replacing it with impersonal, deterministic, and verifiable constraints. The aim is not to eliminate governance, but to narrow the scope in which discretionary decisions can influence outcomes.

This approach implies a clear separation between two layers. The first concerns substantive decisions those that inherently require judgment, negotiation, and contextual reasoning. These remain irreducibly human. The second concerns procedural mechanics ordering, rotation, prioritization, equivalence resolution, where discretionary intervention often introduces contestability without providing corresponding benefits. It is primarily at this procedural layer that anti-arbitrary mechanisms can play a stabilizing role.

Procedural legitimacy, in this sense, does not depend on the perfection of the rules. It depends on their non-appropriability. When outcomes can be traced to known, impersonal constraints rather than to ad hoc intervention, disagreement becomes easier to absorb. An unfavorable result attributed to a transparent rule is more likely to be accepted than the same result perceived as the product of discretionary choice.

Public deterministic references whether mathematical, protocol-defined, or historically fixed can serve this function. Their value does not lie in optimality or randomness, but in the fact that they cannot be selectively chosen to favor particular actors. By removing the ability to decide ex ante which rule applies, such references reduce the space for strategic manipulation.

This logic does not abolish conflict, nor does it remove the need for governance. Instead, it clarifies boundaries. Governance remains responsible for defining objectives and adapting high-level rules, while process-level constraints limit how those rules can be exploited or reinterpreted opportunistically. Anti-arbitrariness is therefore not an end in itself, but a condition that allows governance processes to operate on a more stable and legible foundation.

By linking governance and legitimacy to verifiable procedural constraints, distributed systems can move beyond the aggregation of instantaneous preferences toward forms of

coordination that are more resilient to capture and contestation. This perspective prepares the ground for examining how these principles are instantiated, in a limited and experimental form, within NOORCHAIN.

7. NOORCHAIN as an Early Instantiation

At the time of writing, NOORCHAIN remains an experimental system under active development, operated in controlled environments and at limited scale. It is neither a broadly deployed network nor a production-hardened infrastructure. The discussion that follows should therefore be read as illustrative rather than definitive.

The arguments developed in this paper are primarily conceptual. They aim to clarify a class of limitations and to outline a direction of inquiry, not to prescribe a finalized architecture. Within that context, NOORCHAIN is presented as an early and partial instantiation of certain principles associated with a process-oriented approach to trust.

A central design choice in NOORCHAIN is the explicit separation between two domains that are often conflated in blockchain systems: ledger security and social or institutional legitimacy. Consensus is treated as a mechanism for technical integrity ensuring continuity and correctness of the shared state. Questions of recognition, contribution, and social validation are addressed in a distinct layer, deliberately decoupled from consensus.

This layer, referred to as Proof of Social Signal (PoSS), relies on signals emitted over time and aggregated into periodic snapshots. These snapshots are not intended to assert an instantaneous truth. Instead, they reflect temporal dynamics: repeated participation, continuity of engagement, and validation by designated actors referred to as curators. The standing associated with a signal depends less on its isolated occurrence than on its persistence and coherence across periods. Concretely, standing accumulates only through signals validated in distinct, dated snapshot periods; a contribution history cannot be batched into a single period or backdated, and curator validation is itself a dated, recorded act so both the participant's and the curator's trajectories remain auditable *ex post*.

The role of curators illustrates this orientation. Curators do not merely review data correctness; they help review whether certain processes appear to have been followed under explicit constraints. Their function is exercised over time, renewed through repeated action, and exposed to the possibility of withdrawal or challenge. The legitimacy associated with this role cannot be obtained instantly; it emerges from a trajectory that remains observable, even if imperfectly.

It is important to emphasize the limits of this approach. NOORCHAIN does not claim to fully formalize human processes or to eliminate ambiguity from social validation. The

mechanisms employed are necessarily approximate and depend on specific design choices. Several aspects such as fine-grained measurement of non-compressibility or complete formalization of trajectories remain open and are treated as hypotheses rather than guarantees.

These limitations are intentional rather than incidental. NOORCHAIN positions itself as a practical exploration environment, designed to test whether temporal constraints, persistent roles, and non-discretionary procedures can be combined within an operational distributed system. The objective is not to produce a universal model, but to examine the plausibility of a process-oriented approach under real-world constraints.

In this sense, NOORCHAIN is used to explore the central hypothesis of this paper: moving toward process-based trust does not require a new primitive or a redefinition of consensus. It emerges from a series of modest architectural choices that privilege duration, continuity, and resistance to strategic shortcuts. These choices do not replace human judgment; they seek to stabilize and bound some of its most manipulation-prone dimensions in a verifiable manner.

8. What This Paper Is Not

Given the scope and nature of the arguments advanced so far, it is important to clarify explicitly what this paper does not attempt to be. This clarification is not ancillary; it is essential to avoid misinterpretation, particularly in a field where conceptual proposals are often read as technical roadmaps, economic claims, or normative prescriptions.

First, this paper does not introduce a new cryptographic primitive. It does not propose alternative signature schemes, hashing functions, zero-knowledge constructions, or proof systems. Existing cryptographic tools are assumed to be sufficient for the guarantees they are designed to provide. No claim is made that cryptographic innovation alone would resolve the limitations discussed here.

Second, this work is not a proposal for a new consensus mechanism. It does not seek to redefine how distributed systems reach agreement on a shared state, nor does it aim to improve throughput, latency, finality, or adversarial tolerance. Consensus is treated as a solved, or at least well-scoped problem relative to the questions of process and legitimacy addressed in this paper.

Third, this paper is not an economic or financial model. It does not articulate incentives, token distributions, yield mechanisms, or value capture strategies. It makes no claims regarding returns, efficiency of markets, or optimal incentive alignment. Economic considerations appear only insofar as they intersect with governance or institutional behavior.

Fourth, this work should not be read as a comprehensive solution to decentralized governance. The mechanisms discussed do not eliminate conflict, power asymmetries, or the need for human judgment. They are not intended to replace deliberation, negotiation, or contextual decision-making. Instead, they aim to constrain specific procedural dimensions where arbitrariness and opportunistic behavior tend to concentrate.

Fifth, this paper is not a technical specification, implementation guide, or deployment plan. It does not define interfaces, protocols, or system parameters. Any reference to NOORCHAIN serves an illustrative purpose only and should not be interpreted as a finalized design or a commitment to a particular implementation path.

Finally, this document is not a manifesto. It does not argue for a single correct architecture, nor does it prescribe a universal direction for blockchain systems as a whole. Its objective is more limited and more exploratory: to propose a shift in framing, to articulate a set of constraints that are often implicit or neglected, and to open a space for further research.

By delineating these boundaries, the paper positions itself as a conceptual contribution rather than a prescriptive one. It seeks to refine the questions that blockchain systems ask of themselves, rather than to prematurely fix the answers.

9. Open Questions and Research Directions

The framework outlined in this paper is intentionally incomplete. Its purpose is to identify a structural limitation in prevailing blockchain architectures and to suggest a direction in which that limitation might be addressed. It does not offer a closed theory or a finalized design. As such, several foundational questions remain open and warrant systematic investigation.

A first set of questions concerns the formalization of trajectories. While the notion of a time-constrained process is central to the argument, it is difficult to formalize without collapsing it into a mere sequence of states. Which temporal constraints can be expressed and verified without requiring exhaustive observation? How can a system distinguish a compliant trajectory from an opportunistic one when both lead to the same final state? These questions point to the limits of what distributed systems can attest without introducing unrealistic assumptions about observability or enforcement.

A second line of inquiry relates to non-compressibility. The claim that certain processes should not be accelerable is conceptually clear but operationally subtle. Which forms of acceleration meaningfully undermine institutional significance, and which constitute acceptable efficiency gains? At what point does optimization erase the informational



content carried by duration? Addressing these questions requires a dialogue between systems engineering, institutional economics, and social theory.

A third area involves partial and indirect auditability. Human processes cannot, and should not, be rendered fully transparent. The challenge lies in identifying which minimal traces, signals, or artifacts are sufficient to support ex post verification of procedural constraints without exposing participants to undue surveillance or rigidity. This problem spans technical design, ethical considerations, and institutional norms.

Relatedly, the interaction between on-chain and off-chain components remains largely unexplored in this context. Process-oriented mechanisms inevitably unfold partly outside the chain. How can off-chain dynamics be linked to on-chain guarantees without reintroducing opaque intermediaries or centralized authorities? What forms of coupling preserve accountability while respecting the limits of formalization?

Governance itself raises further questions under a process-based lens. If certain procedural aspects are rendered non-discretionary, how can systems remain adaptable? How can rules governing processes evolve without invalidating the trajectories built under earlier regimes? This tension between stability and adaptability is familiar in institutional design, but it takes on new dimensions in distributed systems where rules are often embedded directly in code.

Finally, there is the broader question of scope. Not all systems benefit from strong temporal constraints or trajectory-based validation. Identifying the domains in which process-based trust provides clear institutional value distinct from those where it adds unnecessary complexity is a critical task for future work. Without such discrimination, the framework risks being applied where its costs outweigh its benefits.

These questions cannot be resolved by a single protocol or discipline. They call for interdisciplinary research at the intersection of distributed systems, applied cryptography, institutional economics, and social sciences. The contribution of this paper is to articulate a shared vocabulary and set of concerns that can guide that research, rather than to close it.

10. Conclusion , Beyond State

This paper began from a simple observation: contemporary blockchain systems are highly effective at securing states, yet they encounter structural limits when asked to support forms of legitimacy that depend on time, continuity, and adherence to process. This limitation is neither accidental nor the result of insufficient engineering. It follows directly from a conceptual framework in which truth and security are primarily defined in terms of verifiable end states.



By distinguishing explicitly between the validity of a state and the legitimacy of a trajectory, the paper argues for a shift in perspective. The issue is not that existing systems fail to function as designed, but that certain human and institutional properties commitment, persistence, and non-accelerability cannot be adequately represented when verification is confined to instantaneous outcomes.

The notion of process-based trust introduced here is intended as a framing device rather than a finished theory. It suggests that, in specific domains, trust can only emerge when time is treated as a first-class constraint and when the path taken matters independently of the result produced. Under this view, duration is not incidental; it carries information that cannot be replicated by point-in-time proofs alone.

This perspective does not call for abandoning established tools. Cryptography, smart contracts, and consensus mechanisms remain central to the security of distributed systems. Their role, however, changes when they are used to support the attestability of processes rather than the validation of states alone. In that context, they contribute to making certain forms of behavior particularly those involving continuity and sustained engagement observable and contestable over time.

The discussion of NOORCHAIN as an early instantiation illustrates both the potential and the constraints of this approach. It shows that formalizing human processes does not aim at exhaustiveness, but at constraining critical dimensions most notably time, persistence, and the absence of strategic shortcuts that escape purely state-based models. At the same time, it underscores that such systems remain experimental and that many design questions are unresolved.

More broadly, the contribution of this paper is not to propose a universal solution, but to clarify a boundary. Blockchains are powerful instruments for producing shared, tamper-resistant states. They are less naturally equipped to secure trajectories that derive their meaning from duration and procedure. Acknowledging this boundary is a prerequisite for any serious attempt to extend distributed systems into institutional and social domains.

Moving beyond state does not imply rejecting existing architectures. It requires recognizing their limits and resisting the temptation to obscure those limits through increasing complexity. By taking the role of process and time seriously, future systems may develop forms of coordination that are better aligned with the realities they seek to support without claiming to replace human judgment or to eliminate institutional uncertainty.

In this sense, the argument advanced here is deliberately open-ended. It invites further work on what distributed systems can reasonably be expected to prove, what they should refrain from claiming, and how they might coexist with the irreducible procedural and temporal dimensions of human coordination.