Name: Permissionless consensus from a common random string explainer

Category: Consensus

Paper title: '*Permissionless Consensus from a Common Random String'*

Authors: Damiano Abram, Marshall Ball, Juan Garay, and Aggelos Kiayias

Paper: https://eprint.iacr.org/2026/1179.pdf

Published: August, 2026 at CRYPTO 2026

Research explainer: Permissionless consensus from a common random string

Permissionless consensus allows a group of participants to agree on a shared decision without knowing in advance who is taking part. There are no identity checks or membership lists, and no central authority. The only constraint is an upper bound on the number of participants. This is a form of Byzantine agreement, in which parties reach agreement even if some are dishonest. Bitcoin is the best-known example of a system that uses this approach to reach consensus.

Since Bitcoin's introduction, researchers have studied how to achieve this kind of agreement securely, even when an adversary controls some of the participants' computational power using idealized models of computation. One common such example is the *random oracle model*, which models a hash function as an idealized random function. Another example is the availability of a randomness beacon which serves as a trusted source that reveals a fresh random value to all parties running the protocol.

A 2025 result by Khovratovich, Rothblum, and Soukhanov raised concerns about the security of several systems that rely on this model. This does not establish a flaw in Bitcoin or other permissionless consensus protocols, but it highlights why researchers want to establish security under more concrete assumptions.

What the research unlocks

The open question is simple: can permissionless consensus be achieved without relying on the random oracle model or requiring a randomness beacon? The authors show that it can. Building on earlier work that established a route to permissionless consensus with a shared randomness beacon, they develop a protocol that removes this requirement. Instead, their construction uses a common random string and two new cryptographic tools.

Without a beacon, it becomes harder to ensure that computational puzzles are unpredictable when participants receive them. A dishonest participant might prepare in advance, gaining an advantage over others. The paper introduces two tools to address this challenge:

New cryptographic tools

The first tool is a **d-wise independent distributed sampler**. It allows participants to generate random values collaboratively while providing security guarantees even when some participants are dishonest. The construction is designed so that the security of the generated values does not depend on every participant being honest.

The second tool is a **multi-verifier signature of work (MV-SoW)**. It allows a participant to demonstrate that they have performed a certain amount of computational work. Multiple verifiers can check the proof independently, without relying on a shared public-key infrastructure (PKI).

The proof is designed to make computational effort costly to fake or multiply: producing more valid proofs requires proportionally more work.

From building blocks to consensus

The authors combine these tools to construct a complete permissionless consensus protocol. Each participant can use a multi-verifier signature of work to establish their participation without relying on a central authority to distribute or confirm identities. Because participants may accept different proofs, each honest participant can end up with a different view of who is taking part.

The protocol builds on an existing agreement method, adapted to work with these overlapping views of participation. Under the paper's assumptions, it tolerates an adversary controlling up to approximately one-quarter of the participants. It does this without requiring a randomness beacon or relying on the random oracle model.

A matching limit

The authors also establish a limit for the protocol class they study. They prove that these protocols cannot tolerate an adversary controlling one-quarter or more of the participants. Improving on this threshold would require a different approach to determining which participants the protocol recognizes, rather than simply improving the cryptographic tools introduced in the paper

Why this research matters

This is foundational research, not a ready-to-use protocol. The authors demonstrate that permissionless consensus can be constructed without relying on the random oracle model or a randomness beacon, under specific cryptographic and computational assumptions. This gives researchers a foundation for exploring consensus protocols with security guarantees based on explicit assumptions.

The distributed sampler and multi-verifier signature of work may also prove useful for problems beyond consensus.

Learn more

Read the paper '*Permissionless Consensus from a Common Random String*' by Damiano Abram, Marshall Ball, Juan Garay, and Aggelos Kiayias for the full construction and security analysis.