Blockchain Economics and Digital Assets

Lecture 2: Blockchain Economics

Dr Daniele Bianchi

Queen Mary, University of London

Semester B, 2025/2026

Contents

Overview

In Topic 1, we established that blockchain is a distributed ledger maintained by a network of computers without central control. But we left a crucial question unanswered: how do these computers agree on what the ledger contains? When thousands of strangers around the world maintain copies of the same database, and some of them might be dishonest, how do they reach agreement?

This is the **consensus problem**, and solving it is what makes blockchain possible. The solution is not merely technical—it is fundamentally economic. Blockchain networks align incentives so that honest behaviour is profitable and cheating is costly. Understanding these incentive structures is essential for understanding both the value and limitations of blockchain technology.

This lecture examines the two dominant approaches to consensus: Proof-of-Work (PoW) and Proof-of-Stake (PoS). We explore the economic logic underlying each, their trade-offs in terms of security and resource consumption, and the fundamental constraints that limit all blockchain designs. We conclude by examining Layer 2 solutions—attempts to scale blockchain without sacrificing its core properties.

1 The Consensus Problem

1.1 Why Consensus Matters

Recall from Topic 1 that a blockchain is a shared ledger with no central authority. This immediately raises a fundamental question:

> How do strangers agree on the state of a shared database when some of them might be lying?

In traditional systems, a central authority—a bank, a clearinghouse, a government registry—maintains the authoritative record. Disputes are resolved by reference to this authority. But blockchain deliberately eliminates this central point of control. The consensus protocol must somehow ensure that all honest participants agree on the same transaction history, even in the presence of adversarial actors.

1.2 The Double-Spend Problem

The challenge is particularly acute for digital money. Physical cash has a natural scarcity: you cannot hand the same banknote to two people. Digital information, by contrast, can be copied perfectly.

Consider this scenario: Alice has 1 BTC. She simultaneously sends it to Bob (Transaction A) and to Carol (Transaction B). Both transactions are valid copies—each correctly signed with Alice's private key, each referencing the same unspent output. Which transaction counts?

With physical cash, this is impossible by the laws of physics. With digital cash, both transactions appear equally legitimate. **Someone must decide which one to accept**, and that decision must be binding on all participants. Without such a mechanism, digital currency is worthless—any recipient might discover their payment was already spent elsewhere.

Traditional payment systems solve this through centralised ledgers. Your bank maintains the authoritative record of your balance; it simply rejects the second transaction. Blockchain's innovation is solving the double-spend problem through decentralised consensus.

1.3 Byzantine Fault Tolerance

The consensus problem is a modern instance of what computer scientists call the **Byzantine Generals Problem**. Imagine generals surrounding a city who must coordinate an attack. They can only communicate by messenger, and some generals may be traitors who send conflicting messages. How can the loyal generals reach agreement?

The blockchain parallel is direct. Thousands of nodes must agree on transaction history. Some nodes might be offline, malfunctioning, or actively malicious. The network must produce a single consistent ledger despite these failures. A consensus protocol that works correctly even when some participants are adversarial is called **Byzantine fault tolerant**.

The theoretical results are sobering: Byzantine agreement requires that fewer than one-third of participants be malicious in most settings. Blockchain's contribution is creating systems where participating honestly is economically rational, so that the one-third threshold is rarely approached.

1.4 Properties of Good Consensus

A well-designed consensus protocol should deliver:

Property	Meaning
Safety	Honest nodes agree on the same transaction history
Fault tolerance	The system continues operating despite some nodes failing or attacking
Decentralisation	No single party can control which transactions are confirmed
Efficiency	Resources are not wasted unnecessarily

These properties often conflict. Greater security may require more computational work. More decentralisation may slow consensus. The design of a consensus mechanism embodies choices about which trade-offs to accept.

2 Proof-of-Work

Before Bitcoin (2009), every attempt at decentralised digital cash had failed because of the double-spend problem. Satoshi Nakamoto's insight was to make adding blocks to the chain **computationally expensive**.

The basic idea is simple. Participants called "miners" compete to solve a cryptographic puzzle. The winner gets to propose the next block and receives a reward—newly minted coins plus transaction fees. Other nodes verify the solution (which is trivial) and append the valid block to their chains.

This is **Proof-of-Work**: to add a block, you must prove that you expended computational resources. The name is apt—the computational work is the proof.

Why does this solve the double-spend problem? Because attacking the network requires outcomputing all honest miners. An attacker who wants to reverse a confirmed transaction must create an alternative version of history that is longer (contains more accumulated work) than the honest chain. If honest miners control the majority of computational power, the honest chain grows faster than any attacker's chain, making successful attacks exponentially less likely as more blocks are added.

2.1 How Mining Works

The puzzle miners solve is finding a number (called a "nonce") that, when hashed together with the block's data, produces an output below a certain threshold. Since hash outputs are unpredictable, there is no shortcut—miners must try billions of random values until one works.

The process works as follows:

1. Collect pending transactions from the mempool into a candidate block
2. Try random nonces until finding one that produces a valid hash
3. Broadcast the winning block to the network
4. Other nodes verify the solution (one hash computation) and append the block
5. The winner receives the block reward plus all transaction fees in the block

The asymmetry is crucial: finding a valid nonce requires billions of attempts; verifying requires exactly one. This asymmetry makes the system work—miners invest heavily in security, but anyone can cheaply verify their work.

2.2 Difficulty Adjustment

Bitcoin targets one block approximately every 10 minutes. But if more miners join the network, blocks would be found faster, potentially destabilising the system.

The solution is **difficulty adjustment**. Every 2,016 blocks (roughly two weeks), the protocol automatically adjusts the difficulty target based on how quickly the previous blocks were found. If blocks came too fast, difficulty increases; if too slow, it decreases.

This creates an **arms race**. Adding more computational power improves your *relative* chance of winning the next block, but it does not increase the total rate of block production. The difficulty simply adjusts to absorb the additional capacity. More hashpower means more security for the network—an attacker needs to overcome more computational power—but no more blocks per hour.

2.3 Block Rewards and the Halving

Miners receive two types of compensation:

1. **Block subsidy**: Newly created bitcoins (currently 3.125 BTC per block)
2. **Transaction fees**: Payments from users seeking transaction inclusion

The block subsidy follows a predetermined schedule called the **halving**. Every 210,000 blocks (approximately four years), the subsidy is cut in half:

Date	Block Subsidy	Event
January 2009	50 BTC	Genesis
November 2012	25 BTC	1st halving
July 2016	12.5 BTC	2nd halving
May 2020	6.25 BTC	3rd halving
April 2024	3.125 BTC	4th halving
∼2028	1.5625 BTC	5th halving (projected)

This creates a **deflationary supply schedule**: the total supply asymptotically approaches 21 million BTC, reached around 2140. The economic implications of this fixed supply—and whether

transaction fees alone will provide sufficient security incentives as block subsidies decline—remain actively debated.

2.4 The Longest Chain Rule

When two miners find valid blocks nearly simultaneously, the network temporarily has two competing versions of the chain—a "fork." Nakamoto's solution is the **longest chain rule**: nodes follow whichever chain has the most accumulated proof-of-work.

Miners choose which branch to extend. Eventually, one branch gets ahead, and the other is abandoned ("orphaned"). Transactions in orphaned blocks return to the mempool for re-inclusion in the winning chain.

This mechanism handles benign forks from network latency. More importantly, it makes attacks difficult. An attacker trying to reverse a confirmed transaction must build an alternative chain faster than all honest miners combined. If honest miners control more than 50% of computational power, the probability of success decreases exponentially with each additional block built on the honest chain.

This is why users are advised to wait for multiple **confirmations** before considering a transaction final. The Bitcoin convention is 6 confirmations (approximately one hour), though fewer may suffice for smaller amounts.

2.5 The Security Model

Proof-of-Work protects against specific attacks:

- **Double-spending**: Would require rewriting confirmed blocks, which becomes exponentially harder over time
- **Censorship**: Any miner can include any valid transaction; no single miner controls the mempool
- **Counterfeit coins**: Nodes independently verify all blocks; invalid transactions are rejected regardless of who proposes them

The **51% attack** represents the theoretical vulnerability. An attacker controlling majority hashpower could double-spend their own transactions, prevent specific transactions from confirming, and generally disrupt the network. However, they still could not steal others' coins (private keys remain required) or create coins beyond the protocol rules (other nodes would reject such blocks).

The cost of such an attack on Bitcoin is estimated in the billions of dollars for hardware and electricity, making it economically irrational against a well-functioning network. Smaller proof-of-work chains with less hashpower are more vulnerable—several have suffered successful 51% attacks.

3 Proof-of-Stake

3.1 Beyond Proof-of-Work

Proof-of-Work achieves security but at significant cost:

- **Energy consumption**: Bitcoin uses an estimated 100–150 TWh annually, comparable to some countries
- **Hardware waste**: Mining equipment becomes obsolete rapidly

- **Centralisation pressure**: Economies of scale favour large operations with access to cheap electricity

These costs prompted a fundamental question: could we achieve similar security without the energy expenditure? **Proof-of-Stake** offers an alternative approach: instead of proving you expended computational resources, you prove you have "skin in the game" by locking up cryptocurrency as collateral.

3.2 How Proof-of-Stake Works

The basic mechanism:

1. Validators deposit ("stake") cryptocurrency as collateral
2. The protocol randomly selects validators to propose blocks, with selection probability proportional to stake size
3. Other validators "attest" (vote) that the proposed block is valid
4. Validators earn rewards for honest participation

The enforcement mechanism is **slashing**: if a validator misbehaves—for example, by proposing conflicting blocks or making false attestations—a portion of their stake is automatically destroyed. This makes attacks economically painful without requiring ongoing energy expenditure.

The key insight is that security comes from *economic penalties* rather than *computational cost*. An attacker must risk losing their staked capital, not merely wasting electricity.

3.3 The Ethereum Merge

Ethereum's transition from Proof-of-Work to Proof-of-Stake in September 2022—known as "The Merge"—represents the largest consensus mechanism change in blockchain history.

Before (Proof-of-Work):

- Miners competed with specialised hardware (GPUs, then ASICs)
- Energy consumption: approximately 80–100 TWh annually
- Block time: variable, averaging 13 seconds

After (Proof-of-Stake):

- Validators stake 32 ETH (roughly $100,000 at current prices)
- Energy consumption: approximately 0.01 TWh annually—a **99.95% reduction**
- Block time: fixed 12-second slots

The Merge demonstrated that a major blockchain can successfully change consensus mechanisms, though the transition required years of development and careful coordination.

3.4 Staking Economics

To become an Ethereum validator:

- Stake exactly 32 ETH (the minimum requirement)
- Run validator software continuously (offline validators face penalties)
- Earn rewards for proposing blocks and making correct attestations

- Current yield: approximately 3–5% APR, varying with network activity

Barriers to entry include the substantial capital requirement (32 ETH is roughly $100,000), technical complexity of running validator infrastructure, and historically, the inability to withdraw staked ETH (addressed in 2023).

Liquid staking protocols like Lido and Rocket Pool address these barriers by pooling stake from many users and issuing tradeable tokens representing the staked position. We examine these in Topic 4 on DeFi.

3.5 Proof-of-Work vs Proof-of-Stake

	Proof-of-Work	**Proof-of-Stake**
Security basis	Computational cost	Economic stake
Energy use	Very high	Minimal
Hardware needs	Specialised (ASICs)	Standard servers
Attack cost	Acquire/rent hashpower	Acquire and risk stake
Barrier to entry	Capital + expertise + electricity	Capital (stake requirement)
Primary example	Bitcoin	Ethereum (post-Merge)
Finality	Probabilistic	Can be deterministic

Neither approach is strictly superior. Bitcoin maximalists argue that Proof-of-Work's energy expenditure *is* the security—real-world resources that cannot be faked or borrowed cheaply. Proof-of-Stake advocates argue that equivalent security can be achieved more efficiently by making capital, rather than electricity, the thing at risk.

Proof-of-Stake is not without criticism:

"Rich get richer": Validators with more stake earn proportionally more rewards, potentially increasing concentration over time. Unlike mining, where new entrants can purchase hashpower, staking rewards accrue to existing stake.

Nothing-at-stake problem: In theory, validators could costlessly vote for multiple competing chain forks (since voting doesn't consume resources like mining does). This is addressed through slashing, but adds complexity.

Long-range attacks: An attacker with old validator keys could theoretically rewrite history from an early point. This is mitigated through checkpointing and "weak subjectivity"—new nodes must get a recent trusted checkpoint, not just genesis.

Validator centralisation: Large staking pools (Lido controls approximately 30% of staked ETH) raise concerns similar to mining pool concentration.

The long-term security properties of Proof-of-Stake at scale are still being validated in production. Ethereum's experience since the Merge provides the most significant real-world test.

4 The Blockchain Trilemma

4.1 The Trade-off

Vitalik Buterin articulated a fundamental constraint on blockchain design:

> A blockchain can optimise for at most two of three properties: **decentralisation**, **security**, and **scalability**.

This is not a proven theorem but an observed engineering trade-off that manifests across different design choices.

4.2 Understanding the Tensions

Decentralisation vs Scalability: More nodes mean more communication overhead. Every transaction must propagate to all validators; every block must be verified by everyone. Bitcoin's roughly 15,000 nodes can process about 7 transactions per second. Visa's centralised system handles 65,000 transactions per second at peak capacity.

Security vs Scalability: Faster blocks leave less time for propagation, creating more temporary forks and easier attack opportunities. Larger blocks process more transactions but require more storage and bandwidth, pricing out smaller validators and reducing decentralisation.

Security vs Decentralisation: Stronger security often requires more resources (hashpower or stake). Higher costs push participants toward pools to share variance, concentrating power. Both Bitcoin mining and Ethereum staking exhibit significant concentration.

4.3 Where Different Blockchains Sit

Blockchain	Decentralisation	Security	Scalability
Bitcoin	High	High	Low (∼7 TPS)
Ethereum L1	High	High	Low (∼15 TPS)
Solana	Lower	Medium	High (∼5,000 TPS)
BNB Chain	Low	Medium	High
Hyperledger	Permissioned	High (within trust model)	High

The question is not which design is "best" but which trade-off suits the application. Bitcoin prioritises security and decentralisation for a store of value. Solana sacrifices some decentralisation for throughput suitable for applications requiring many transactions.

5 Layer 2 Solutions

5.1 The Scaling Approach

If Layer 1 blockchains face fundamental scalability constraints, can we build additional layers on top that provide greater throughput while inheriting the underlying security?

> **Layer 2 solutions** process transactions off the main chain but inherit security from it by periodically posting proofs or summaries back to Layer 1.

The idea is to batch many transactions into a single Layer 1 transaction, amortising the cost and throughput constraints across many users.

5.2 Bitcoin's Lightning Network

The **Lightning Network** enables instant Bitcoin payments through payment channels:

- Two parties lock Bitcoin in a shared address (opening a channel)
- They can transact unlimited times between themselves by exchanging signed messages—no on-chain transactions required
- Either party can close the channel at any time, with final balances settled on-chain
- Channels can be chained: Alice pays Carol through Bob, even without a direct channel

Lightning enables micropayments (fractions of a cent) with instant settlement and near-zero fees. Capacity currently exceeds 5,000 BTC locked in channels, with growing adoption for everyday payments in jurisdictions like El Salvador.

5.3 Ethereum Rollups

Ethereum's scaling roadmap centres on **rollups**—Layer 2 systems that bundle hundreds of transactions and post compressed data to Ethereum:

Optimistic Rollups (Arbitrum, Optimism): Assume transactions are valid by default. Anyone can submit a "fraud proof" during a challenge period if they detect invalid transactions. If the fraud proof succeeds, the invalid batch is reverted and the malicious proposer is penalised.

ZK-Rollups (zkSync, StarkNet): Use cryptographic proofs (zero-knowledge proofs) to mathematically guarantee validity. More computationally intensive to generate but provide immediate finality without challenge periods.

Current state:

- Arbitrum and Optimism together hold over $10 billion in total value locked
- Fees: Often 10–100x cheaper than Ethereum mainnet
- Speed: Near-instant confirmation for users, with periodic settlement to L1

5.4 Trade-offs of Layer 2

Layer 2 does not "solve" the trilemma—it shifts the trade-offs:

- Users must trust Layer 2 operators to some degree (though fraud proofs or validity proofs limit this trust)
- Bridging between L1 and L2 introduces complexity and potential attack vectors
- Liquidity fragments across layers
- Different L2s may not interoperate smoothly

Nonetheless, Layer 2 represents the most practical near-term path to scaling blockchain for broader adoption, and understanding these systems is essential for understanding modern blockchain architecture.

6 Economic Implications

6.1 Energy Consumption

The energy debate around blockchain has evolved significantly:

Bitcoin (Proof-of-Work): Estimated 100–150 TWh annually, comparable to countries like Argentina or Norway. Critics point to the environmental impact and carbon footprint. Defenders note that Bitcoin increasingly uses renewable and otherwise-stranded energy, and that the energy expenditure secures a network worth over $1 trillion.

Ethereum (post-Merge Proof-of-Stake): Approximately 0.01 TWh annually—a 99.95% reduction from pre-Merge levels. This is comparable to a few thousand households.

The implication is that energy criticism now applies specifically to Proof-of-Work chains (primarily Bitcoin), not to blockchain technology generally. Whether Bitcoin's energy use is justified

depends on how one values Bitcoin's properties—a debate that extends beyond technical considerations into questions of monetary policy and financial sovereignty.

6.2 Mining and Staking Concentration

Despite decentralisation goals, both mining and staking exhibit concentration:

Bitcoin mining pools: The top 4 pools control approximately 70% of hashpower. Foundry USA, AntPool, F2Pool, and ViaBTC dominate, with geographic concentration in the US, China (despite the ban), and Kazakhstan.

Ethereum staking: Lido controls approximately 30% of staked ETH. Coinbase and Kraken hold another 15% combined. Independent validators represent a minority.

Why does concentration occur? Economies of scale in hardware and electricity for mining; pooling to reduce variance for both mining and staking; technical barriers that favour sophisticated operators.

Important nuance: pool operators do not directly control miners' or stakers' machines. Participants can switch pools, providing some check on pool behaviour. The concentration is concerning but not equivalent to the pool operators controlling the network.

6.3 Wealth Concentration

Cryptocurrency holdings are highly concentrated. Approximately 70% of Bitcoin addresses hold less than $1,000, while roughly 0.01% hold more than $10 million.

Caveats apply: one person can control many addresses; one address (such as an exchange) can represent many people; early adopters naturally hold more, having purchased at $1 rather than $60,000.

Whether this concentration undermines claims of "democratisation" is debated. It may simply reflect early-stage adoption patterns, or it may indicate structural features that favour early participants.

7 Summary and Looking Ahead

This lecture has covered the economics of blockchain consensus. The key takeaways:

Consensus is blockchain's core innovation. It solves the double-spend problem without trusted intermediaries, using Byzantine fault tolerant protocols that work even when some participants are adversarial.

Proof-of-Work achieves security through computational cost. Battle-tested since 2009, it is energy-intensive but provides robust security guarantees. The halving schedule creates a long-term transition toward fee-based security.

Proof-of-Stake achieves security through economic stake. Ethereum's Merge demonstrated a 99.95% energy reduction. Different trust assumptions and centralisation risks apply, and long-term security properties are still being validated.

The trilemma constrains all blockchain designs. Decentralisation, security, and scalability exist in tension. Layer 2 solutions shift rather than eliminate these trade-offs.

Concentration is real. Mining pools, staking pools, and wealth holdings all exhibit significant concentration, complicating the decentralisation narrative.

In the next lecture, we turn to smart contracts and decentralised applications—the programmable layer that transforms blockchain from a payment system into a platform for financial innovation.

Readings

Required:

- Nakamoto, S. (2008). "Bitcoin: A Peer-to-Peer Electronic Cash System." Sections 4–6 on Proof-of-Work, timestamps, and incentives.

Supplementary:

- Buterin, V. (2021). "Why Proof of Stake." Ethereum Foundation blog post explaining the rationale for the Merge.
- Budish, E. (2018). "The Economic Limits of Bitcoin and the Blockchain." NBER Working Paper 24717. A rigorous economic analysis of Proof-of-Work security.
- Saleh, F. (2021). "Blockchain Without Waste: Proof-of-Stake." *Review of Financial Studies*, 34(3), 1156–1190. Economic analysis of Proof-of-Stake incentives.