Lorenzo Naranjo November 11, 2026

Proof of Stake

Proof of work secures a blockchain by making it expensive to add blocks: miners burn electricity, and an attacker would have to burn more. Proof of stake replaces the electricity with collateral. Participants lock up the blockchain's own coins to earn the right to add blocks, and they lose part of that collateral if they cheat. Ethereum, the second-largest blockchain, switched from proof of work to proof of stake on September 15, 2022, in an upgrade known as the Merge, which the Ethereum Foundation estimates cut the network's energy use by about 99.95% (Ethereum Foundation 2025). This chapter explains how proof of stake works, why validators behave honestly, and how its security compares with proof of work.

How proof of stake works

In proof of stake, the participants who maintain the chain are called *validators*. To become one, a participant deposits coins into the protocol, its *stake*; on Ethereum the minimum is 32 ether. Time is divided into slots, 12 seconds each on Ethereum. In every slot the protocol picks one validator to propose the next block, with probability proportional to stake, and asks a committee of other validators to *attest* that the block is valid, by signing a vote for it.

Blocks are linked by hashes exactly as in proof of work, so tampering with history is still visible to anyone who checks. What changes is how the network chooses between competing versions of the chain. Instead of following the chain with the most computational work, nodes follow the chain with the most attestations, weighted by stake. An attacker who controls a share q of the stake adds weight to their version at rate q while honest validators add weight to theirs at rate $1 - q$, so, as with computing power in proof of work, an attacker needs a majority to overtake the honest chain.

Validators earn rewards, paid in newly created coins and in fees, for proposing and attesting. They are *slashed*, losing part of their stake, if they sign two conflicting votes, for example by proposing two different blocks for the same slot or by voting for two

different checkpoints at the same point in the chain. Such double voting is provable: the two signed messages are the evidence, and anyone can submit them. On Ethereum, a validator slashed on its own loses a small part of its stake, but the penalty grows with the share of all stake slashed at about the same time, and reaches the whole stake when a third or more is slashed together. A lone mistake is cheap; a coordinated attack is ruinous.

Why validators behave

Why would validators not cheat when they choose which blocks to approve? In proof of work, an attacker's cost is the electricity and hardware needed to outpace honest miners. In proof of stake, adding blocks to a competing chain costs almost nothing in computation. Early critics called this the *nothing at stake* problem and saw it as a fatal flaw: a validator might as well vote for every competing chain at once.

Saleh (2021) answers with economics. A validator holds the coin it stakes. If validators fail to agree on a single chain, users lose confidence and the coin loses value, so validators lose too. Saleh shows that "a sufficiently modest reward schedule" leads validators to reach consensus quickly and rules out forks that persist, the latter "because PoS, unlike PoW, requires that validators are also stakeholders." Ethereum adds a second, explicit deterrent: slashing the stake of validators who sign conflicting votes.

Finality

Proof of work never makes a block completely final: a reversal becomes less and less likely as confirmations accumulate, but it is never impossible. Ethereum's proof of stake adds *finality*. Every 32 slots, about every 6.4 minutes, the chain has a *checkpoint*. If validators holding at least two-thirds of all stake attest to a checkpoint, and then to the next one, it becomes finalized, which normally takes about 13 minutes.

Buterin and Griffith (2017) prove the key property, called accountable safety: "two conflicting checkpoints cannot both be finalized unless $\geq 1/3$ of validators violate a slashing condition (meaning at least one third of the total deposit is lost)." To reverse a finalized block, then, at least a third of all staked coins must break the rules in a way that can be punished, so the attack puts the value of a third of the stake at risk.

The guarantee is only as good as the punishment. Budish (2025) points out that "a large-enough attacker (specifically, with at least [two thirds of the stake]) can always withdraw all of the capital that signed conflicting transactions before the confirmation of any transactions that slash their stake," because such an attacker controls which transactions the chain confirms. An attacker with more than a third of the stake, too few to reverse anything, "can effectively halt transactions for a significant period of time" by refusing to vote. His reading is that proof of stake "can successfully mimic traditional trust to deter sub-2/3 attackers, but that rule of law must step in in the case of a large-enough attacker."

The economics of staking

Validators stake only if it pays. Let R be the total rewards the protocol pays per year and S the total amount staked. To keep the algebra simple, take R as fixed; on Ethereum it actually grows with the square root of S, which does not change the conclusions below. Rewards are shared in proportion to stake, so every staked coin earns the same gross rate R/S. Net of the costs of running a validator and the risk of penalties, a fraction δ of the stake per year, the staking yield is

$$y = \frac{R}{S} - \delta. \tag{1}$$

The more is staked, the lower the yield, because the same rewards are shared among more coins.

Stake flows in as long as staking pays more than the best alternative use of the same coins, which earn ρ there, and flows out when it pays less. Staked or not, the coins carry the same price risk, so ρ is a return in coins, for example from lending them out, not a dollar interest rate. In equilibrium the gross rate covers both, $R/S = \rho + \delta$, so

$$S = \frac{R}{\rho + \delta}. \tag{2}$$

Higher rewards attract more stake, and with finality, more stake means more capital that an attacker must be willing to lose.

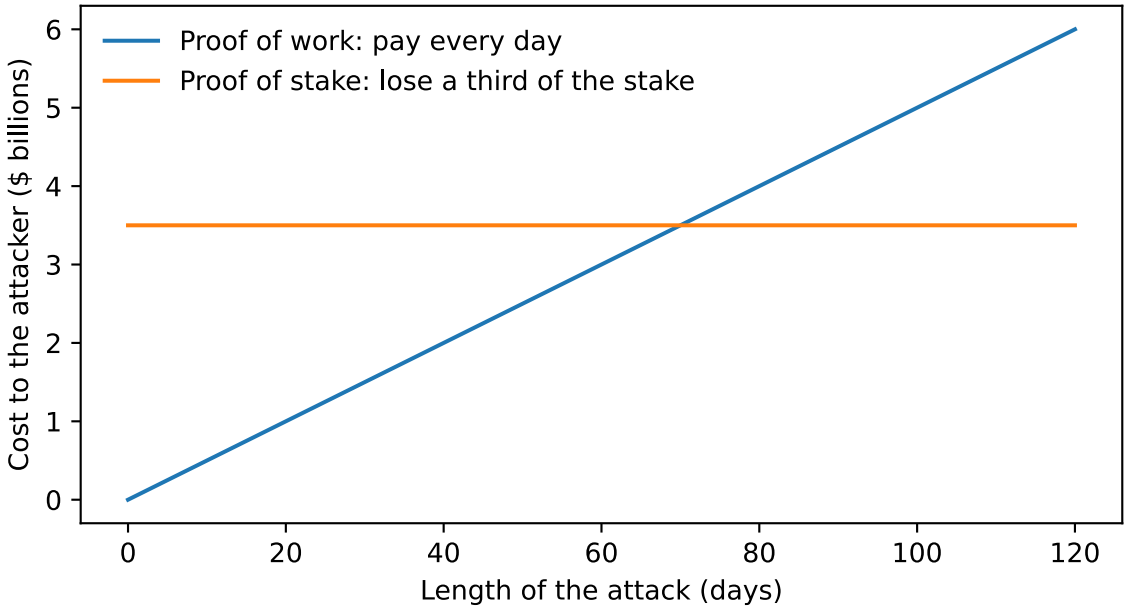


Figure 1: Stylized cost of an attack that lasts a given number of days. Under proof of work the attacker pays for computing power every day; under proof of stake the attacker loses one third of the stake once. The numbers are illustrative.

Proof of work and proof of stake compared

The two systems make attacks expensive in different ways. In proof of work, an attacker pays for computing power for as long as the attack lasts: the cost is a *flow*, which grows with the length of the attack. In proof of stake, the cost of reversing a finalized block is a *stock*: at least one third of the stake, lost all at once, however long the attack takes.

Figure 1 illustrates the difference with made-up numbers. Which system costs an attacker more depends on the numbers and on how long an attack has to last. John et al. (2025) build an equilibrium model of both systems and, applying real-world parameter values, conclude that proof-of-stake blockchains are more secure than otherwise equivalent proof-of-work blockchains, with the advantage largest for blockchains of large scale.

Proof of stake has its own risks. Stake can concentrate in a few large staking services and exchanges, which stake on behalf of many users. If a few operators control a third or more of the stake between them, finality depends on their behavior. And the coins that

secure the chain are the chain's own coins: if their price falls, so does the value of the stake an attacker would lose.

Practice Problems

These problems give you a chance to practice the concepts introduced in this chapter. Try to solve each one on your own before expanding the solution.

Solutions to all problems can be found at https://lorenzonaranjo.com/class-materials/fintech/.

Problem 1 (Staking Yield). A proof-of-stake protocol pays rewards of 1 million coins per year. Running a validator and the risk of penalties cost 0.5% of the stake per year. Compute the staking yield when 20 million coins are staked and when 40 million are staked. If the alternative use of capital earns 3% per year, how much will be staked in equilibrium?

Problem 2 (The Price of Reversing a Finalized Block). A proof-of-stake blockchain has 30 million coins staked, and each coin trades at $3,000. What is the minimum value an attacker must lose to reverse a finalized block? What happens to that number if the coin's price falls by half?

Problem 3 (Flow versus Stock). Select all true statements.

a. In proof of work, the cost of an attack grows with how long the attack lasts.
b. In proof of stake, adding blocks to a competing chain requires as much electricity as in proof of work.
c. Under Ethereum's finality rule, reversing a finalized block requires validators holding at least a third of the stake to violate a slashing condition.
d. In proof of stake, validators who double vote can be identified and lose part of their stake.

References

Budish, Eric. 2025. “Trust at Scale: The Economic Limits of Cryptocurrencies and Blockchains.” *Quarterly Journal of Economics* 140 (1): 1–62. https://doi.org/10.1093/qje/qjae033.

Buterin, Vitalik, and Virgil Griffith. 2017. “Casper the Friendly Finality Gadget.”. https://arxiv.org/abs/1710.09437.

Ethereum Foundation. 2025. “The Merge.”. https://ethereum.org/roadmap/merge/.

John, Kose, Thomas J. Rivera, and Fahad Saleh. 2025. “Proof-of-Work Versus Proof-of-Stake: A Comparative Economic Analysis.” *Review of Financial Studies* 38 (7): 1955–2004. https://doi.org/10.1093/rfs/hhaf013.

Saleh, Fahad. 2021. “Blockchain Without Waste: Proof-of-Stake.” *Review of Financial Studies* 34 (3): 1156–90.