Do Blockchain Systems Achieve Decentralization?

Christina Ovezik

AGT@Blockchains Workshop - WINE 2024

Based on work with Aggelos Kiayias, Dimitris Karakostas, Daniel Woods

What is a blockchain?

- Digital ledger that satisfies a set of liveness and safety properties, “**without relying on trust**”
- Mechanism design perspective: a protocol with **decentralisation** as an objective

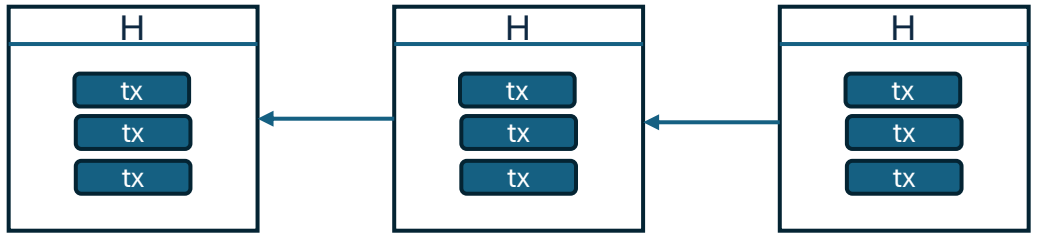

Centralised vs Decentralised systems

- A **single party** has full control over the system (“rules them all”)
- If it misbehaves or crashes, the system is compromised, e.g.:
  - Censorship
  - Inconsistency
  - Data loss / unavailability

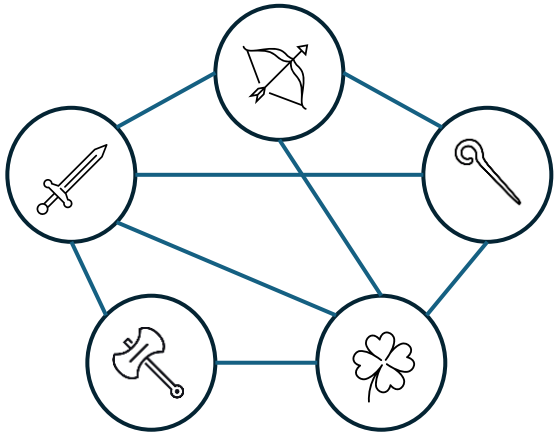

- A number of **independent parties** collectively control the system and guarantee its properties
- If only few nodes are malicious / faulty, the system remains operational
  - No single points of failure

Example of centralised system failure

BBC Sign in Home News Sport Weather iPlayer Sounds

NEWS

Home | InDepth | Israel-Gaza war | War in Ukraine | Climate | UK | World | Business | Politics | Culture

Technology

CrowdStrike IT outage affected 8.5 million Windows devices, Microsoft says

Joe Tidy
Cyber correspondent, BBC News

20 July 2024

Blockchains are multi-layered systems

- **8 layers of** blockchain systems where **(de)centralisation** can occur

- **Centralisation** in some layer **threatens** desirable system properties, such as **safety, liveness, privacy and price stability**

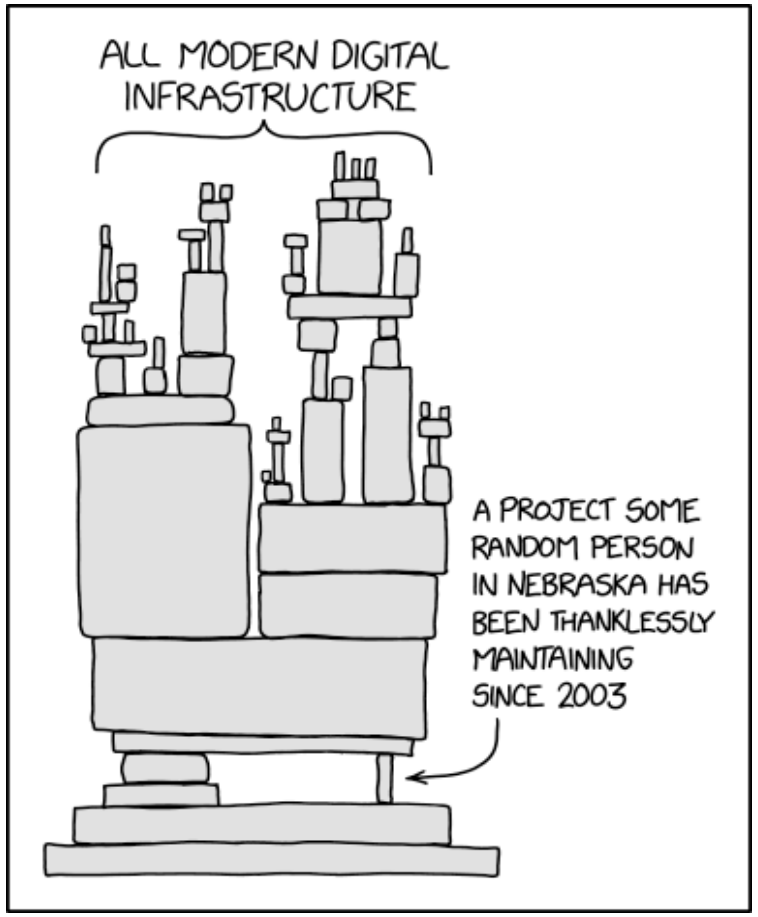

Image credit: https://xkcd.com/2347/

For each layer

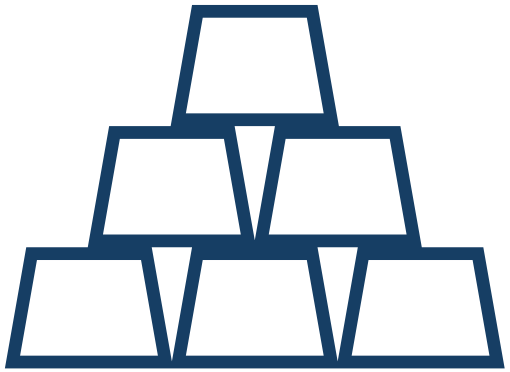

Resource

Relevant parties

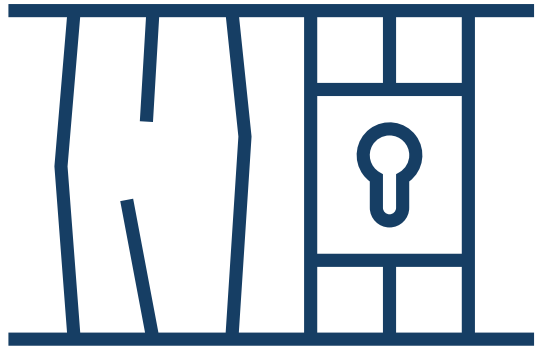

Properties at risk

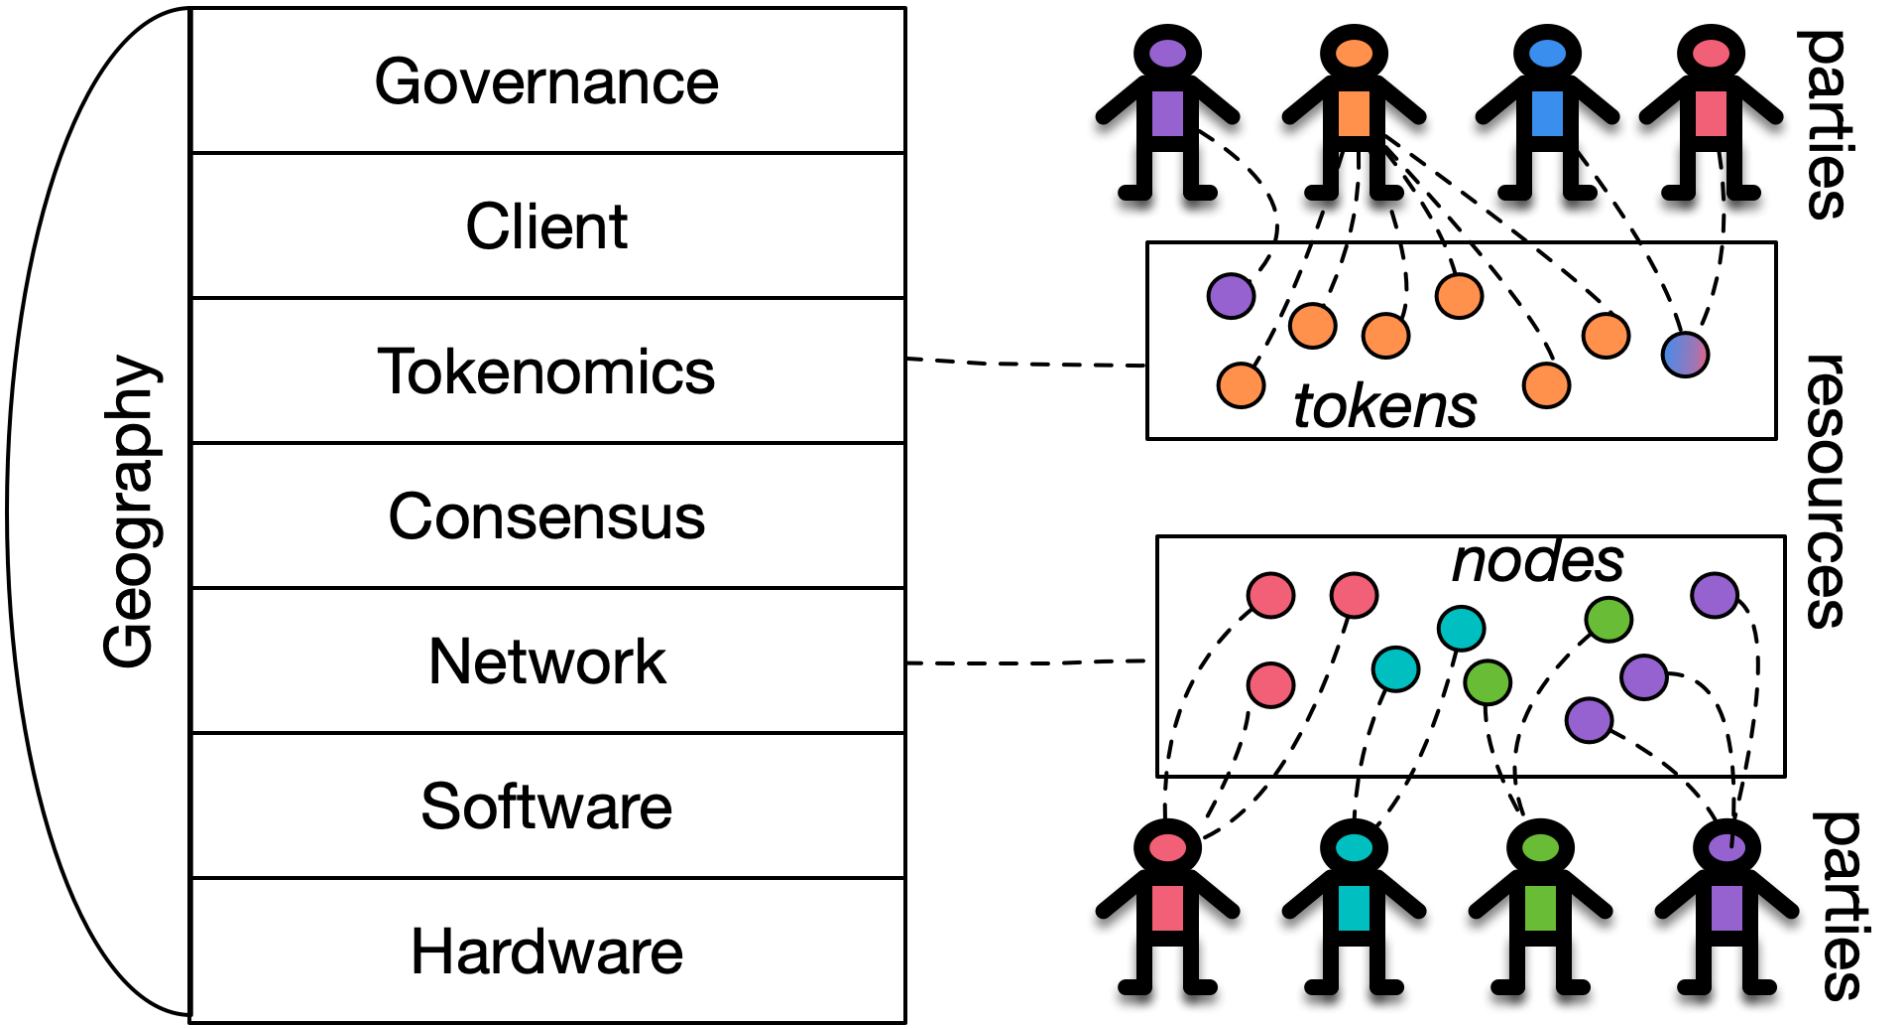
Governance
Client
Tokenomics
Consensus
Network
Software
Hardware
Geography
parties
tokens
resources
nodes
parties

Case study: consensus layer

Measuring the decentralisation of block production

Consensus layer: extending the blockchain

- Block creators:
  - Decide which transactions get included in a block
    - and in what order
  - Receive rewards for each new block
- The more blocks one creates the more influence they have in the system

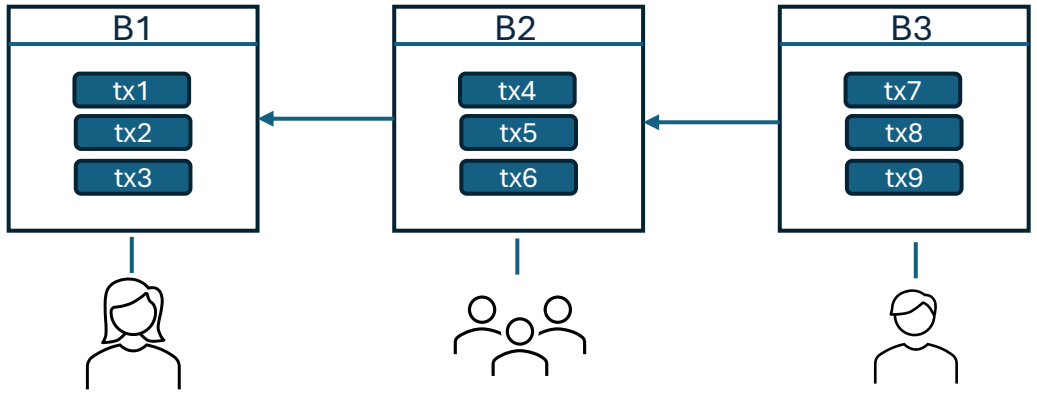

Consensus – Decentralization Analysis

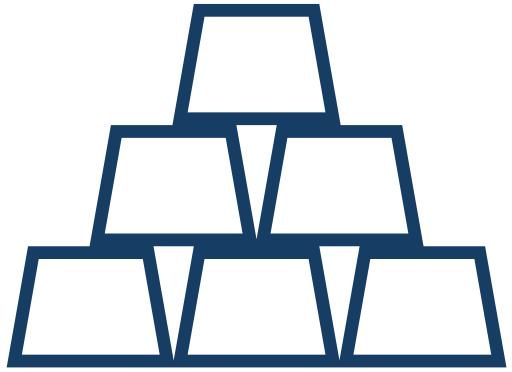

Resource
=
Blocks

Relevant parties
=
Block
producers

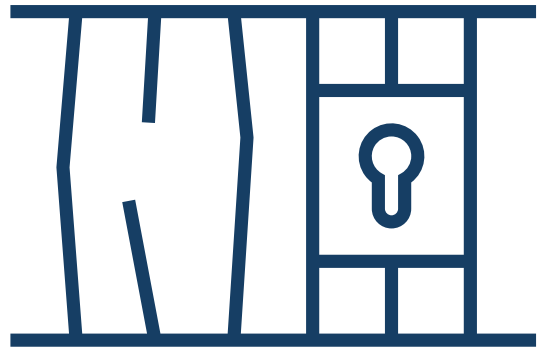

Properties
at risk
=
Safety / Liveness

Mining and the formation of coalitions

- Economies of scale incentivize the formation of coalitions (**mining / stake pools**) with their leaders consolidating disproportionate power
- We treat each such coalition as a single entity

Block production dynamics

Bitcoin

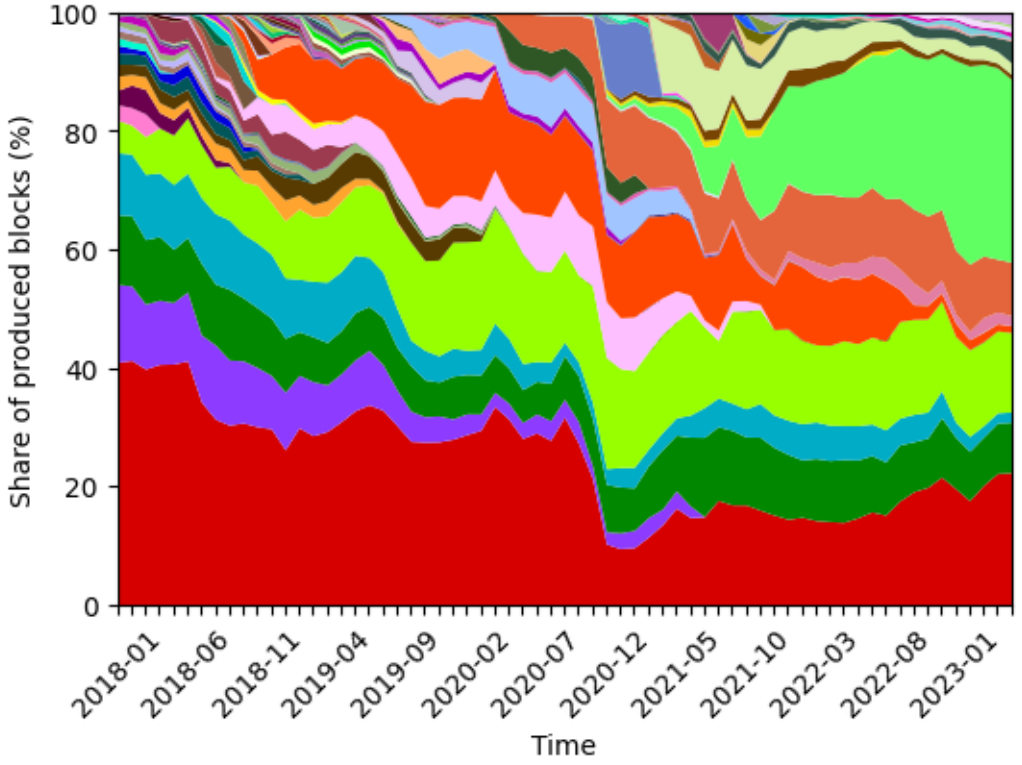

Ethereum

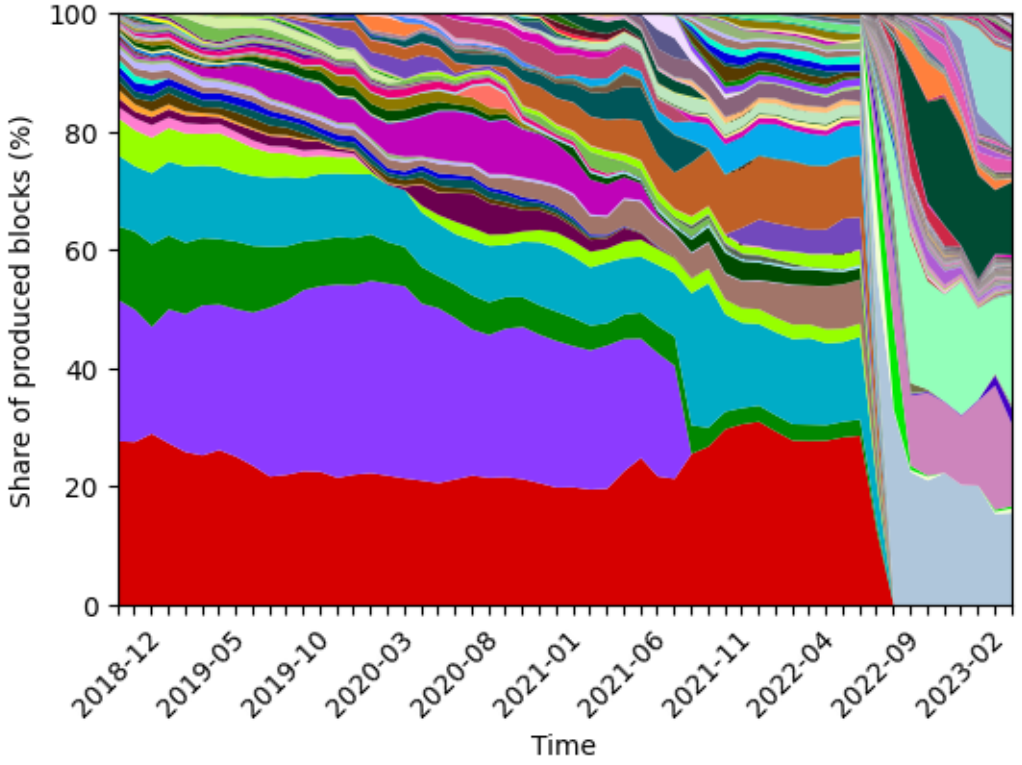

Block production dynamics

Cardano

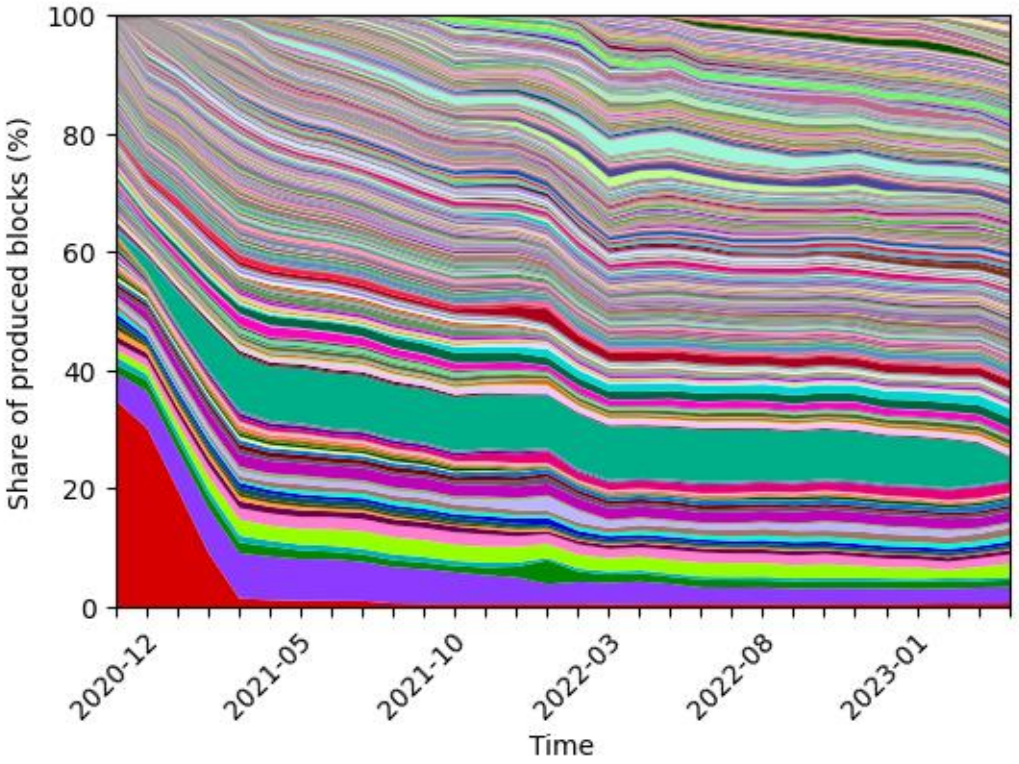

Tezos

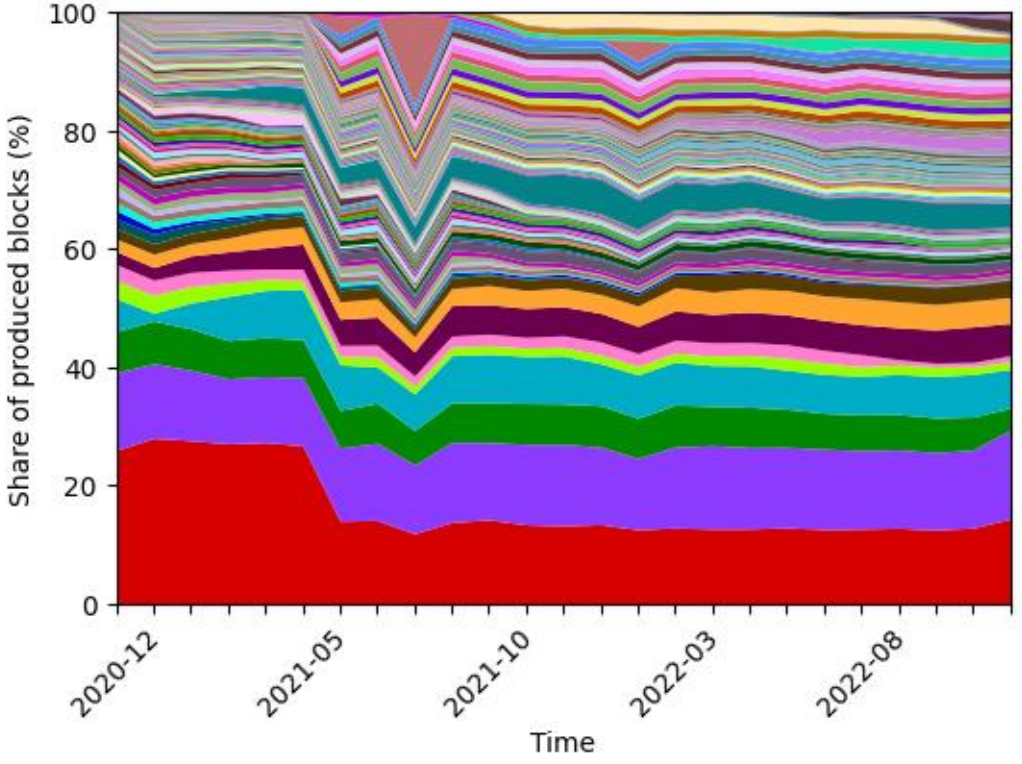

Decentralization metrics

- Assign a value that represents the decentralization of a distribution
- Metrics used in the blockchain decentralization literature:
  - Nakamoto coefficient (NC)
  - Gini coefficient (G)
  - Shannon entropy (SE)
  - Herfindahl–Hirschman index (HHI)
  - Concentration ratios (CR)
  - Number of parties (P)

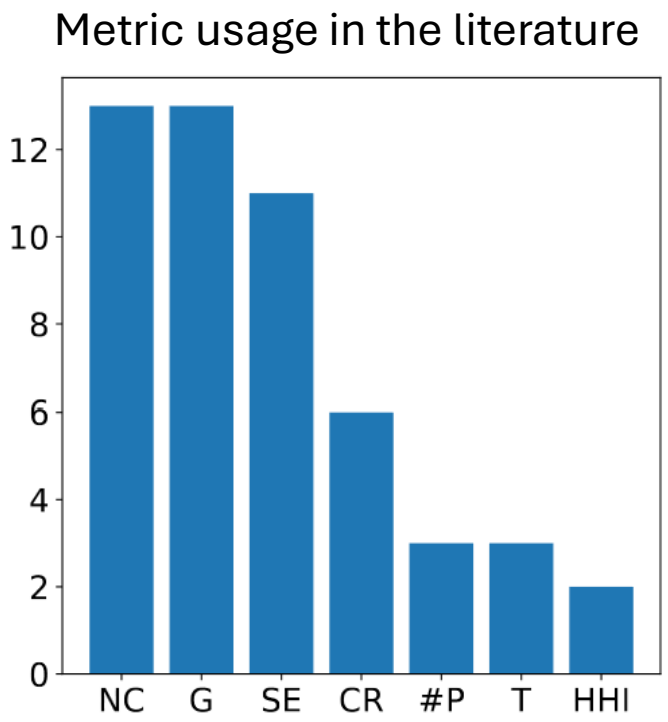

Number of parties

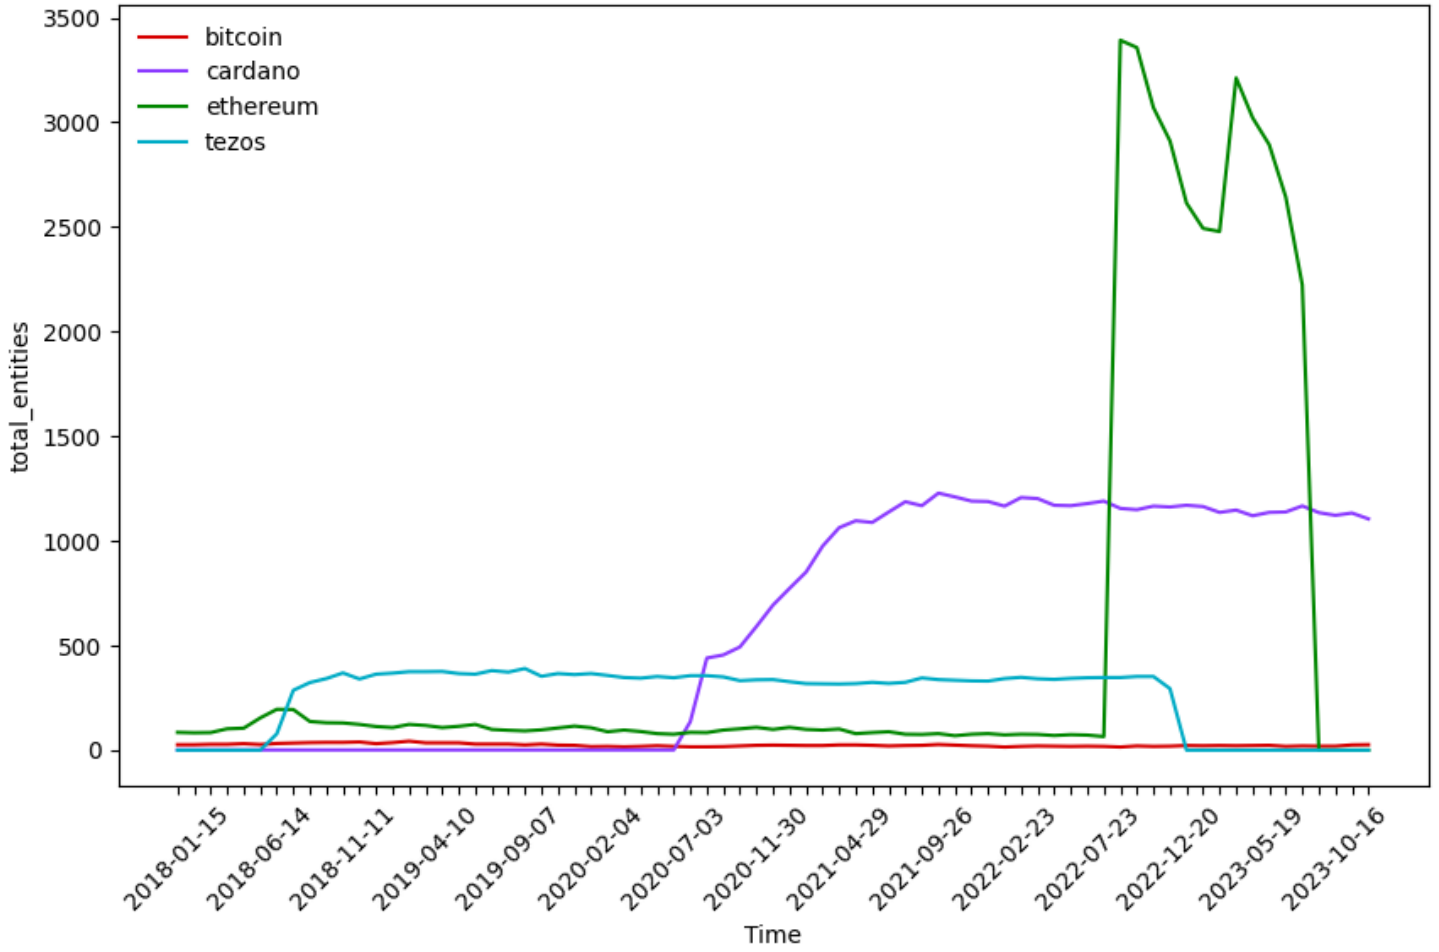

Nakamoto coefficient

- Represents the **minimum number of entities** that collectively control a **majority of resources** (> 50%)

- ..aka the number of parties that need to collude in order to launch a 51% attack

- The higher the Nakamoto coefficient, the higher the resilience to a majority attack

Nakamoto coefficient

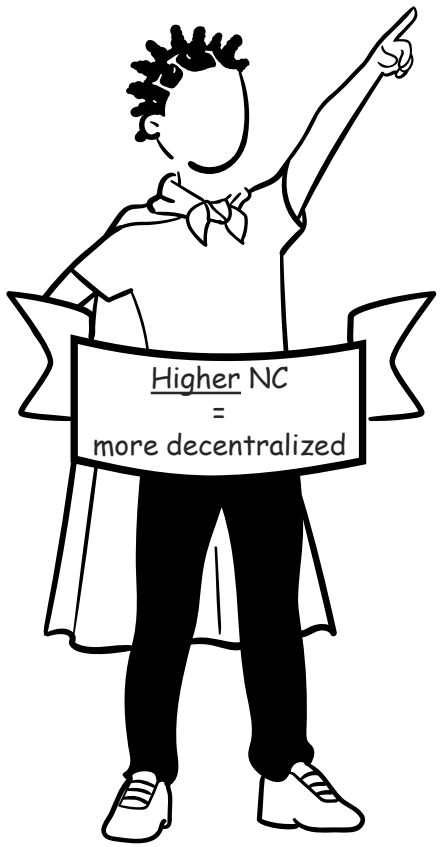

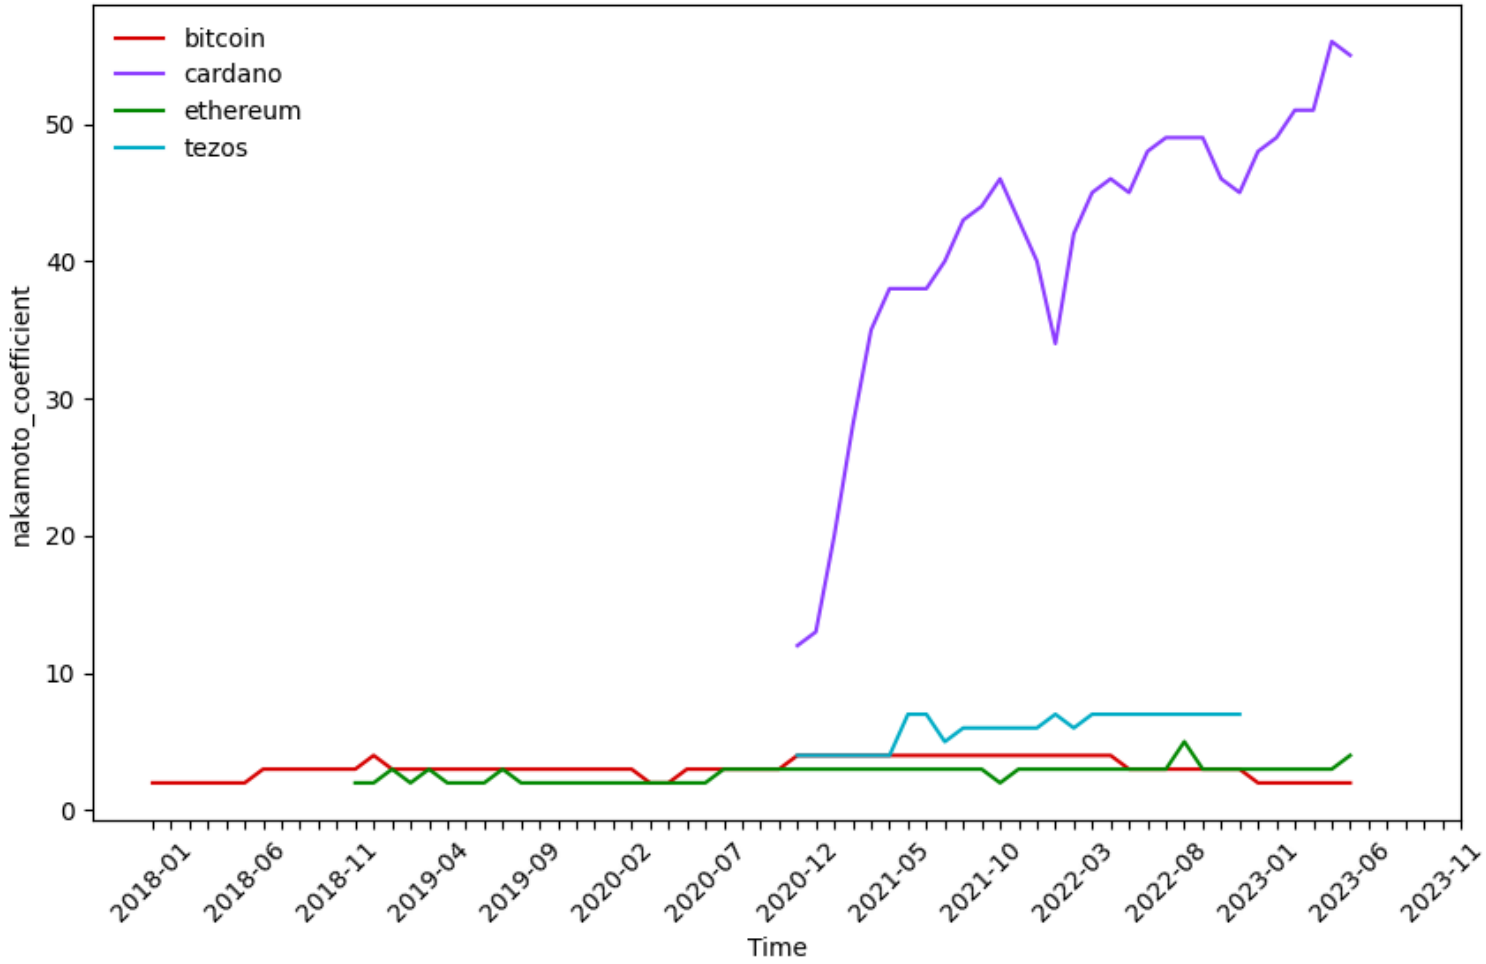

Gini coefficient

- Lorenz curve of resources across entities
  - Points: the cumulative ownership of resources (y-axis) by a percentage of entities (in ascending order) (x-axis)
- Gini coefficient: A / A + B
- Maximum equality ⇒ Gini = 0
  - Every entity holds the same amount of resources
- Maximum inequality ⇒ Gini = 1
  - One entity holds all resources

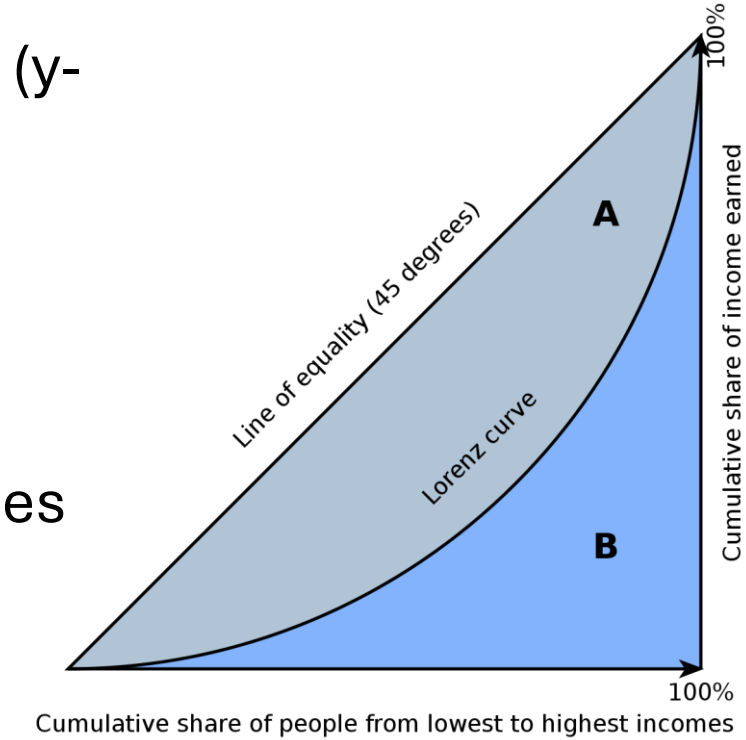

Gini coefficient

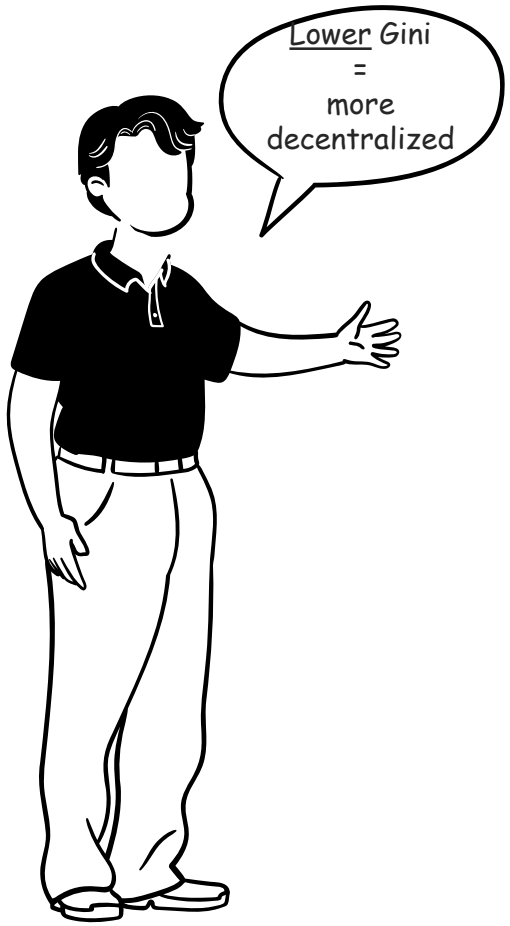

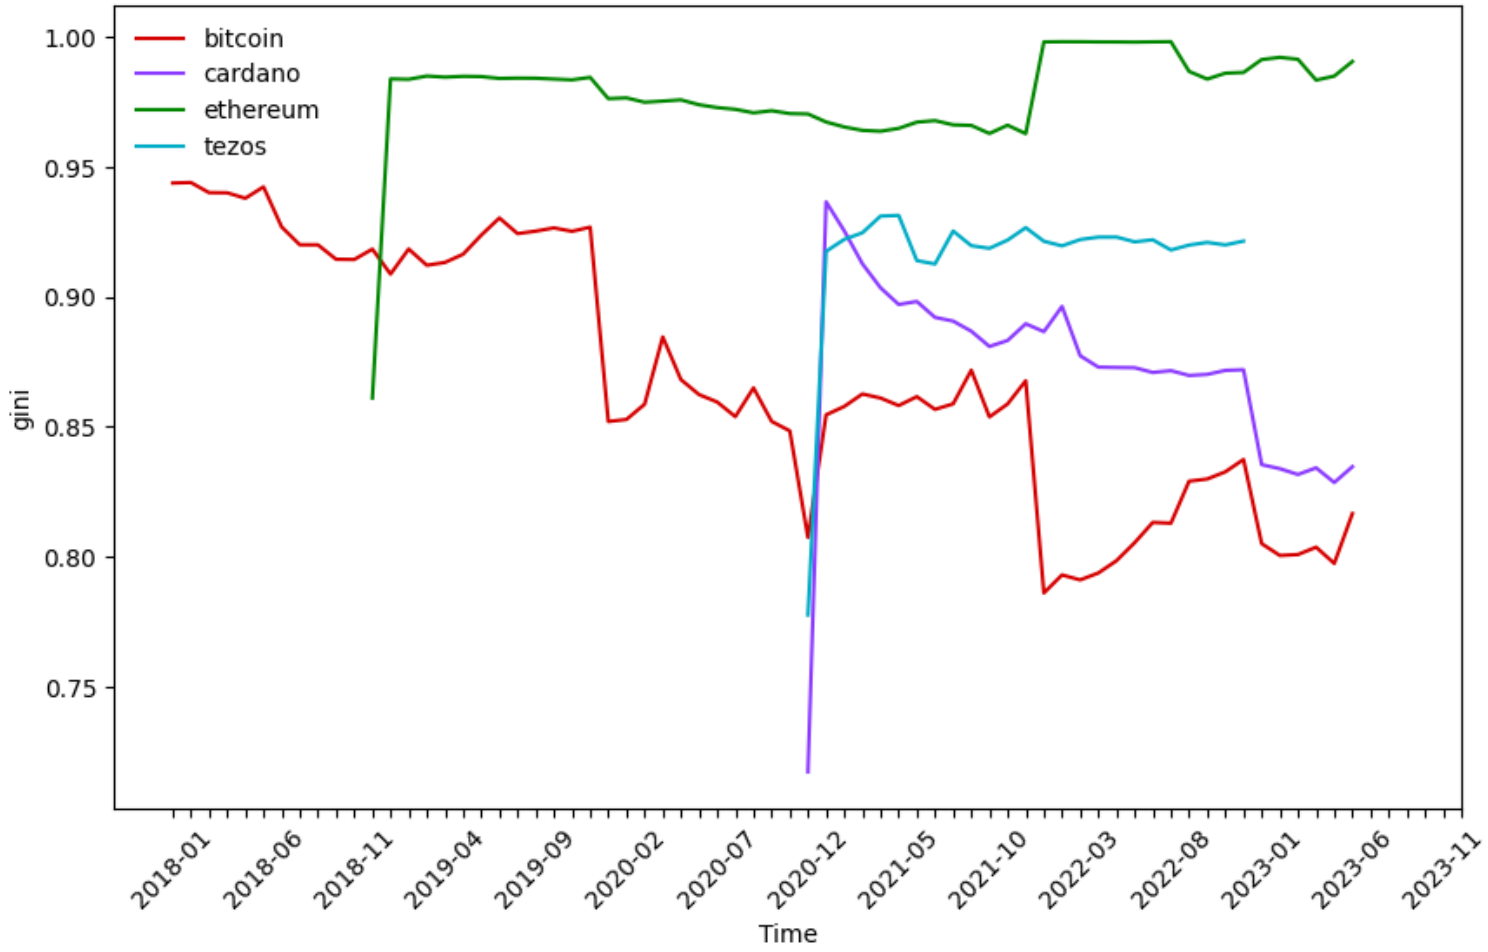

Shannon entropy

- The level (bits) of “information” in a distribution
  - $H(X) = - \Sigma (p(xi) * log2p(xi))$ (X: random variable, p(x): probability of event X=x)
- In our case:
  - $entropy(S) = - \Sigma (f(Si) * log2f(Si))$
  - f(Si): the relative resources of entity Si compared to all resources (as a percentage)
- Resources centralized around a few entities ⇒ Lower entropy
- Resources distributed among many entities ⇒ Higher entropy
  - Max entropy: when resources are evenly distributed among all entities

Shannon entropy

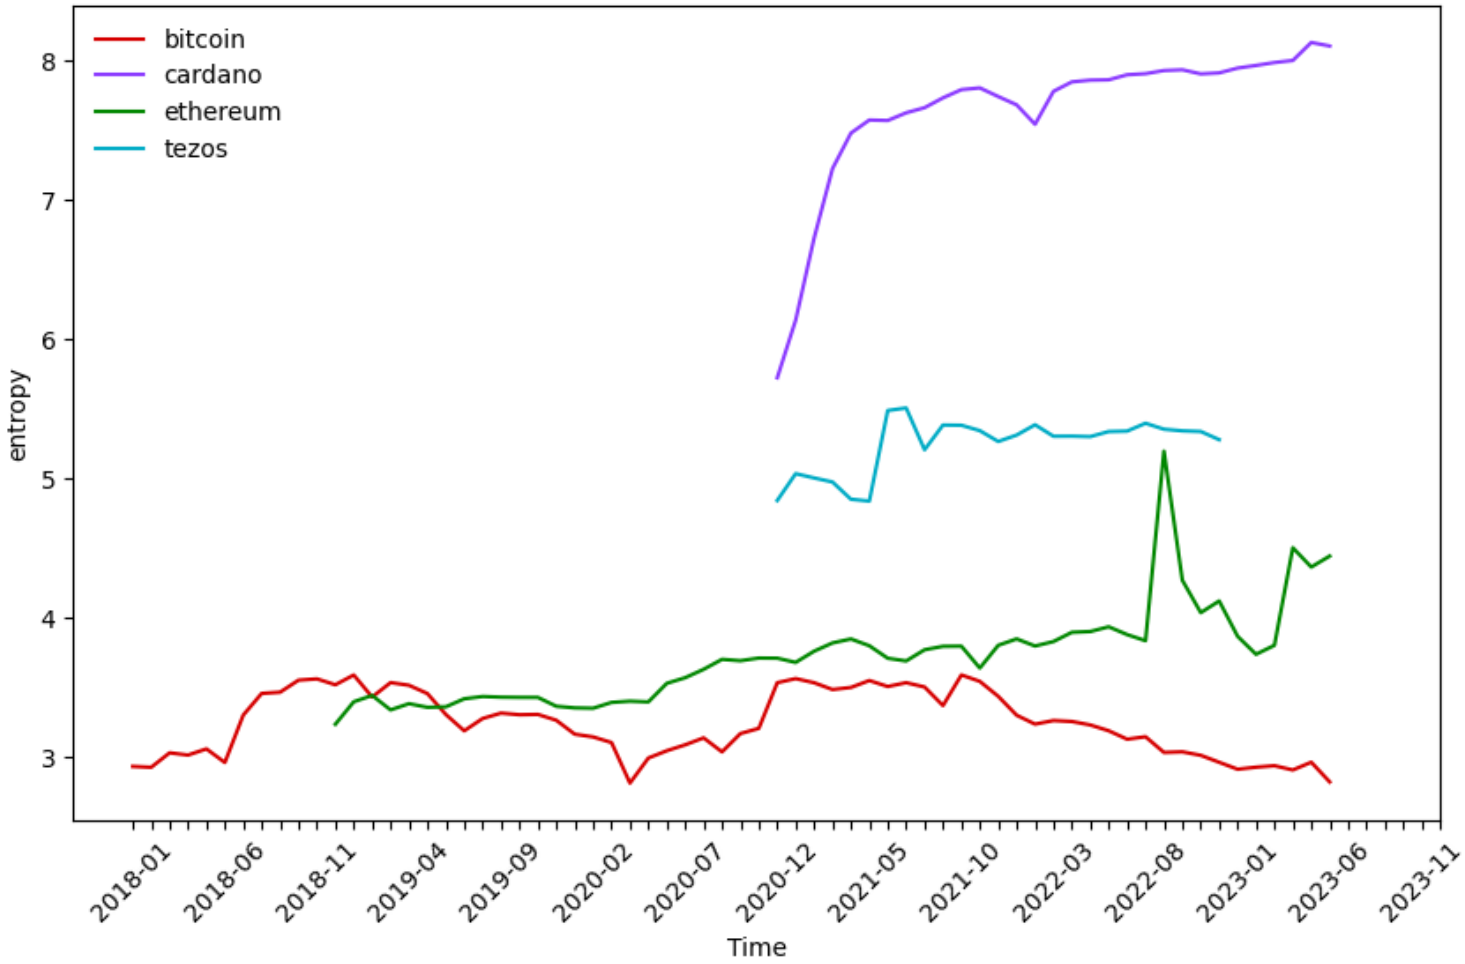

Herfindahl–Hirschman index (HHI)

- Market concentration metric
- Can be calculated as follows:

$$HHI = \sum_{i=1}^{n} s_i^2$$

Where s_i is the market share of firm *i* (as a whole number, e.g. 20 for 20%)

- U.S. Department of Justice guidelines:
  - HHI < 1,500: competitive market
  - 1,500 ≤ HHI ≤ 2,500: moderately concentrated market
  - HHI > 2,500: highly concentrated market
- Values chosen for traditional markets, may need different thresholds for blockchains

Herfindahl–Hirschman index (HHI)

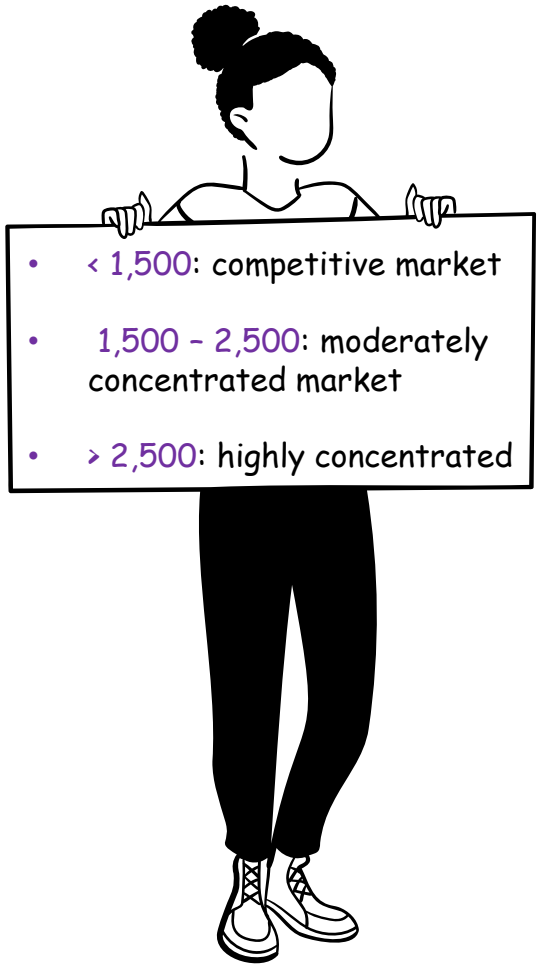

U.S. Department of Justice

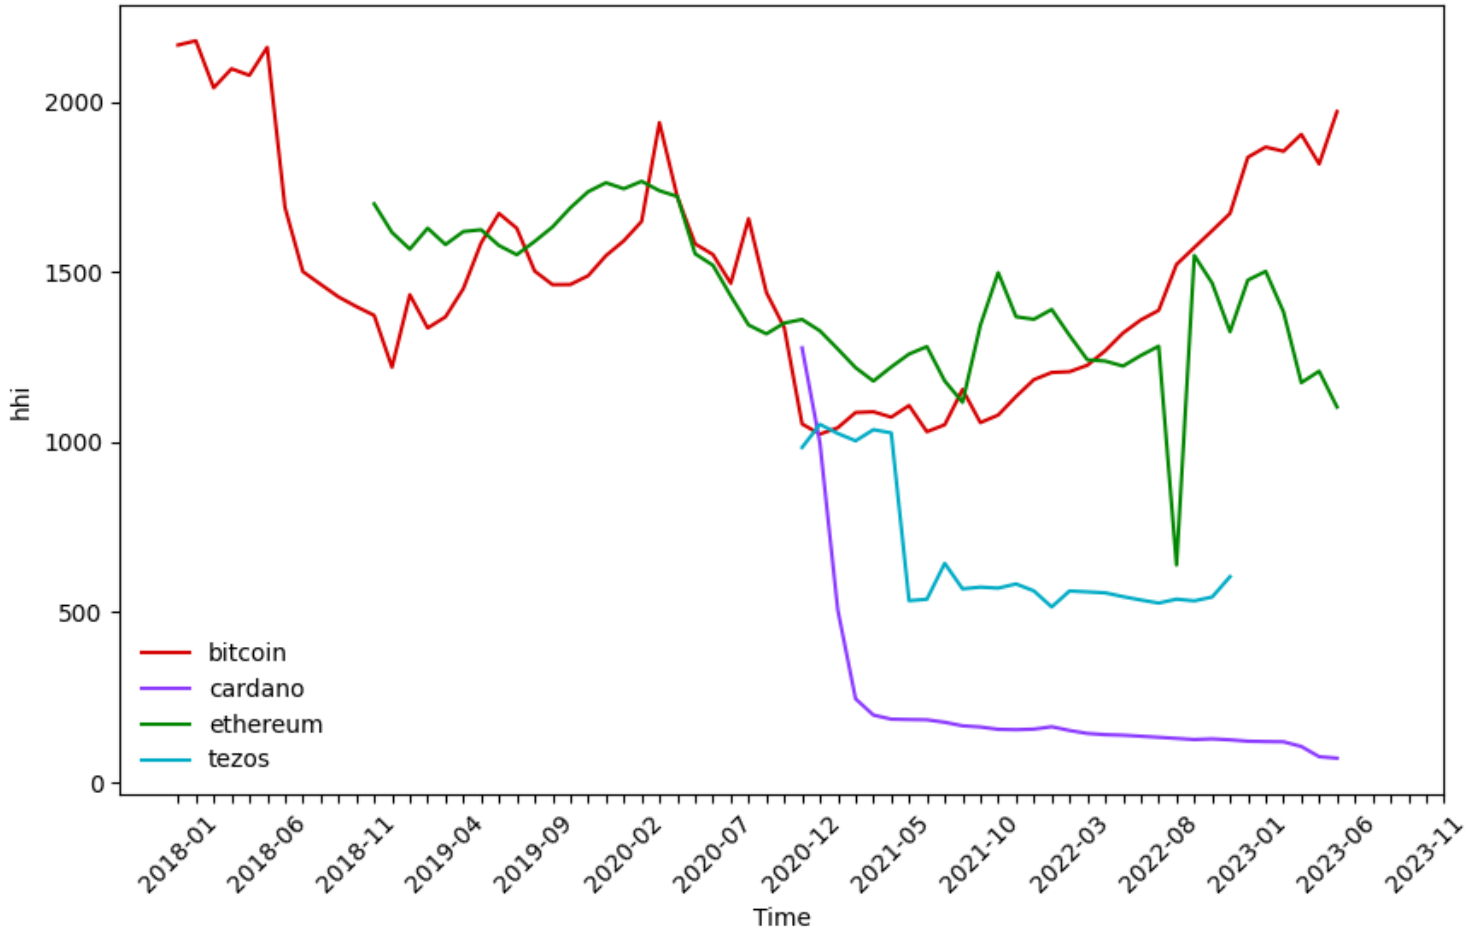

Do all metrics capture the same signal?

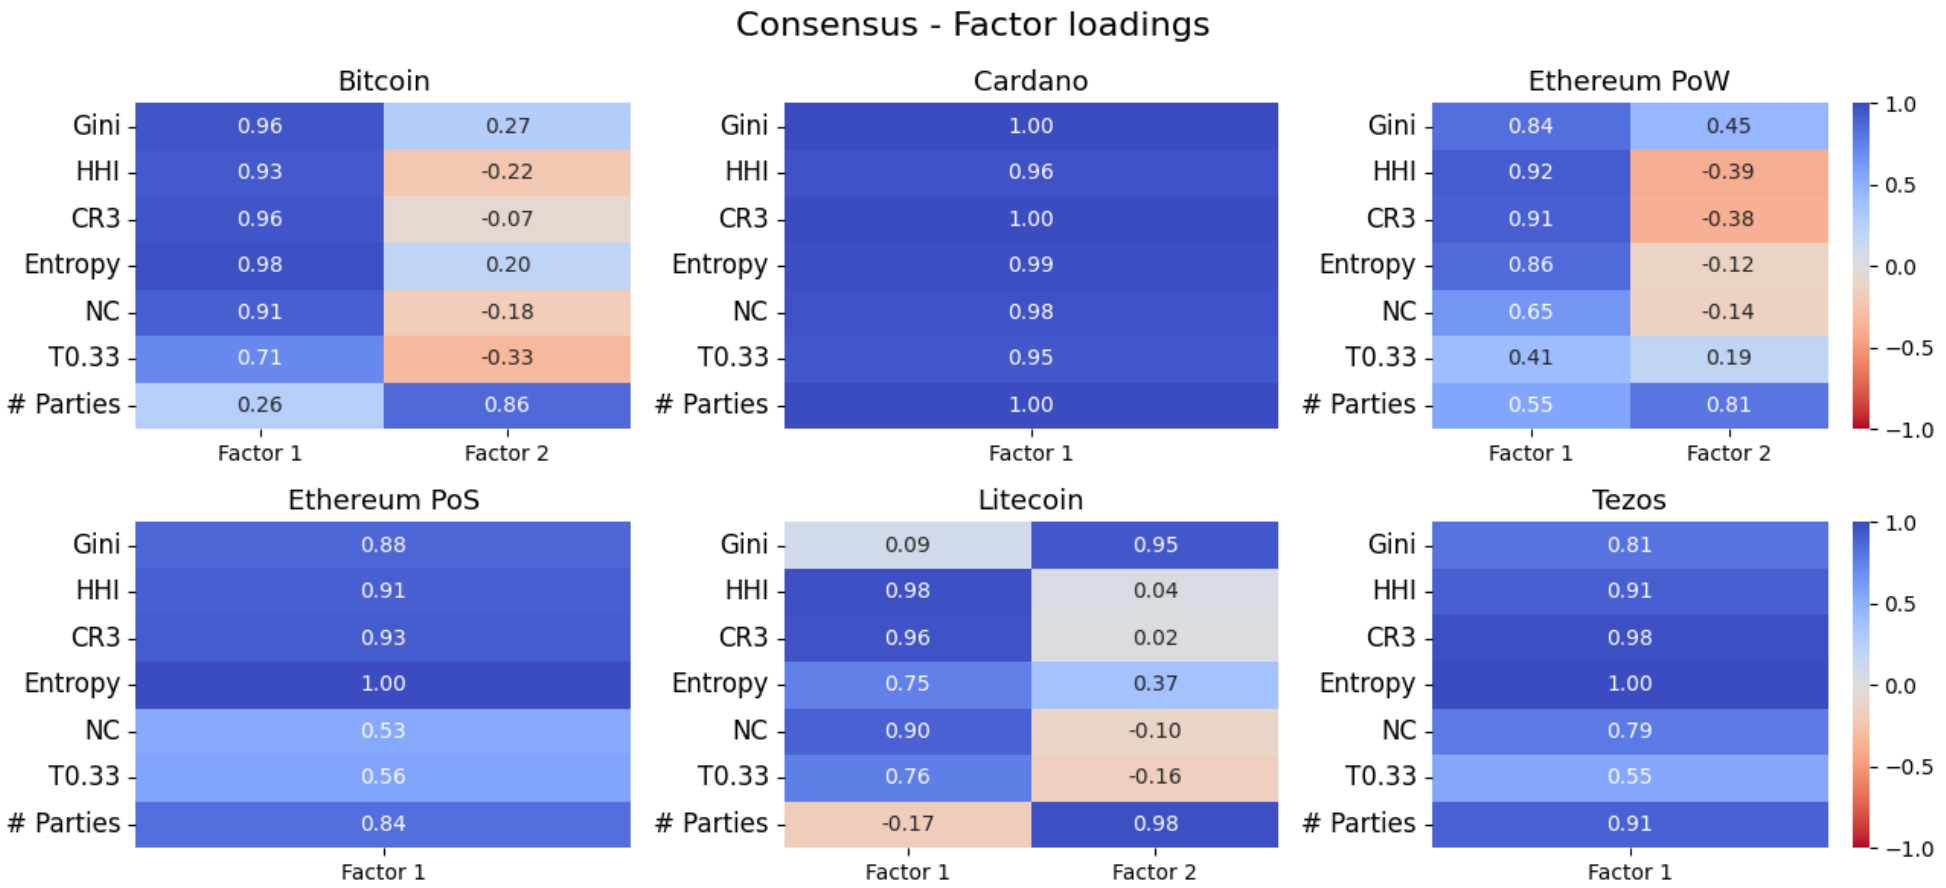

Do all metrics capture the same signal?

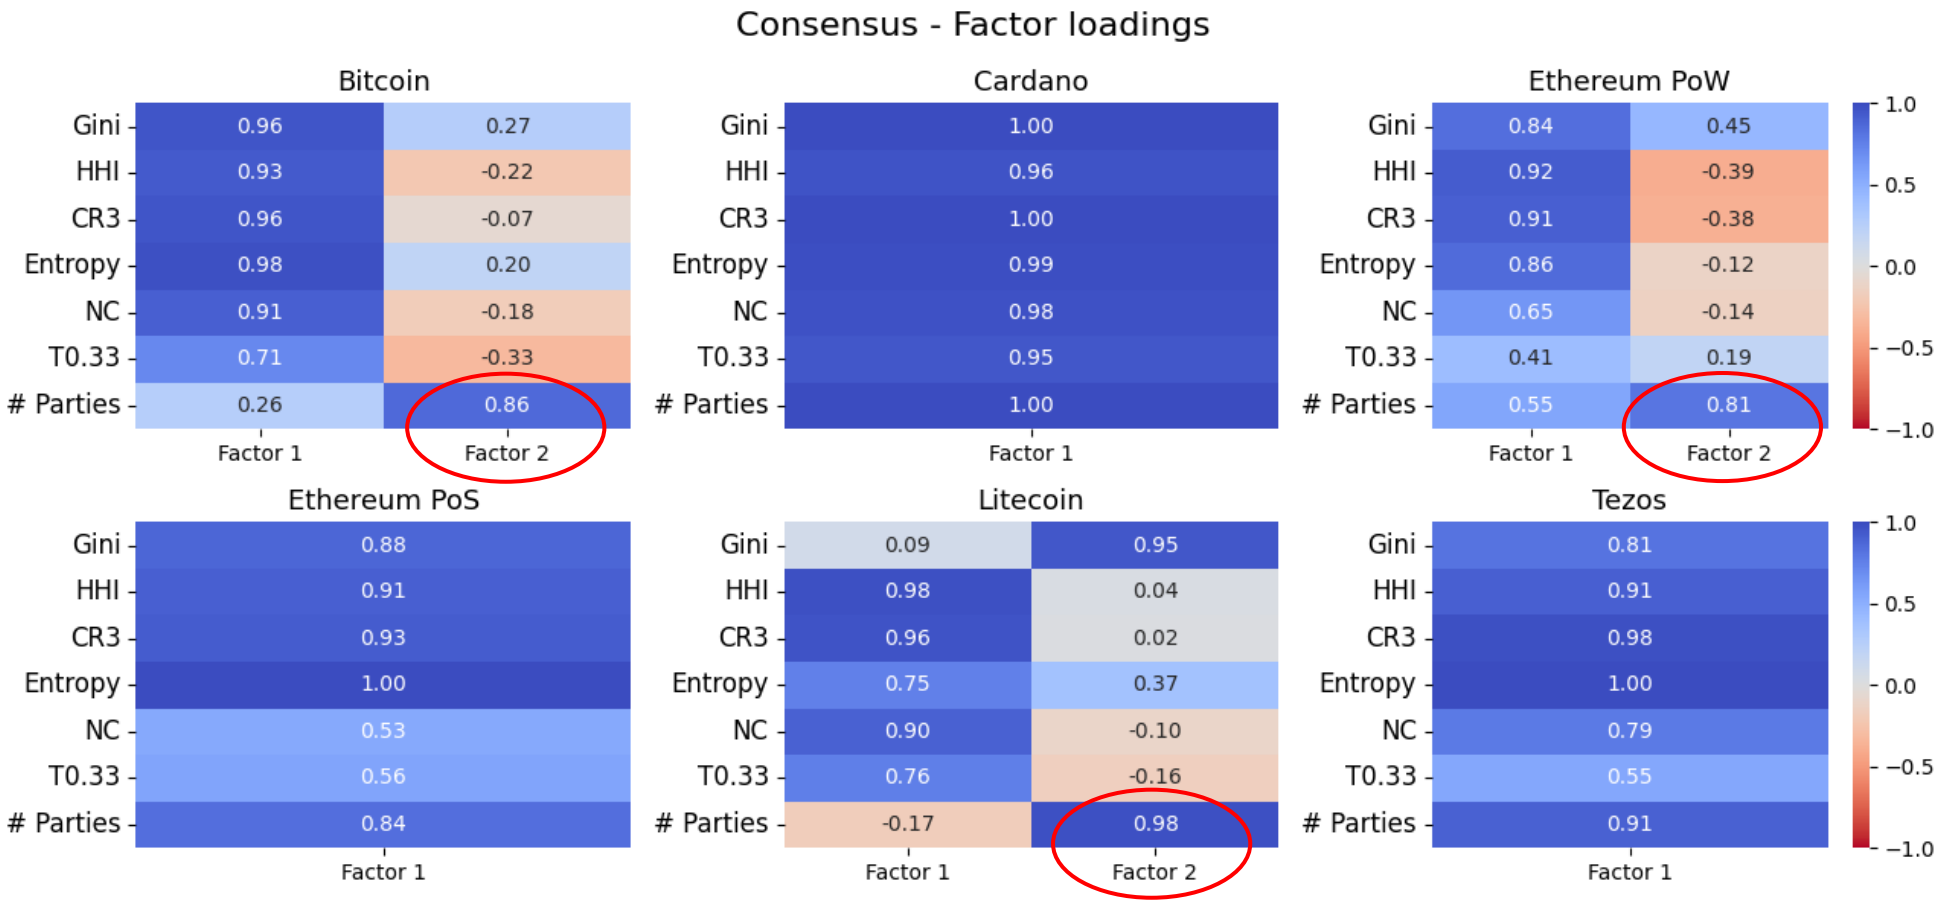

Open questions & future work

- How do different design choices or components impact decentralization?
  - Proof of Work vs Proof of Stake
  - Simple vs sophisticated reward schemes
  - Proposer-builder separation
- Can we merge all relevant metrics and layers into a single index that provides a holistic representation of a blockchain system's decentralization?

Edinburgh Decentralization Index (EDI) Public dashboard

- Live public dashboard for multiple layers and systems: http://blockchainlab.inf.ed.ac.uk/edi-dashboard/

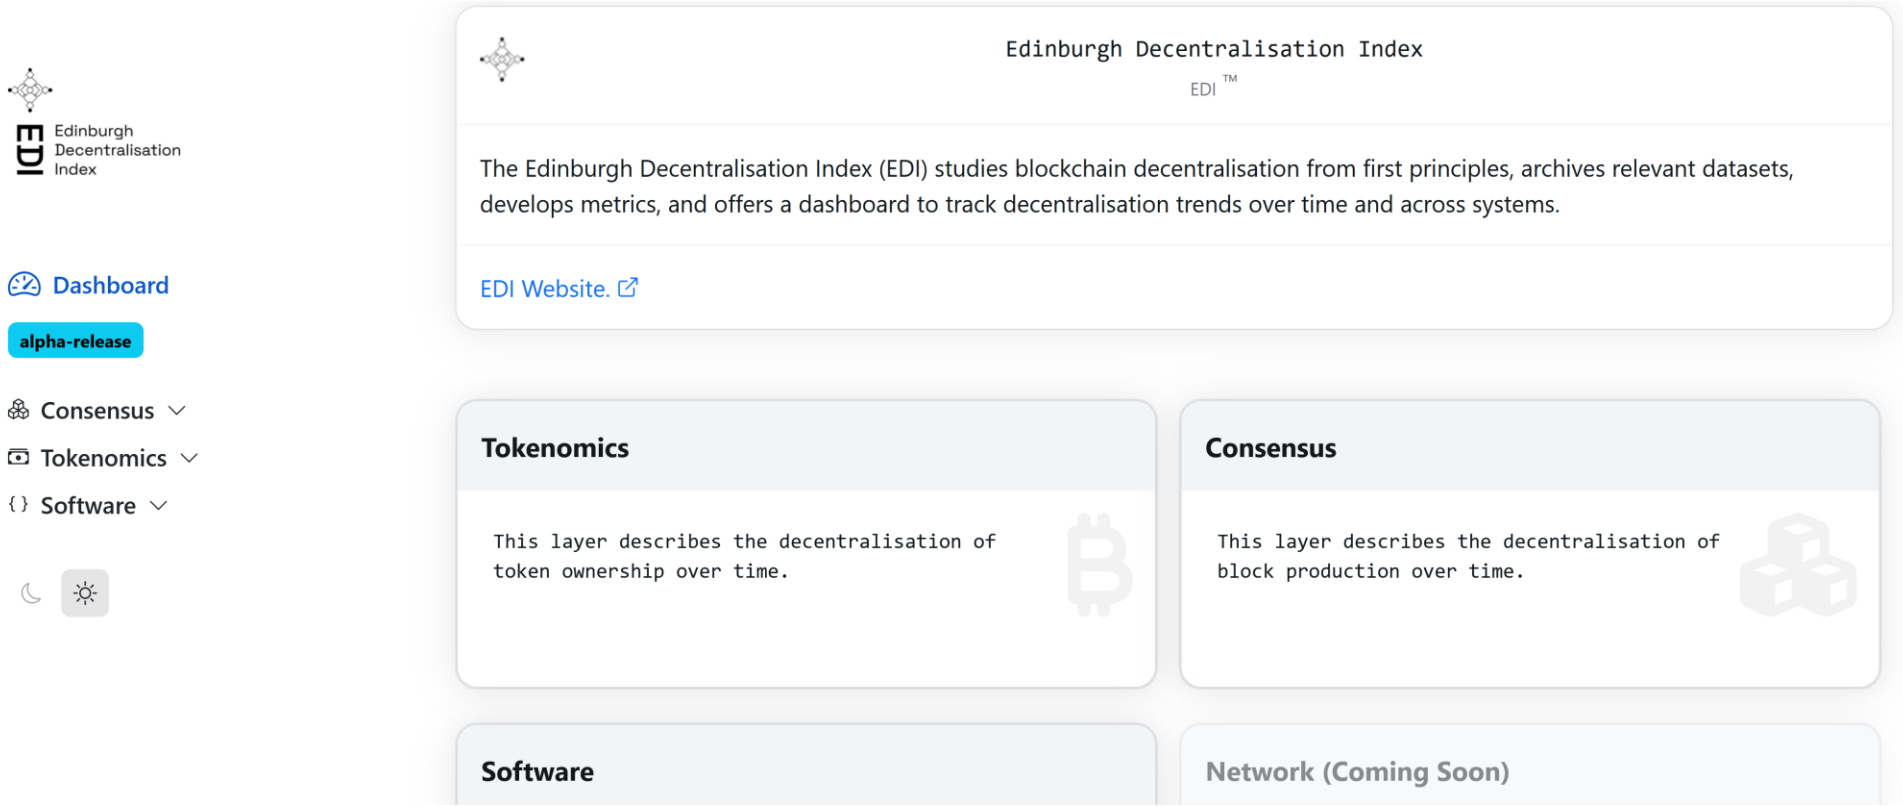

THE UNIVERSITY of EDINBURGH
informatics

Thank you!

Questions?

christina.ovezik@ed.ac.uk