Bobtail: Improved Blockchain Security With Low-Variance Mining

GEORGE BISSIAS BRIAN LEVINE

UNIVERSITY OF MASSACHUSETTS AMHERST

Compressed Review of Blockchains

- We focus on *public / open* blockchains that use proof-of-work (PoW)
- *Decentralized* and *distributed* ledgers
  - Ledger comprises set of transactions
  - Financial, logistical, legal, …
- PoW: not the only approach, but most popular and relatively easy to analyze

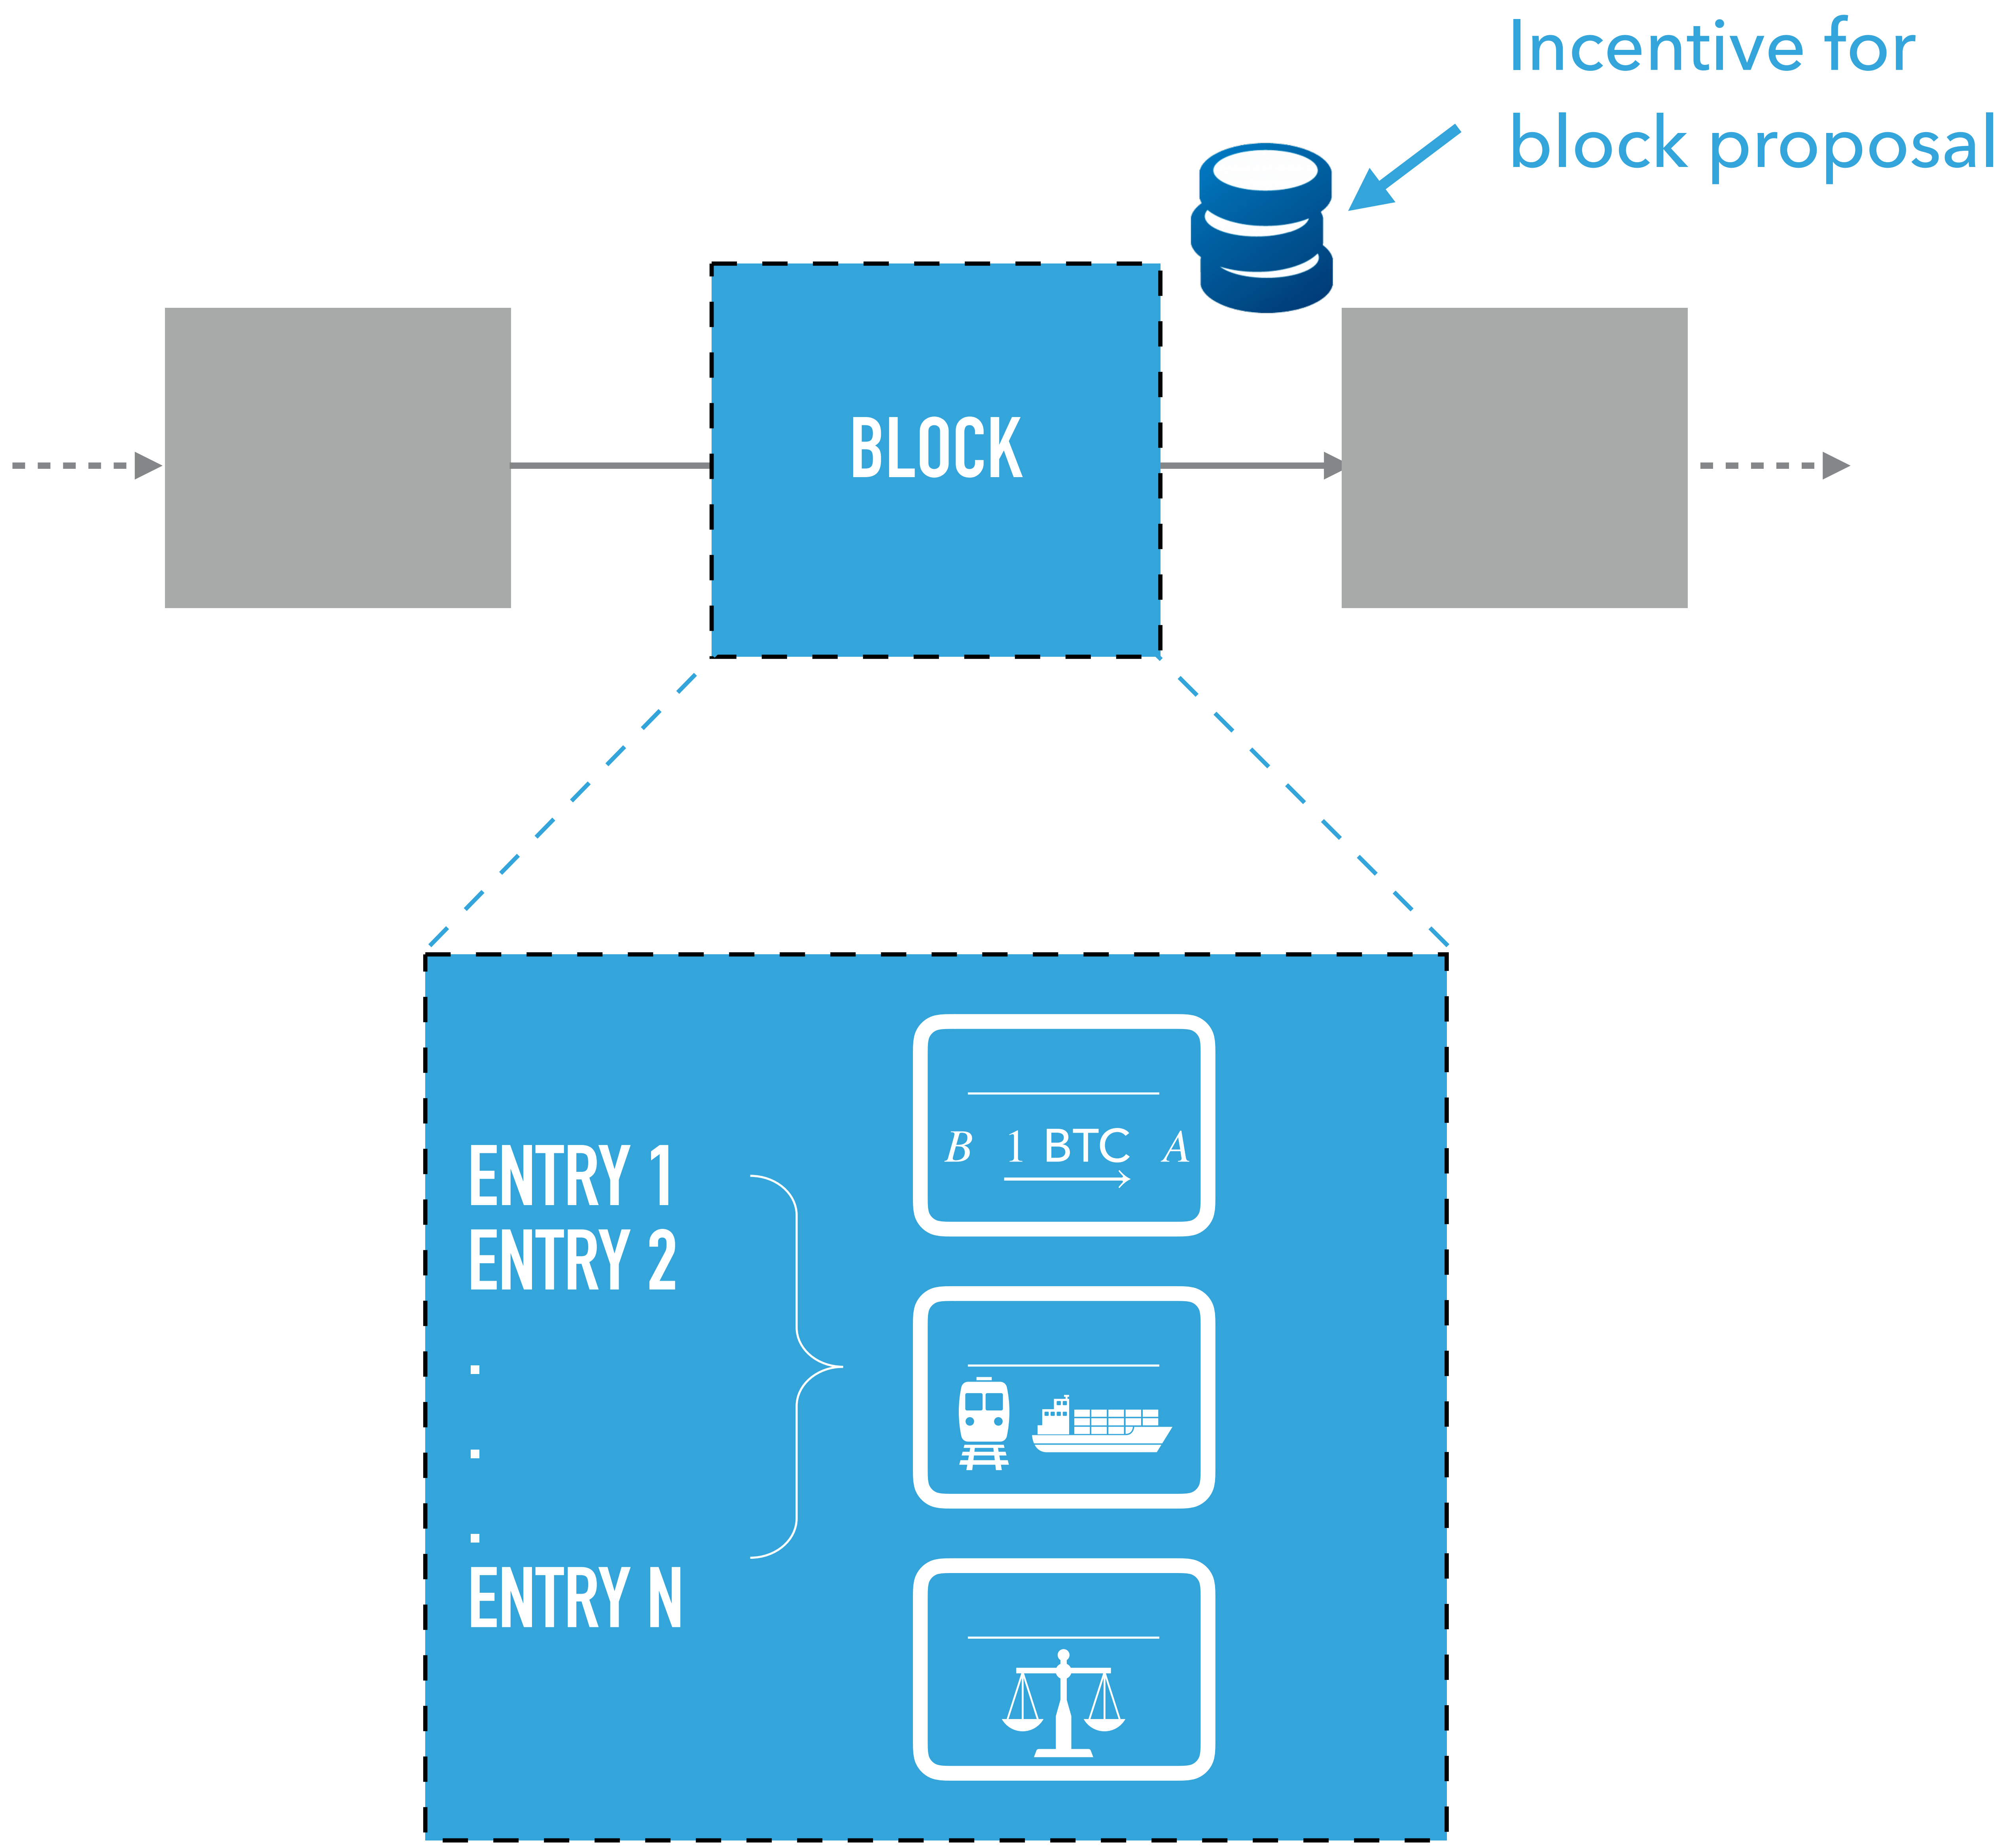

Proof-of-Work Mining Basics

- Miners repeatedly hash block header
- Hashes are within $[0,S]$
- A block is mined when hash falls below t
- Block time T is function of hash rate h (seconds)
- Convention is to extend longest chain

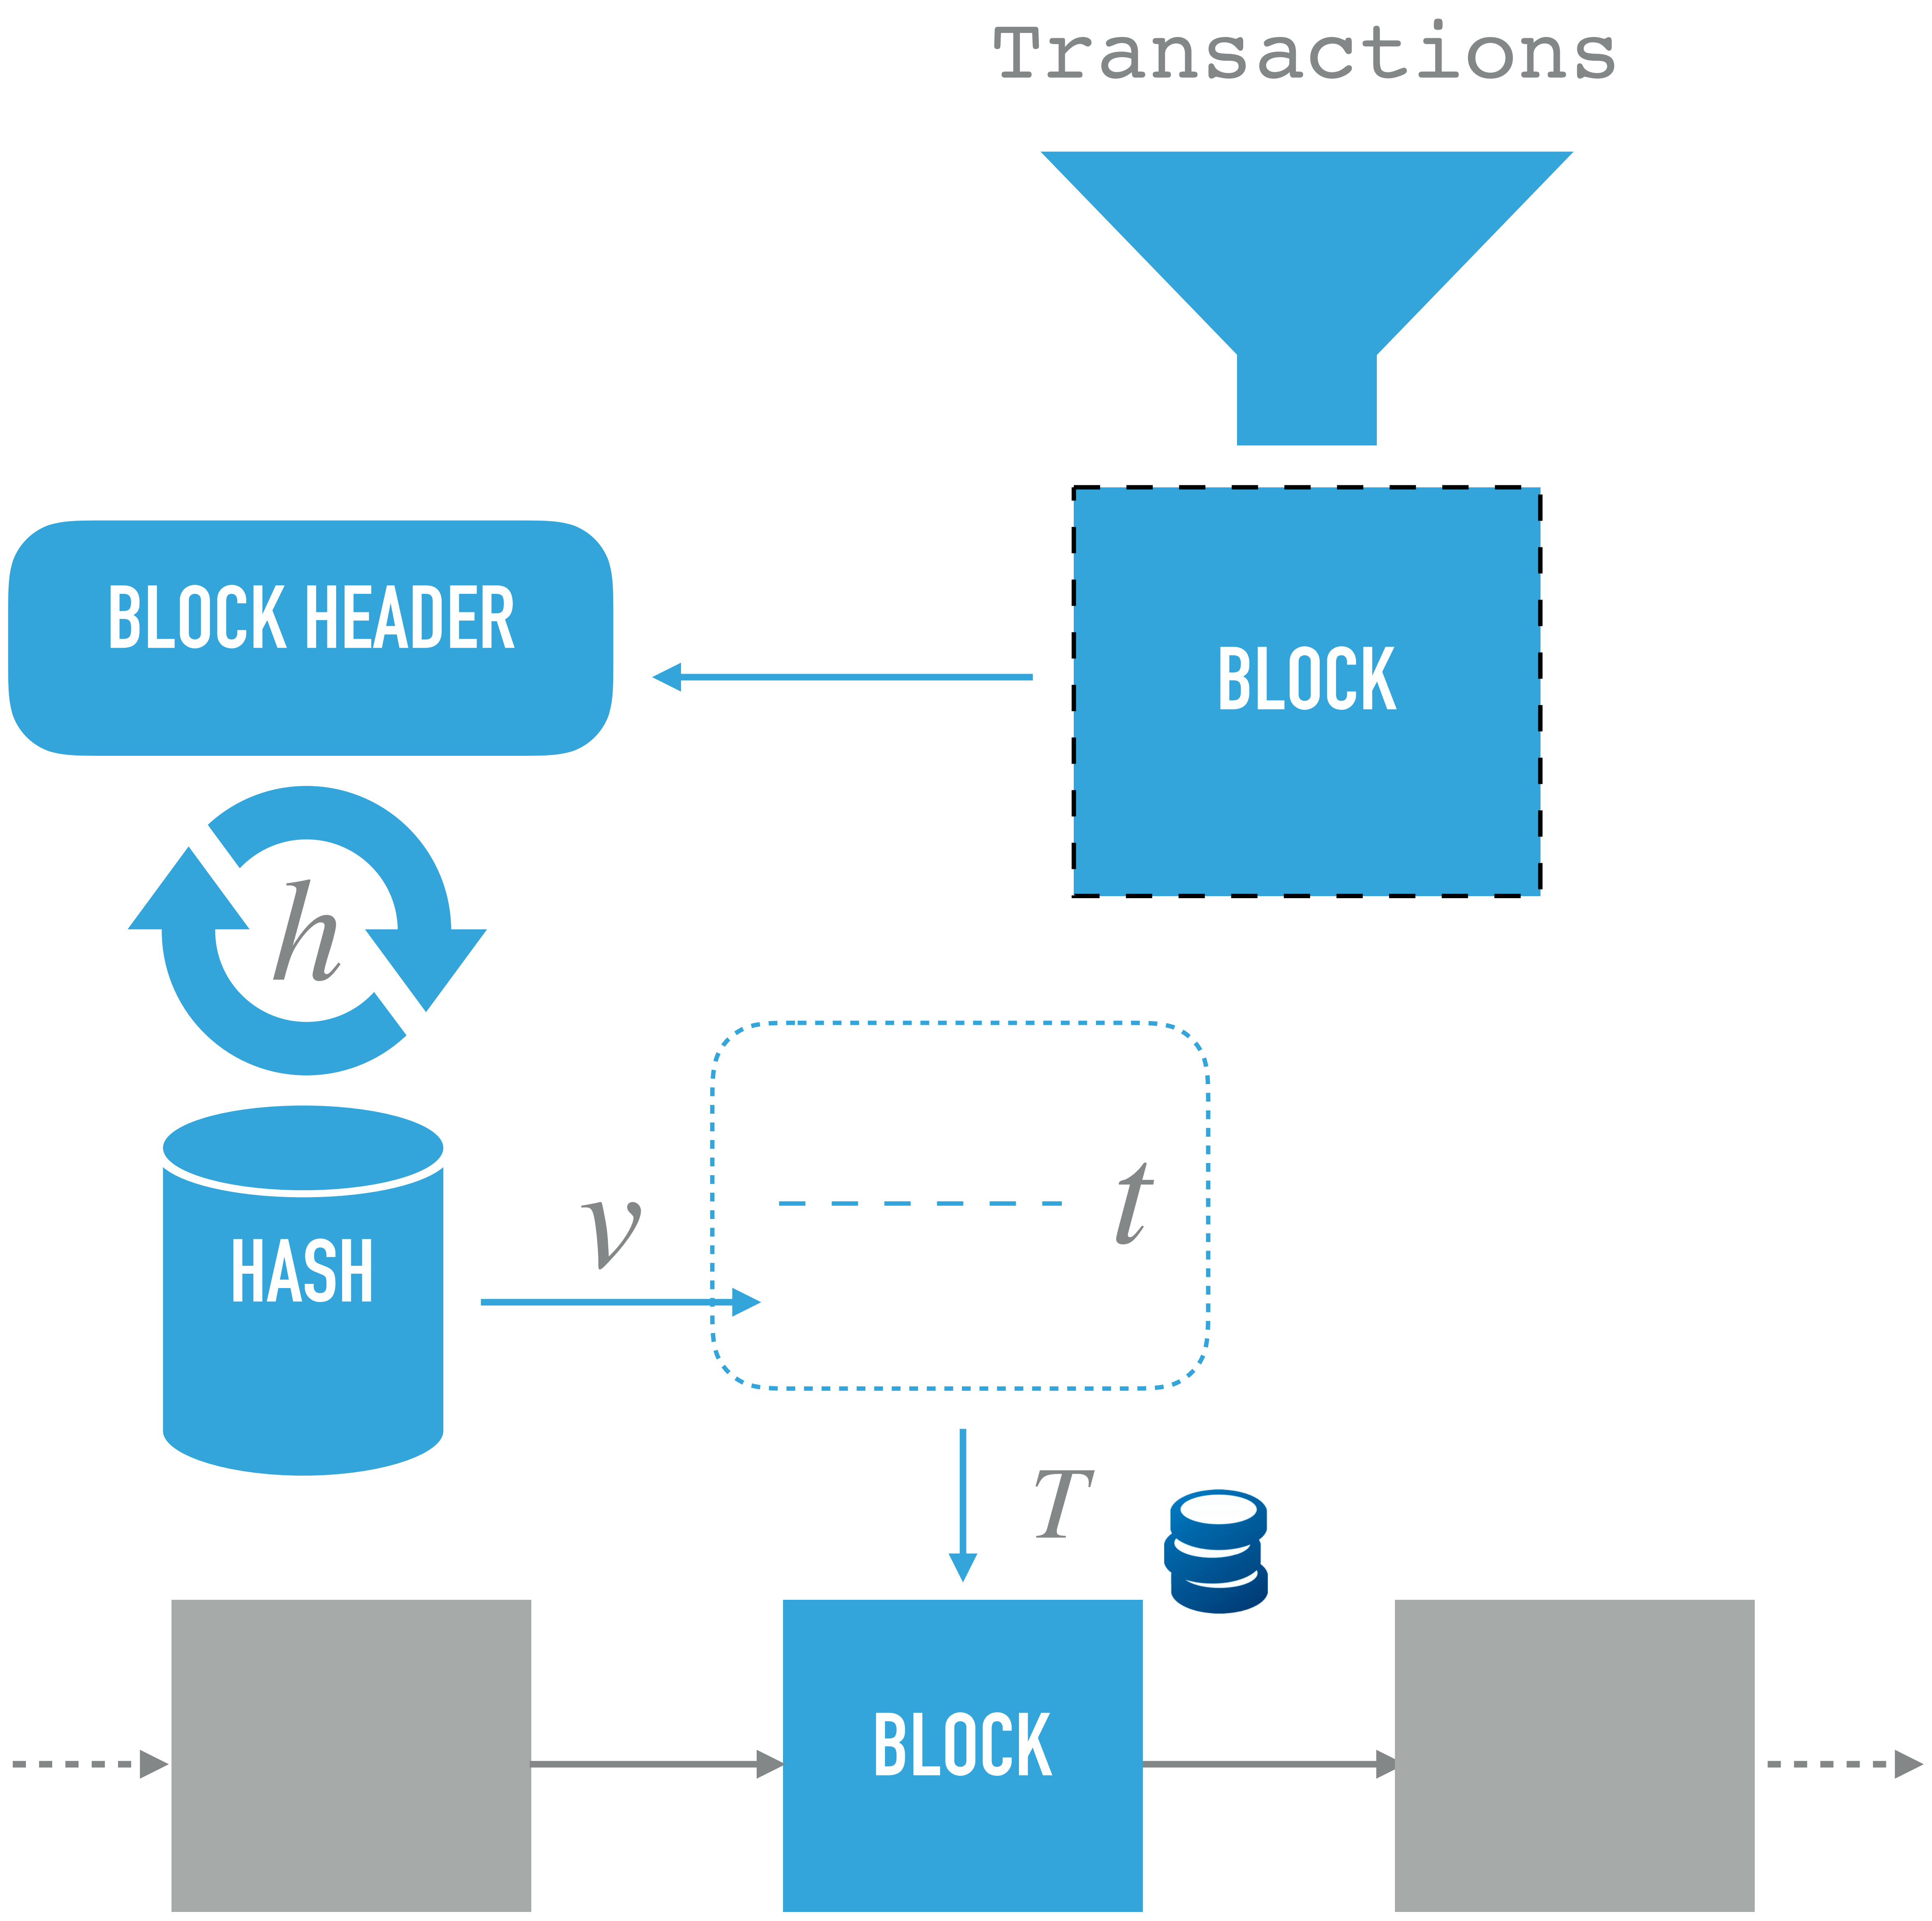

Mining is a Lottery

- Miners “draw” numbers until they cross threshold 5

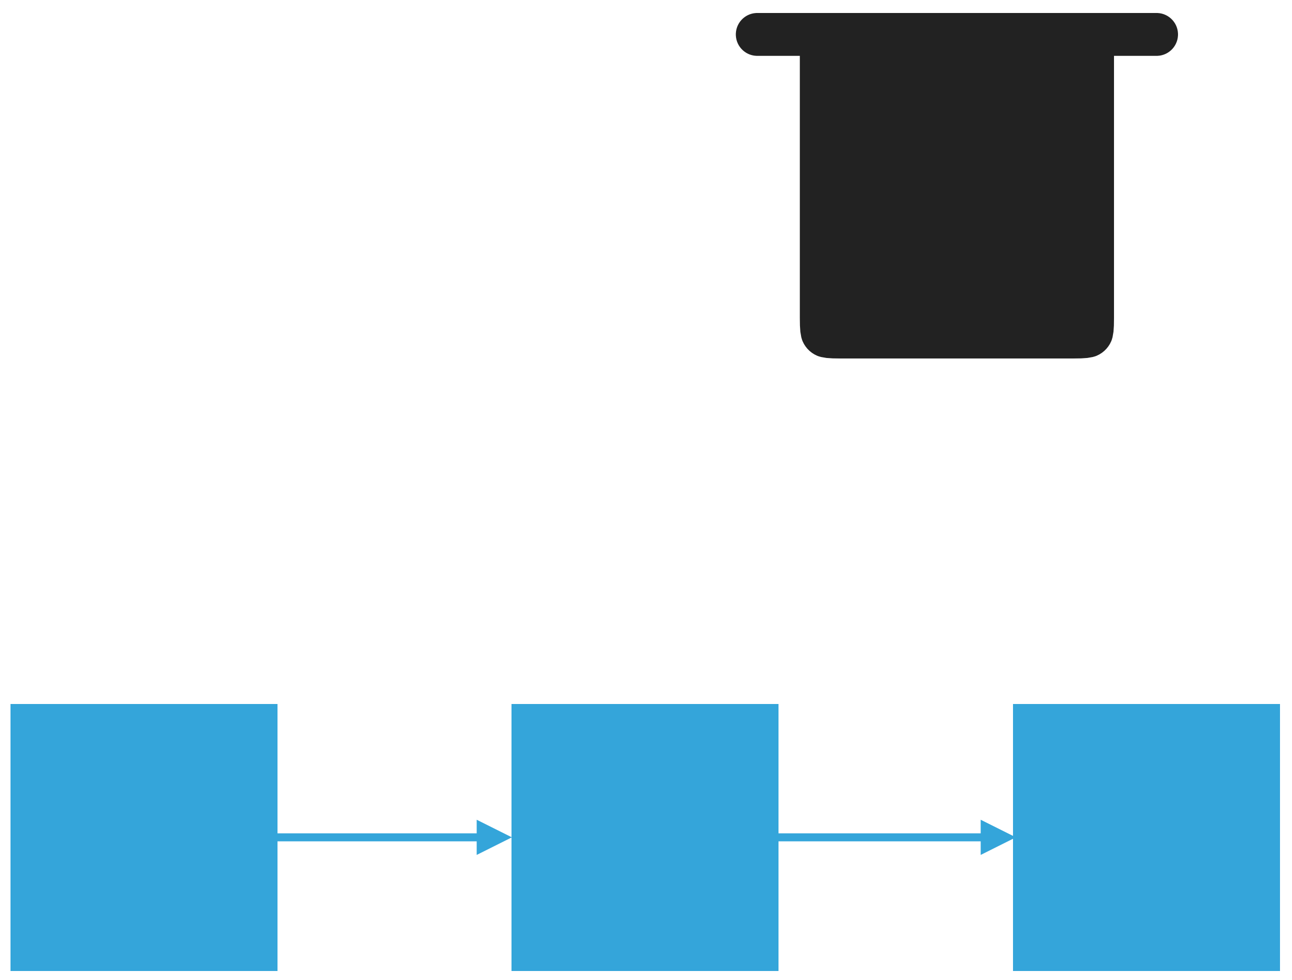

Mining is a Lottery

- Miners "draw" numbers until they cross threshold 5
- Each draw "costs" a hash

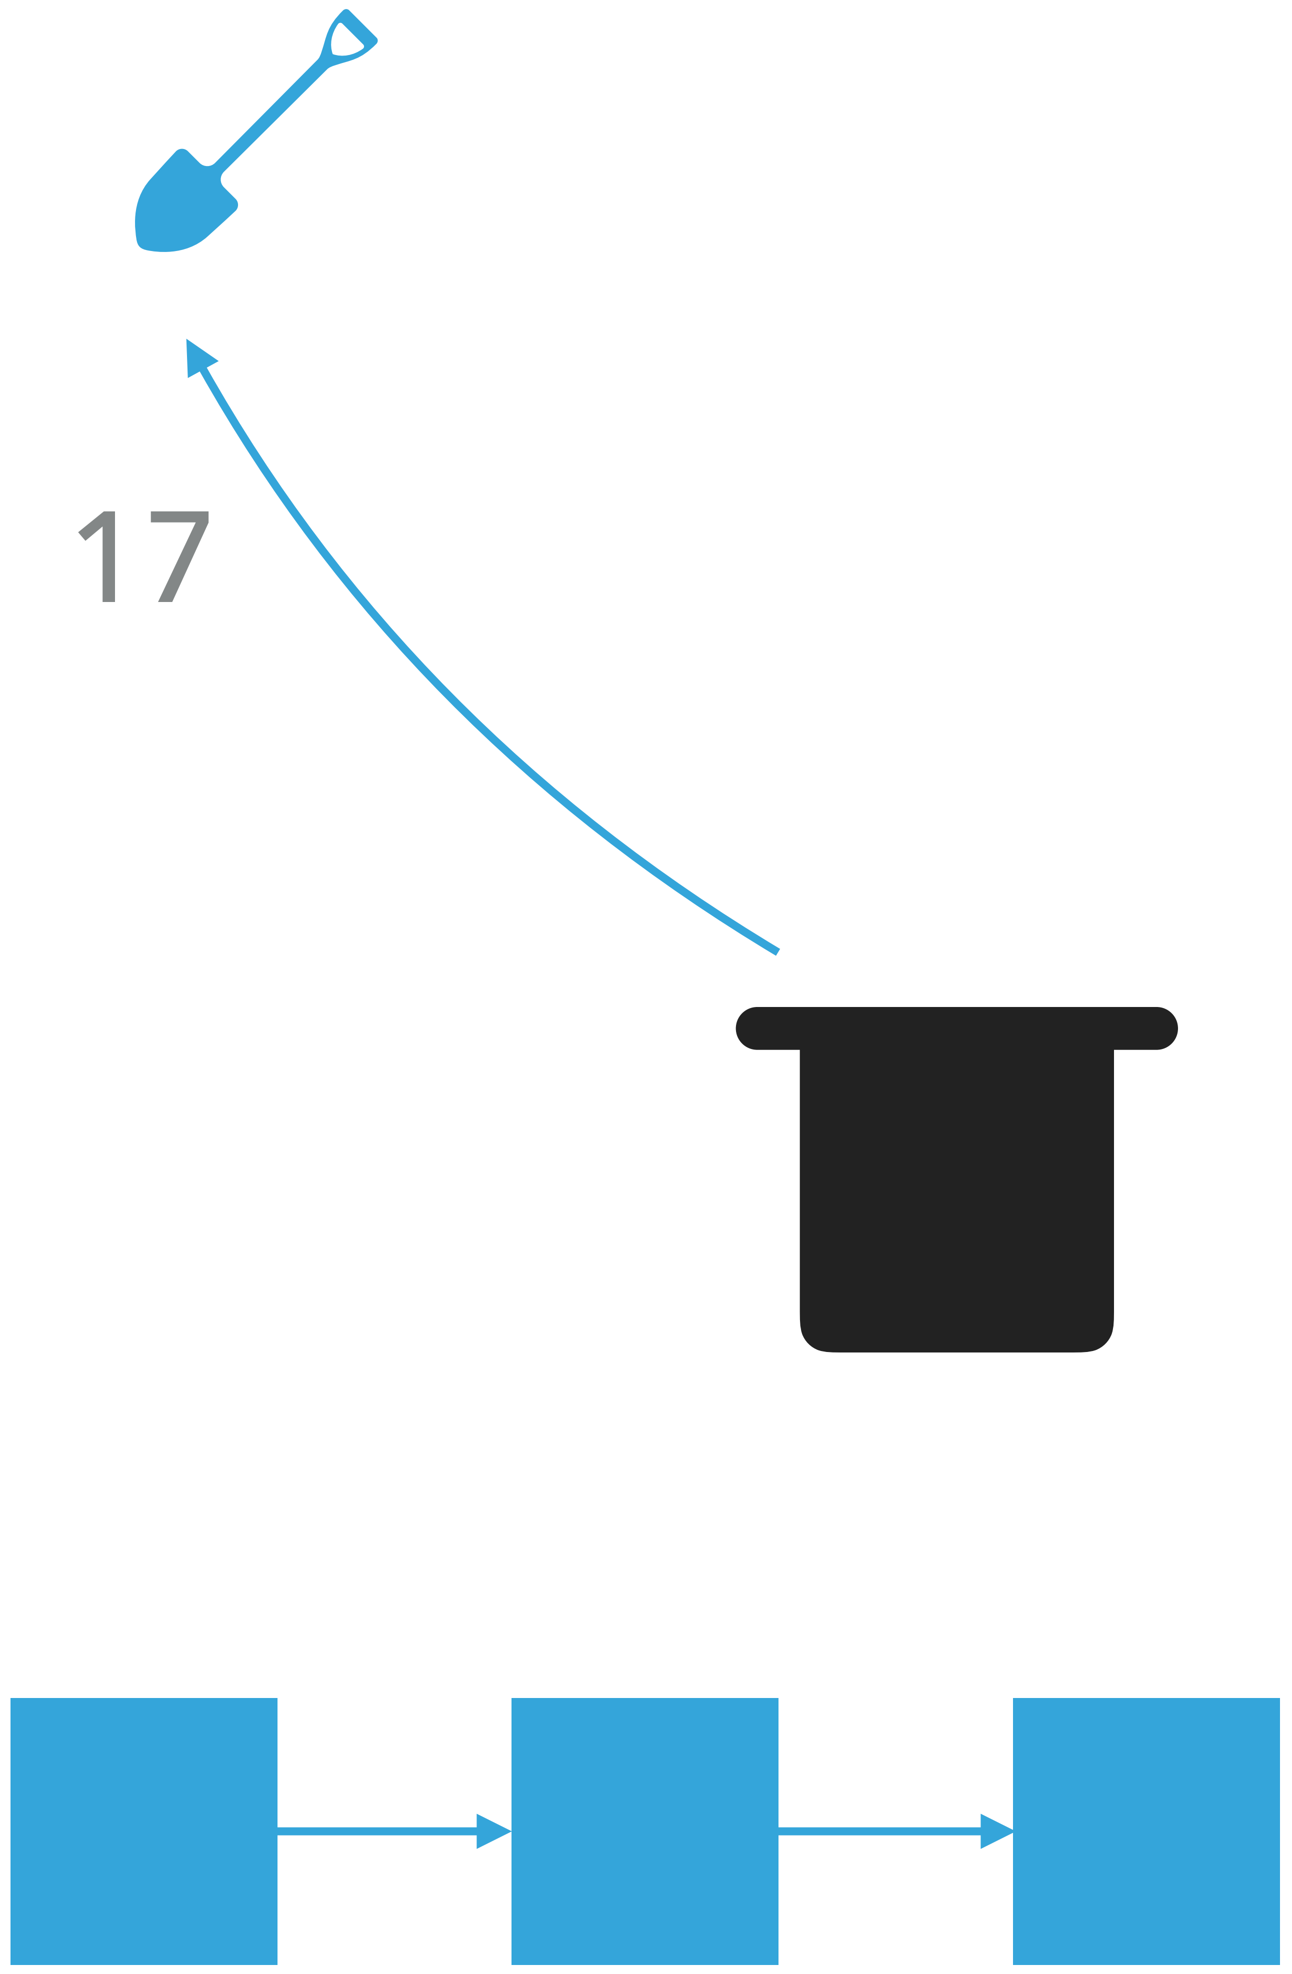

Mining is a Lottery

- Miners “draw” numbers until they cross threshold 5
- Each draw “costs” a hash

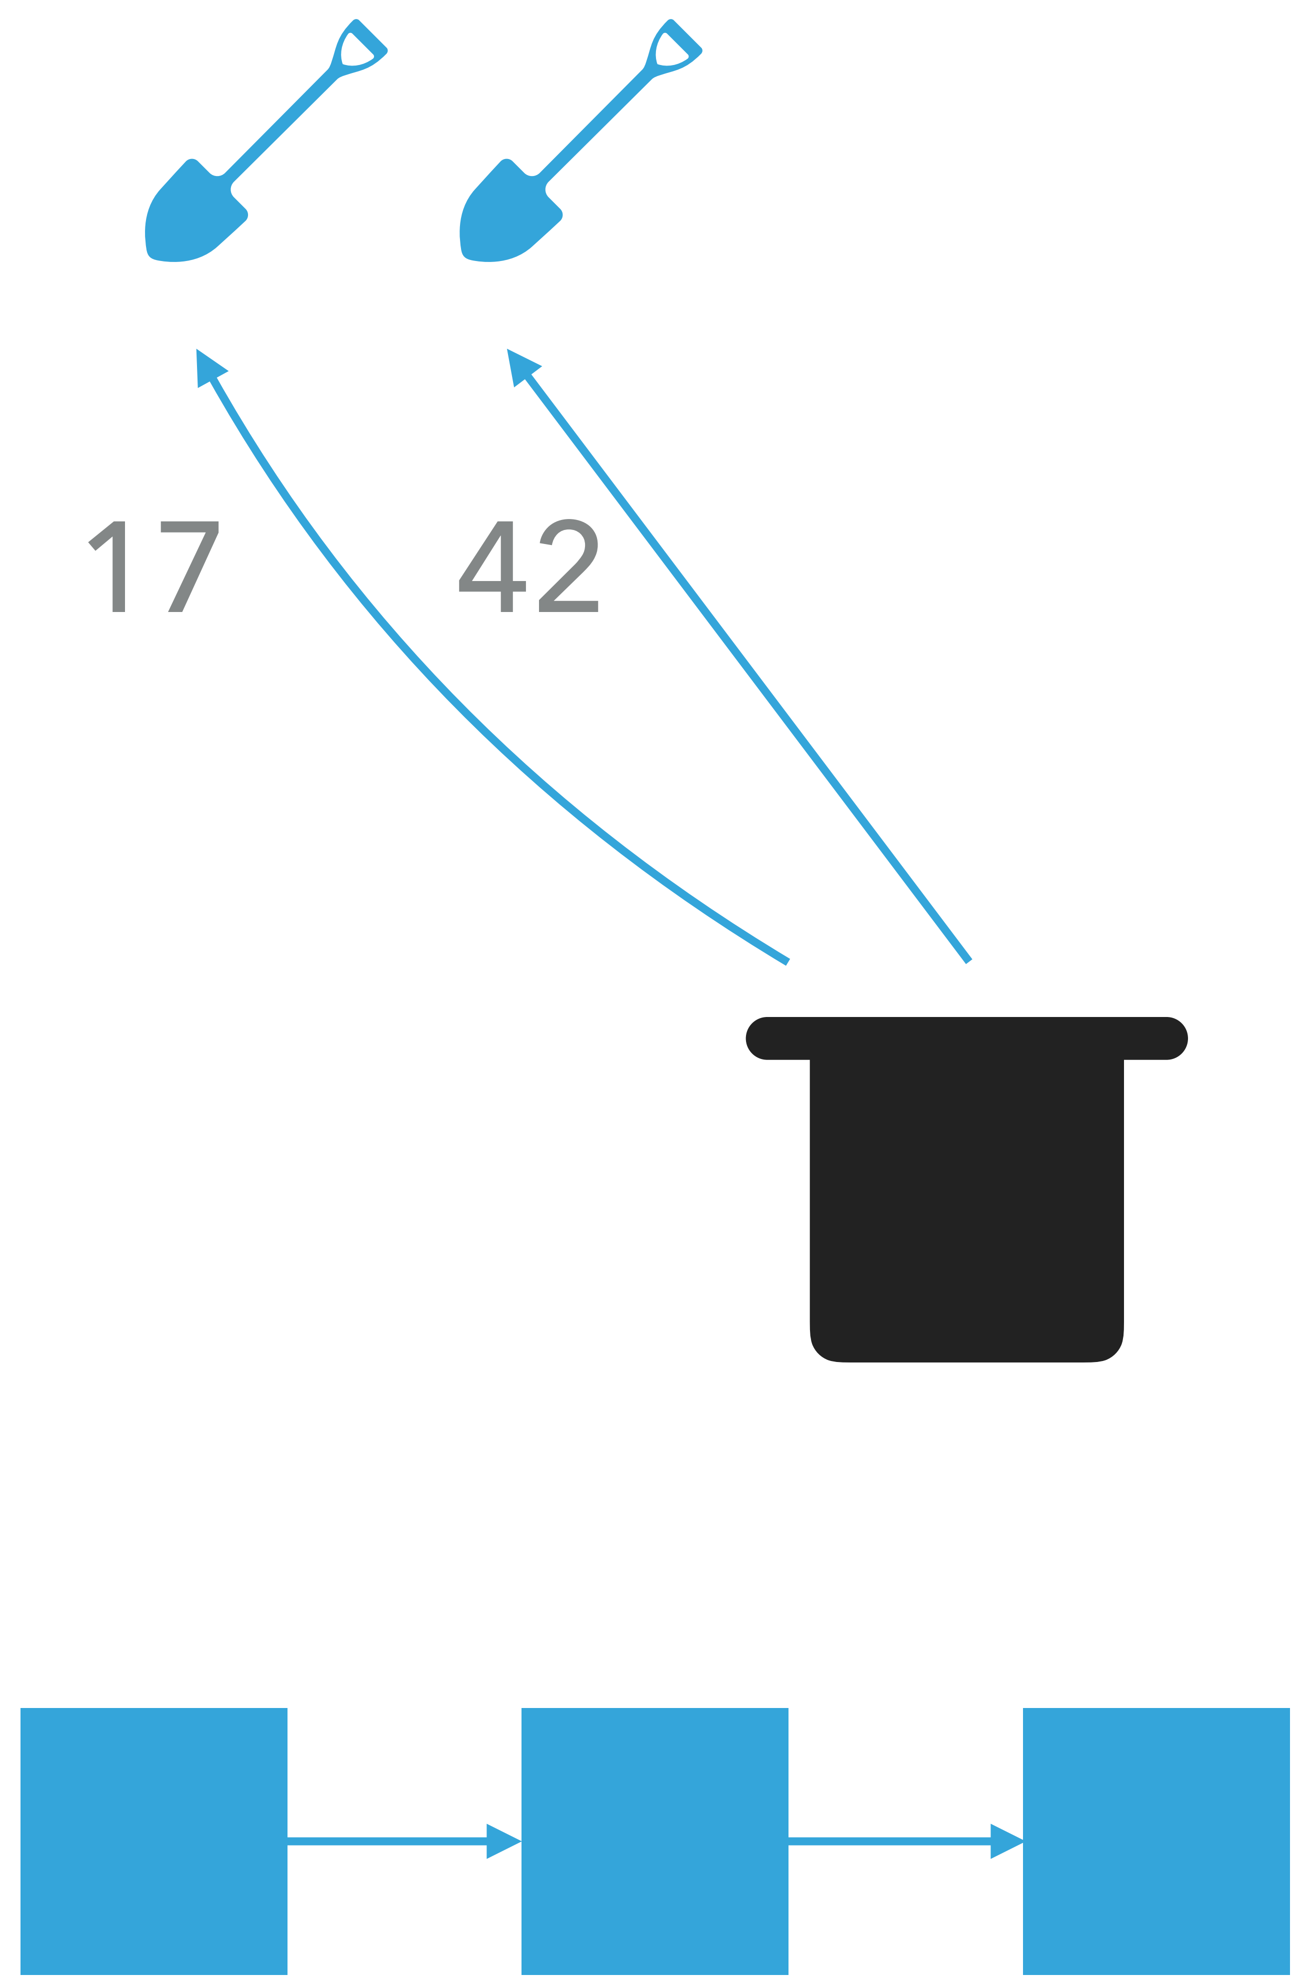

Mining is a Lottery

- Miners "draw" numbers until they cross threshold 5
- Each draw "costs" a hash
- First to cross threshold wins
- Winner receives a reward and proposes a block

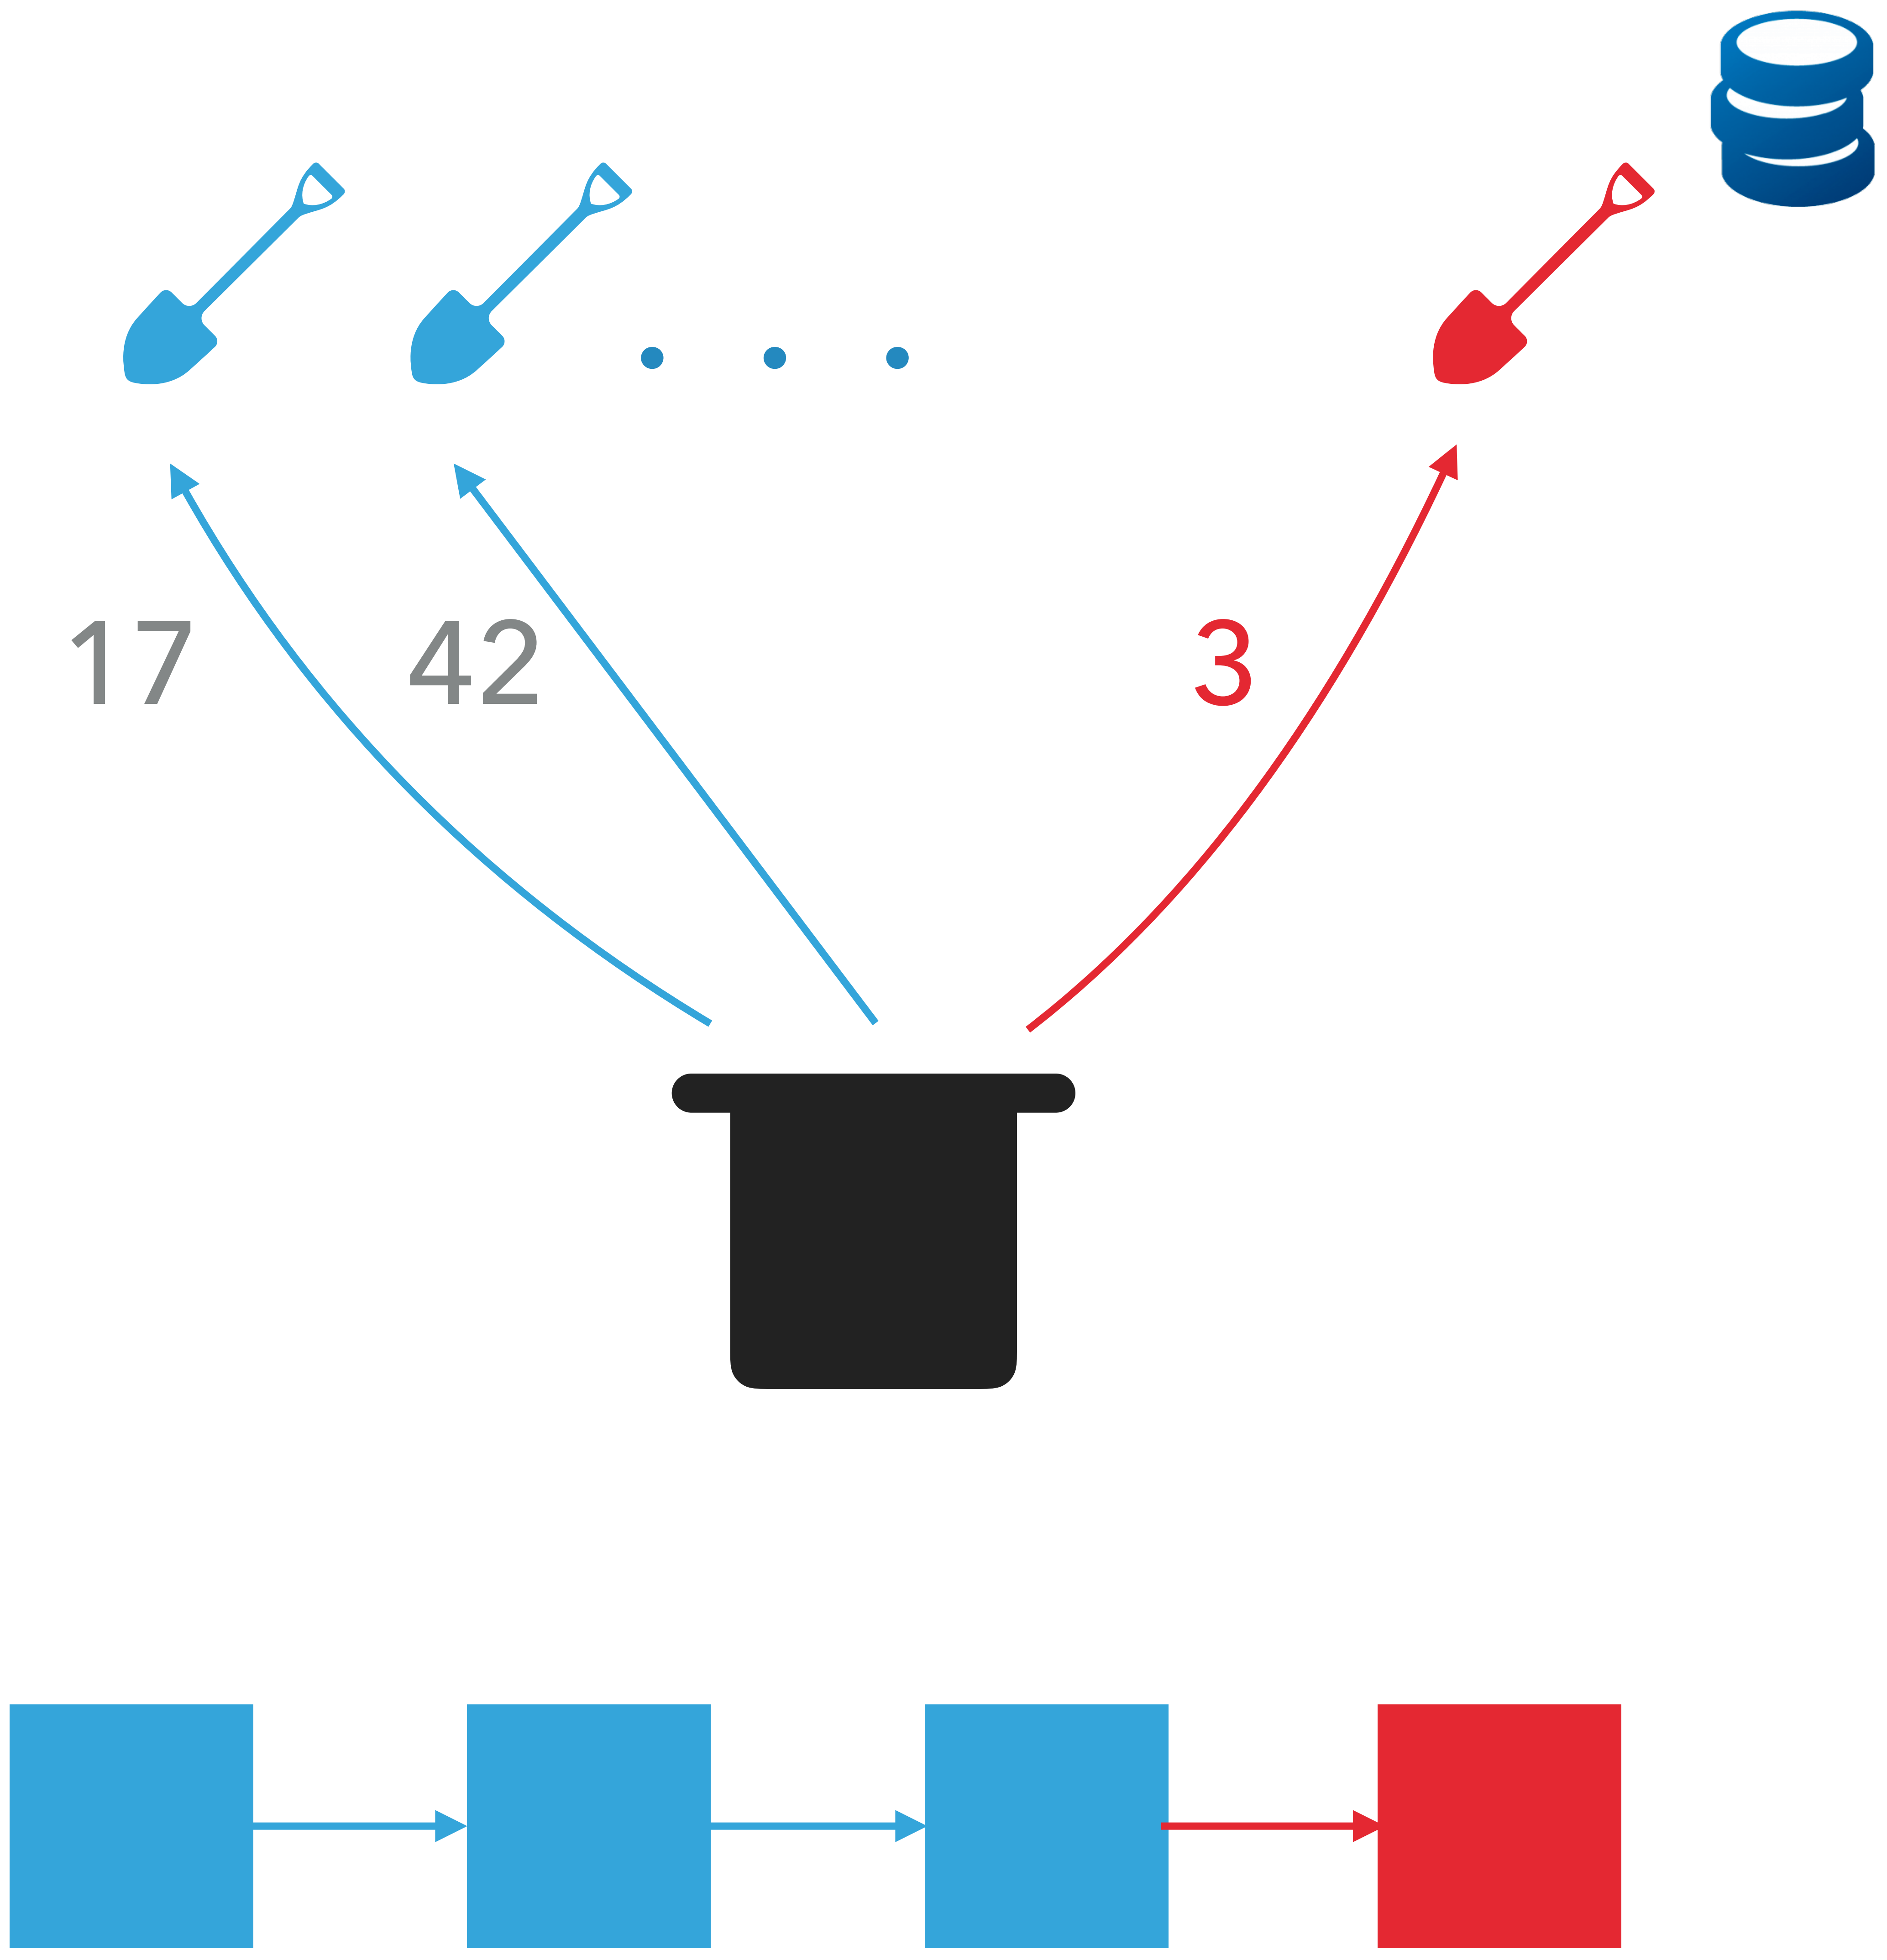

Mining is a Lottery

- Miners “draw” numbers until they cross threshold 5
- Each draw “costs” a hash
- First to cross threshold wins
- Winner receives a reward and proposes a block
- Game repeats

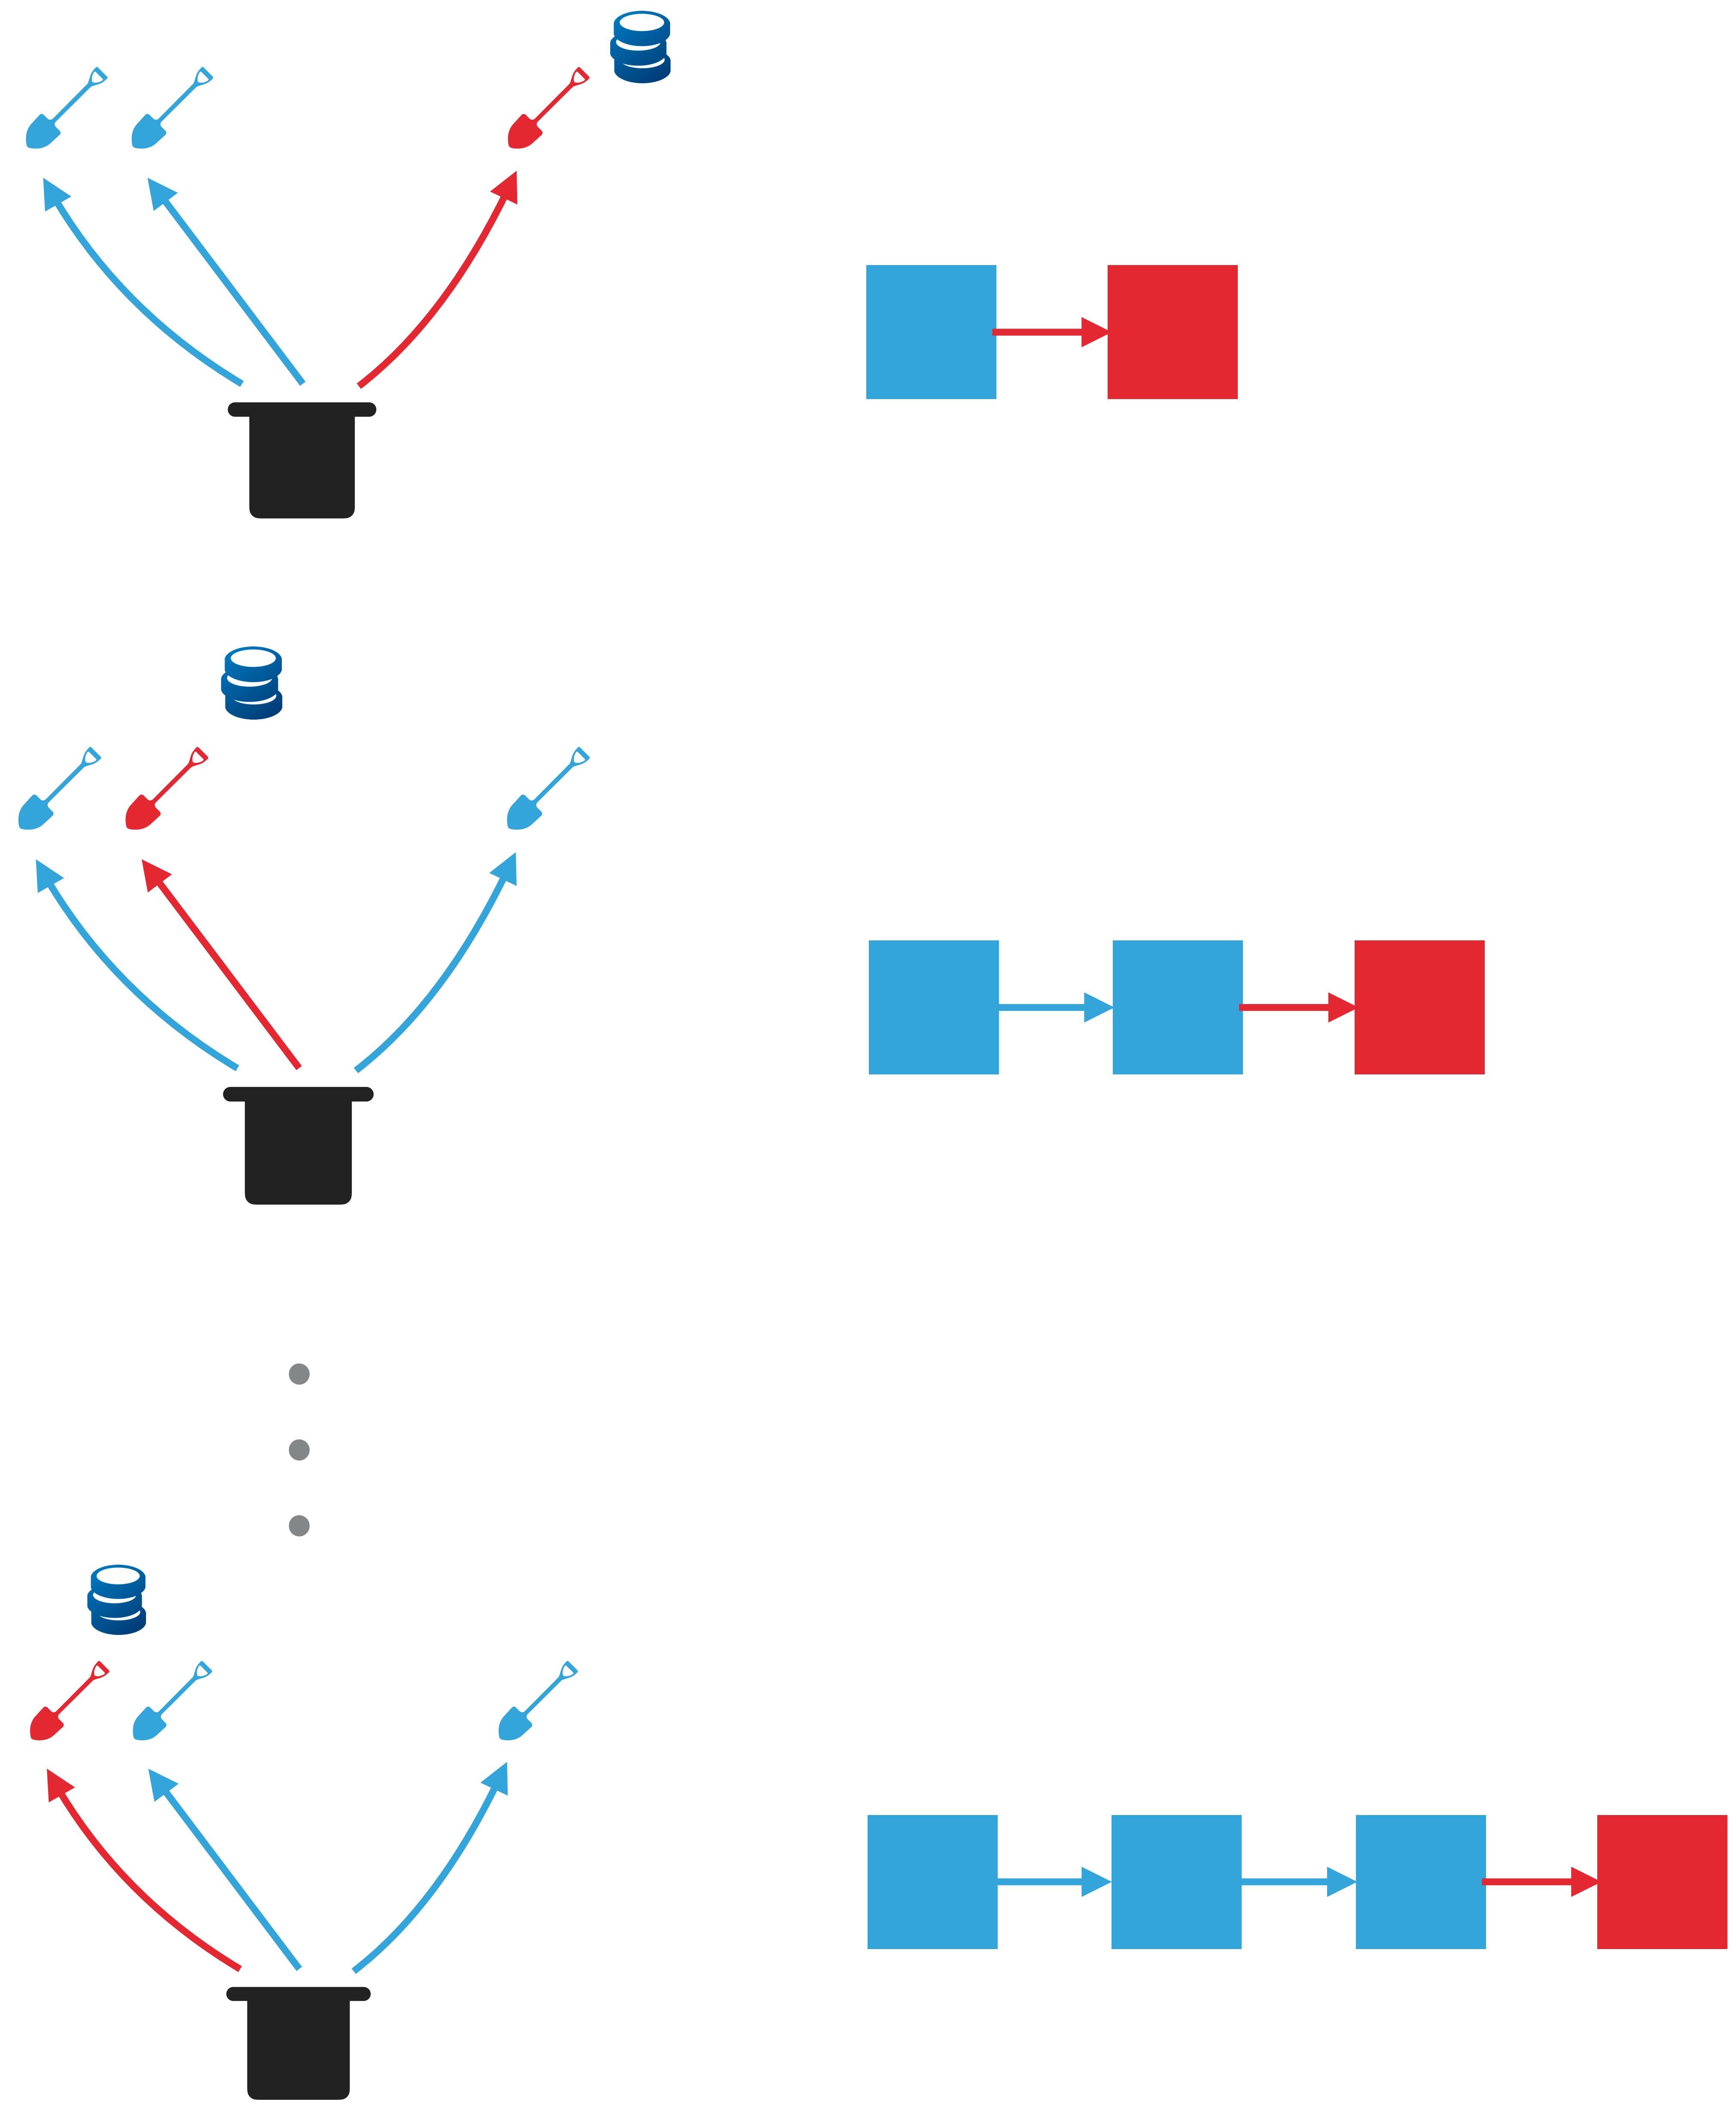

Mining statistics

- Time to draw below threshold is approximately $\texttt{Expon}\left(\frac{T}{q}\right)$
- 20% miner expects to take 4 times as long to mine a block as others

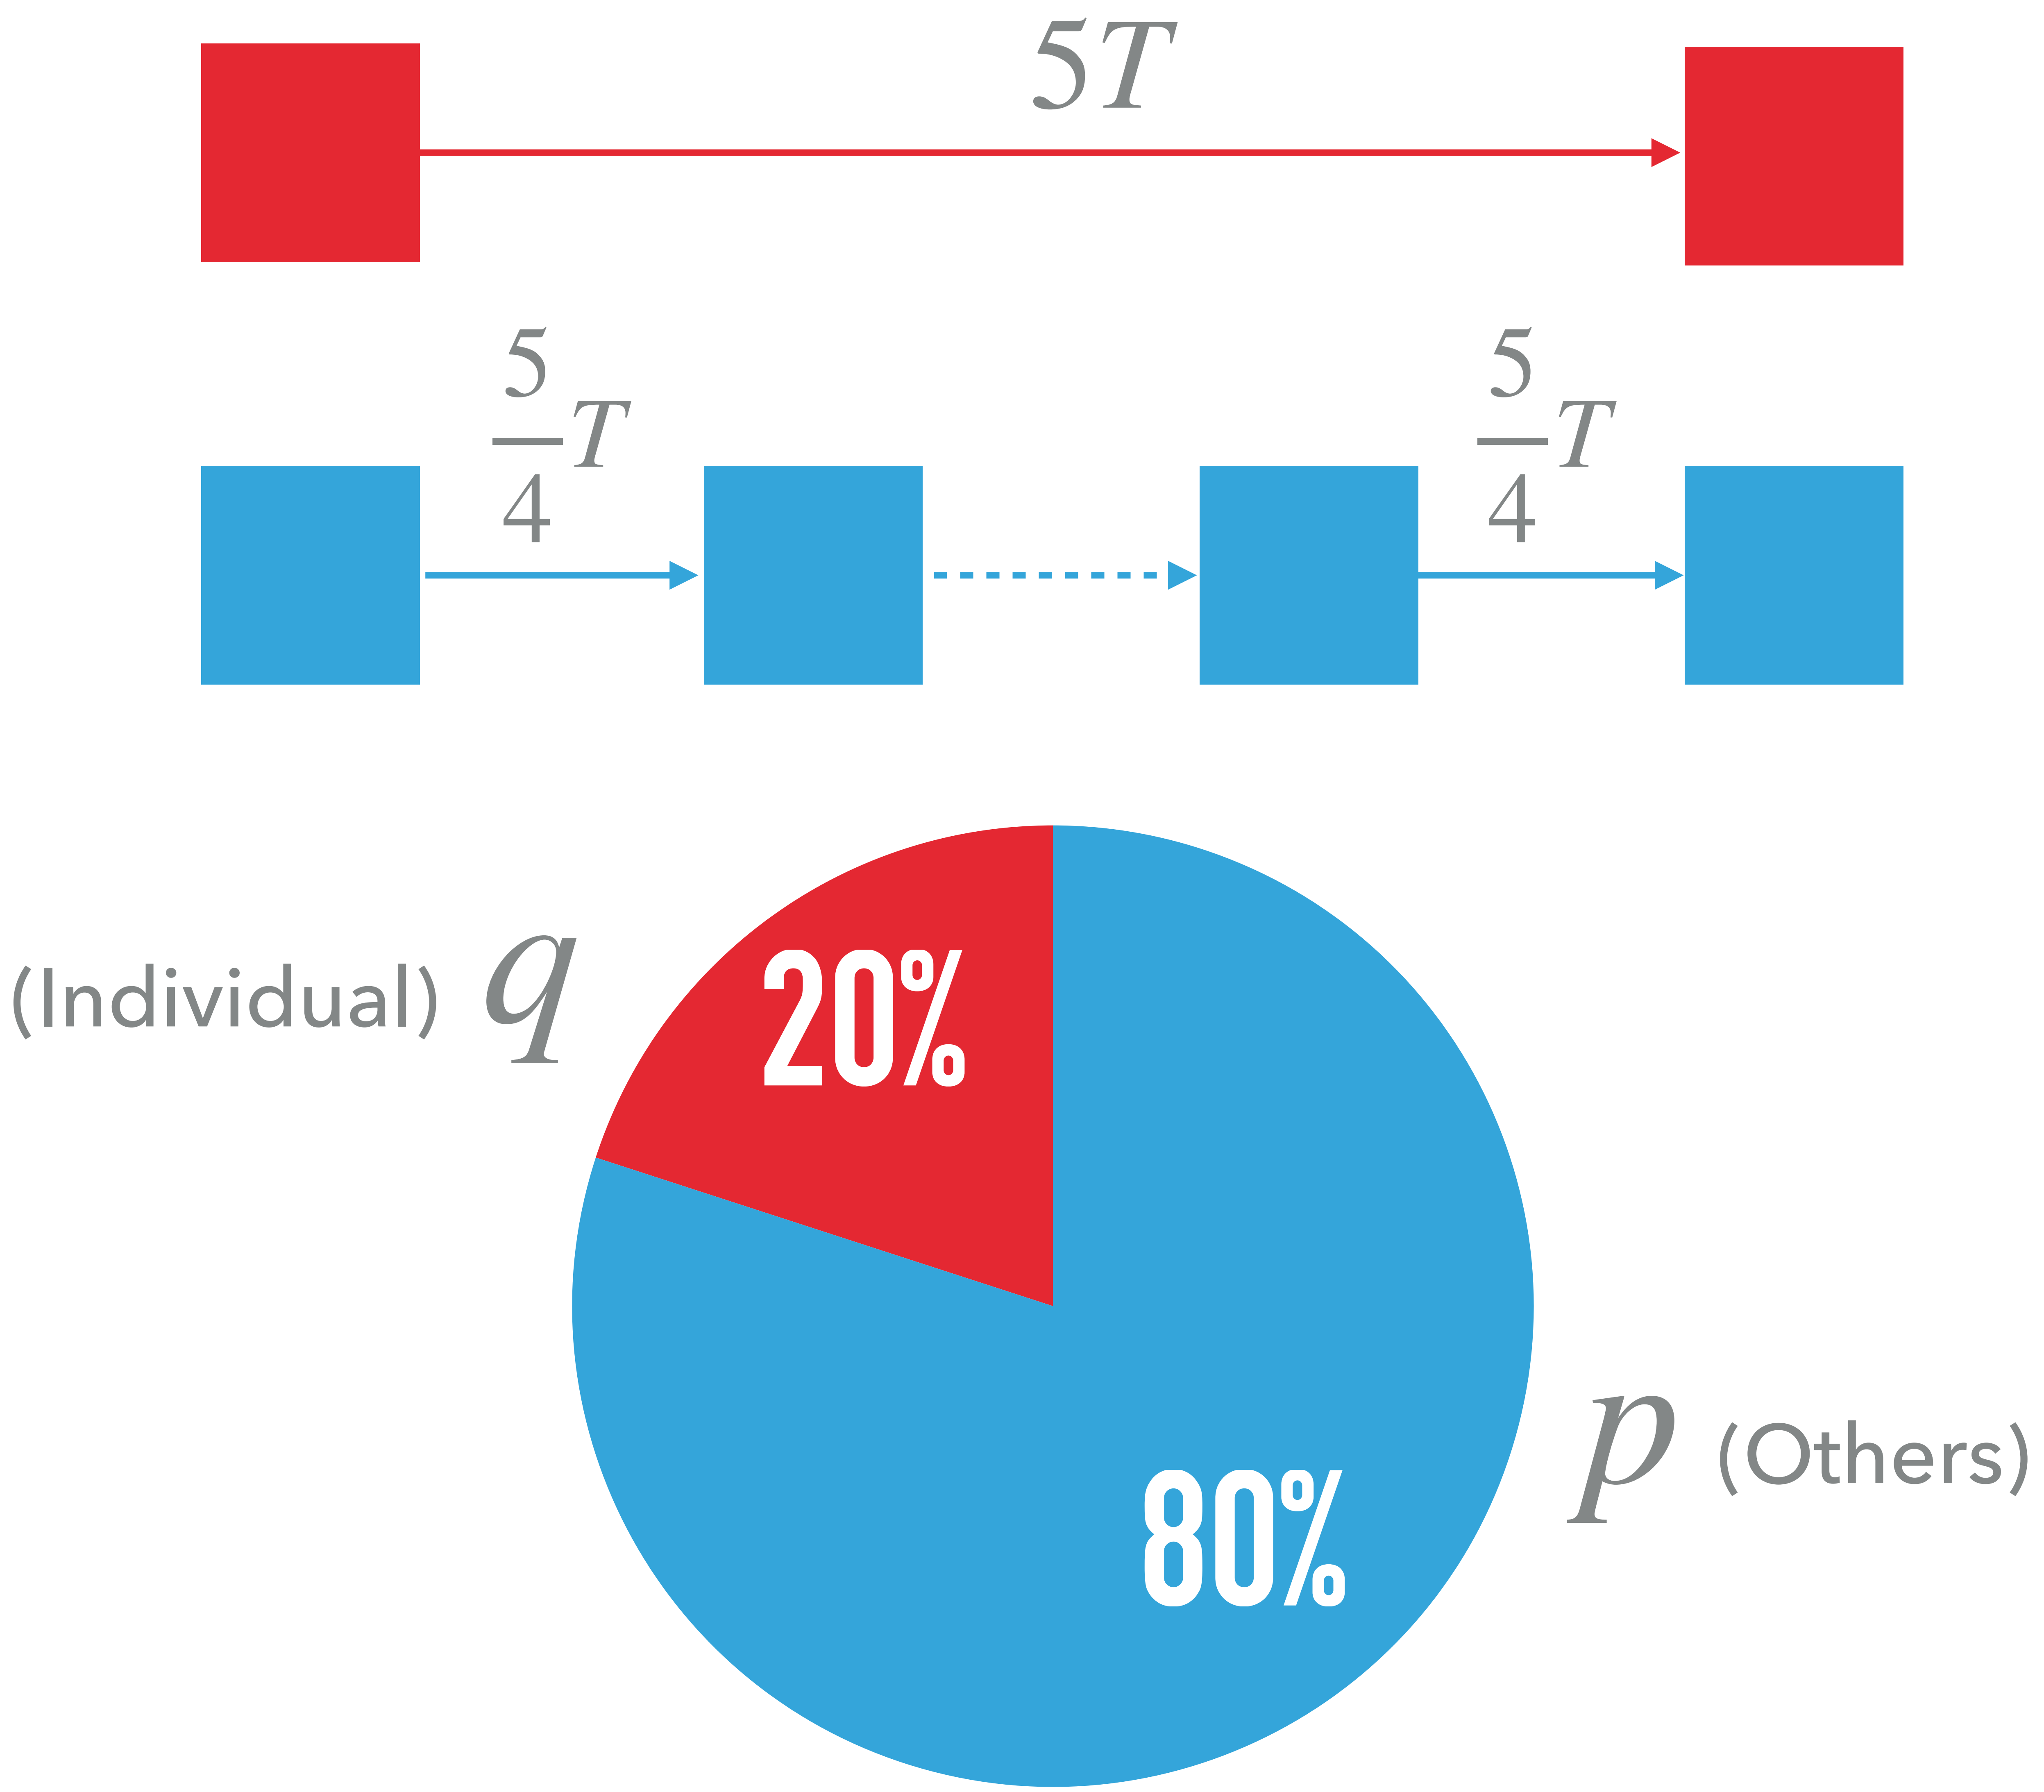

Double-spending Attack

- Alice trades car for 1 BTC
- Transaction appears in block 1
- Assumes majority are mining chain
- Alice knows about law of large numbers
- Goods are released only once payment has z "confirmations"

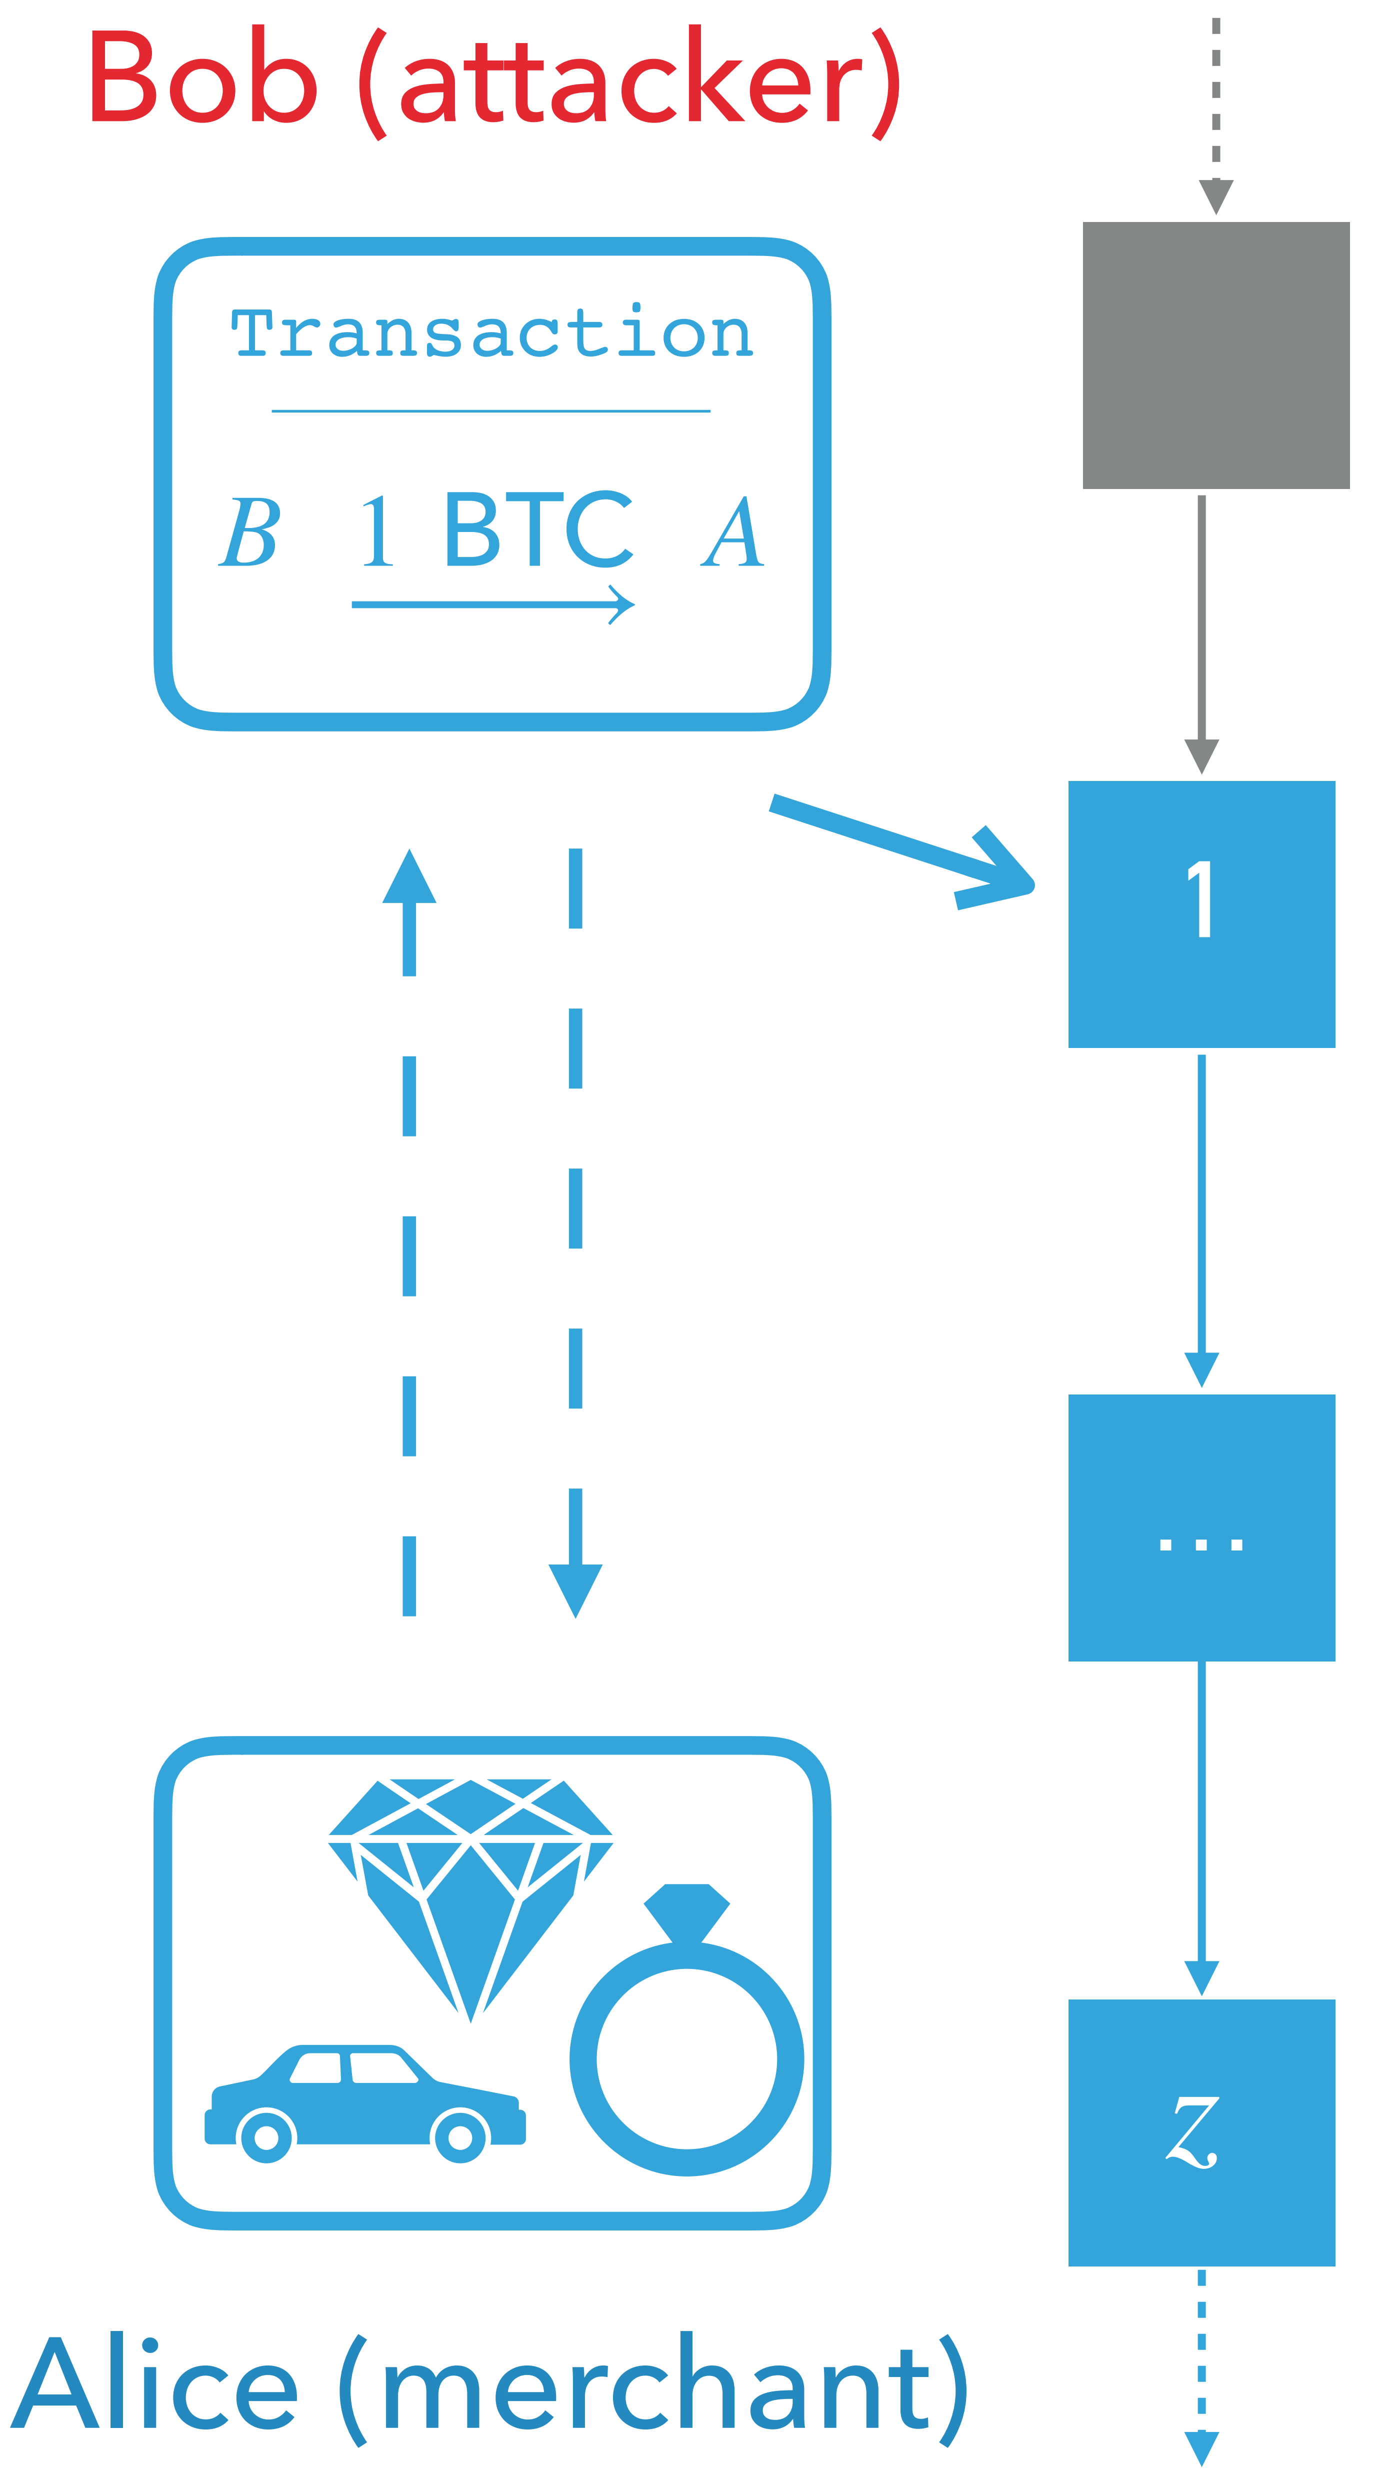

Double-spending Attack

- Bob steals goods if red chain grows longer than blue
- Relies on high variance of the exponential distribution
- **Goods worth more than cost of attack?**

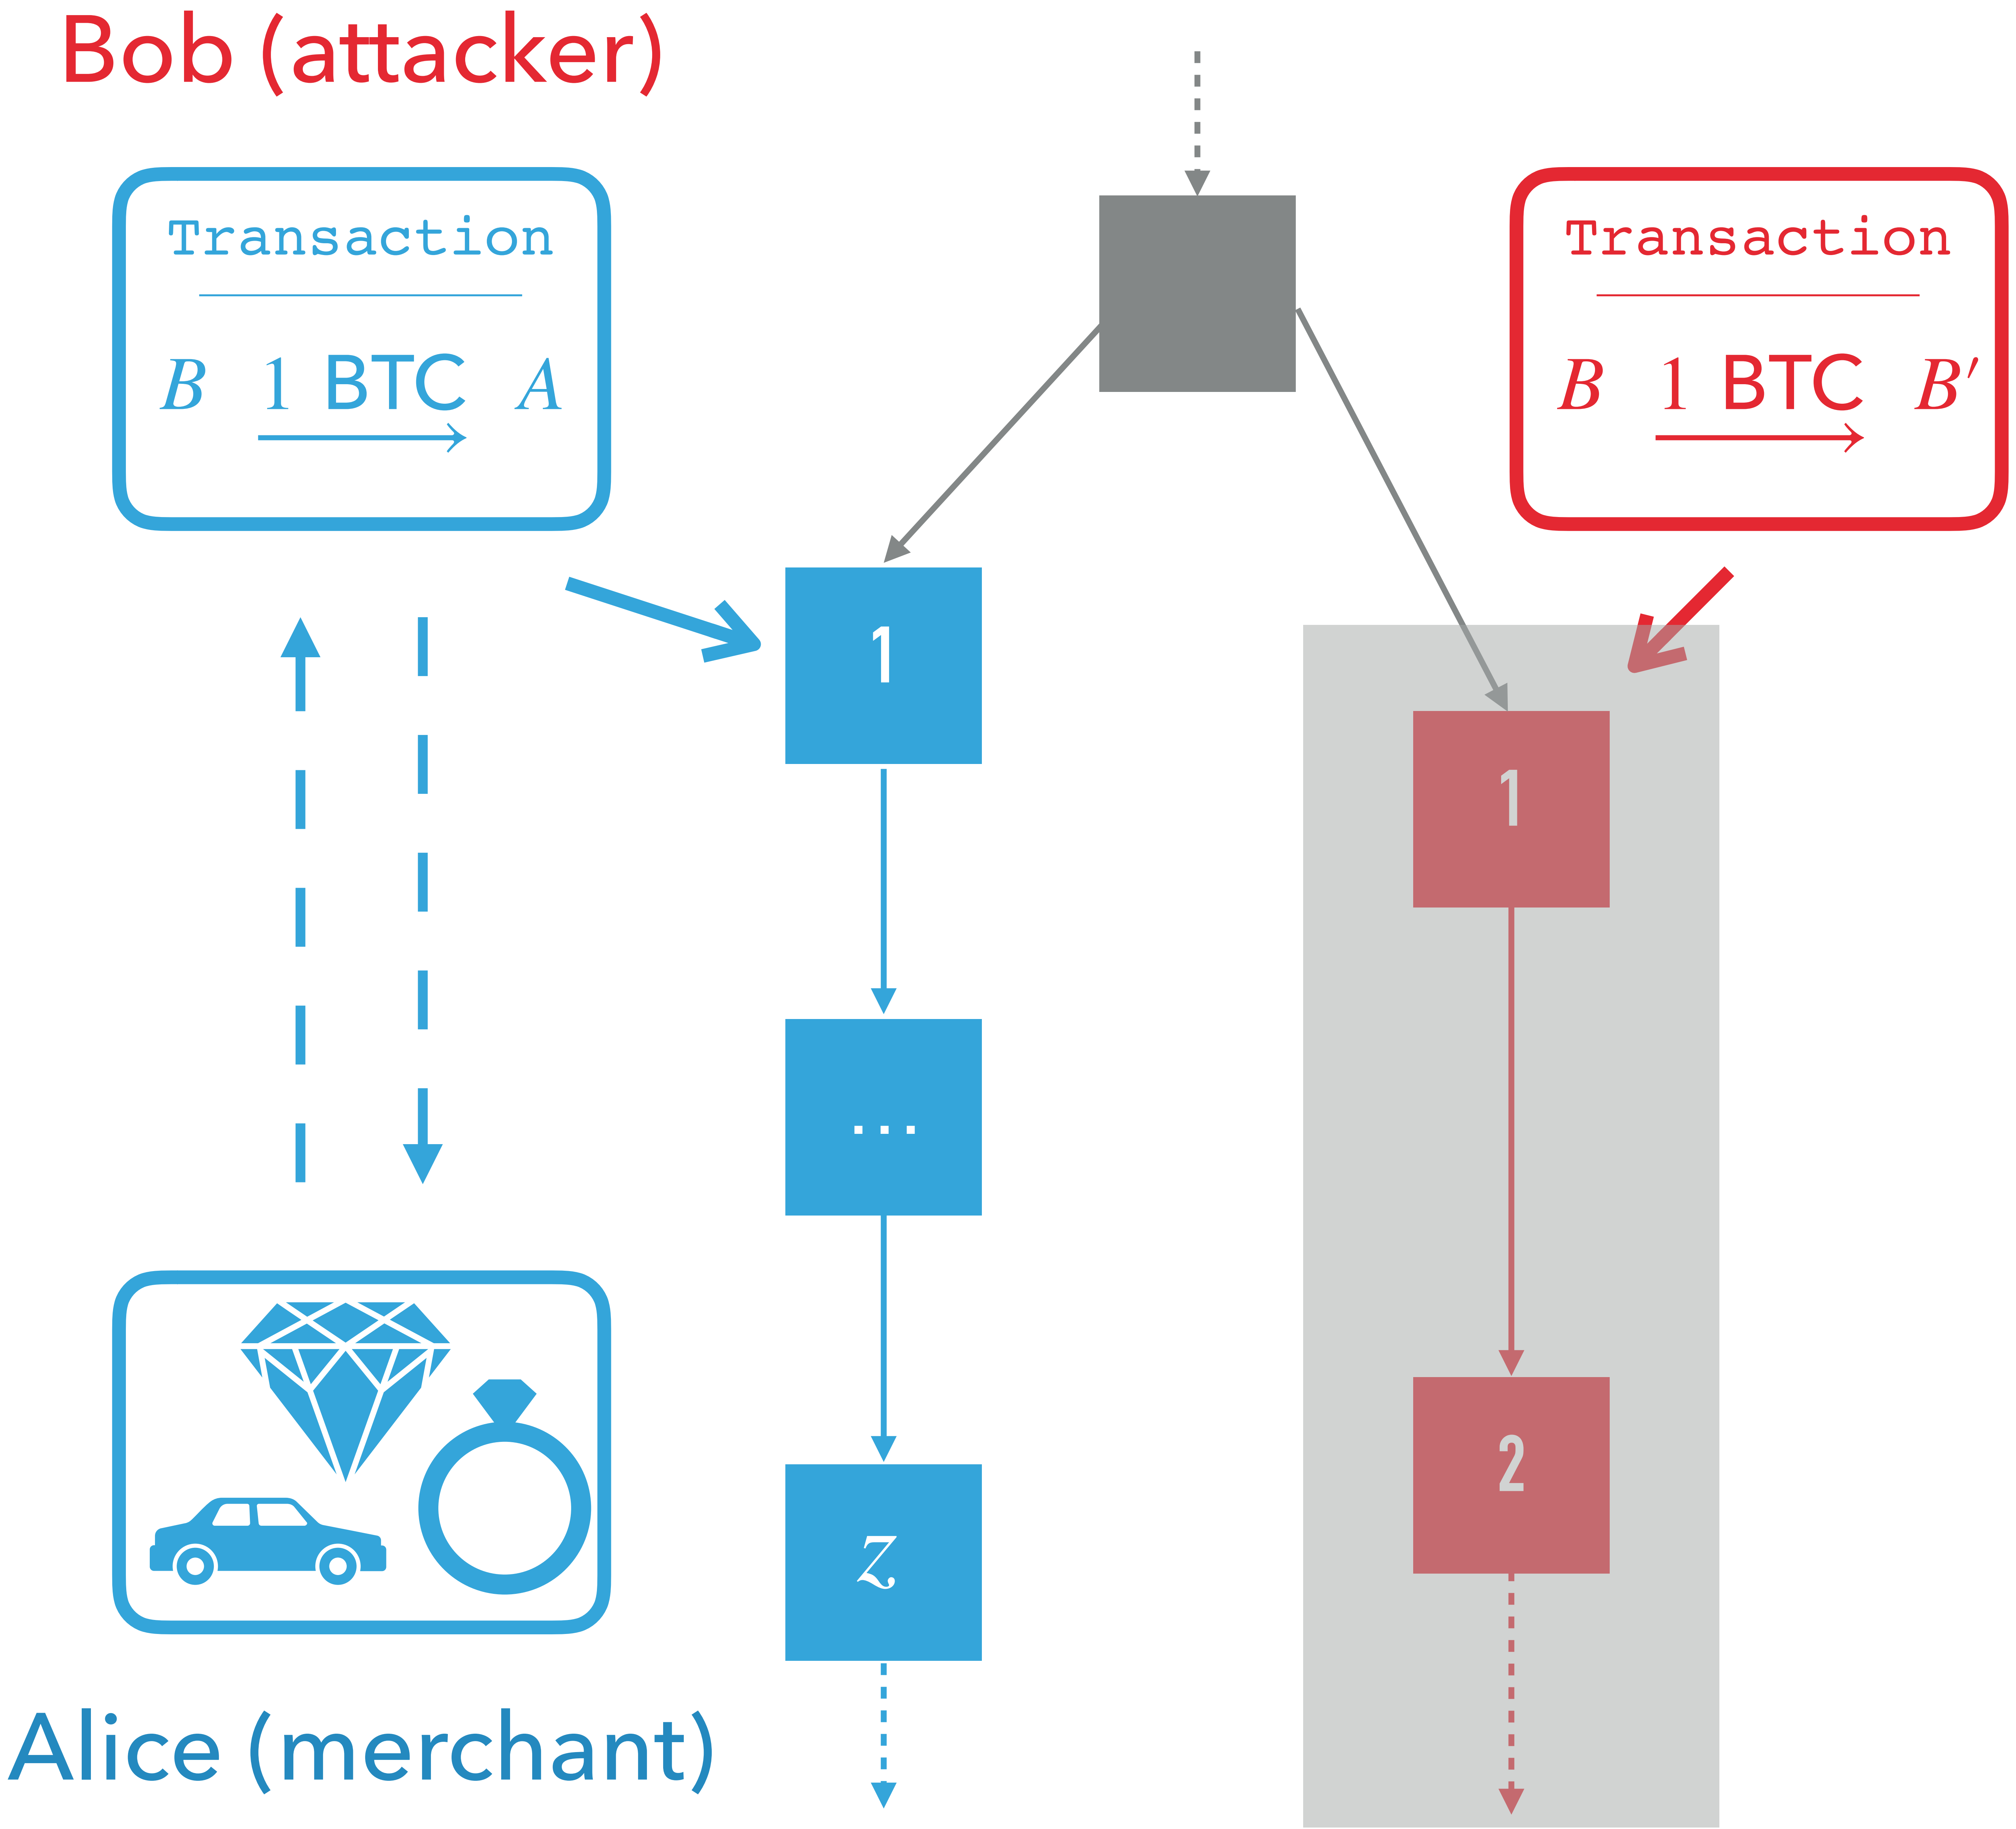

Attack Success Probability

- Attacker needs to get ahead by at least one block sometime after the first z blocks
- Even a 20% miner has 5% chance of winning after 6 blocks

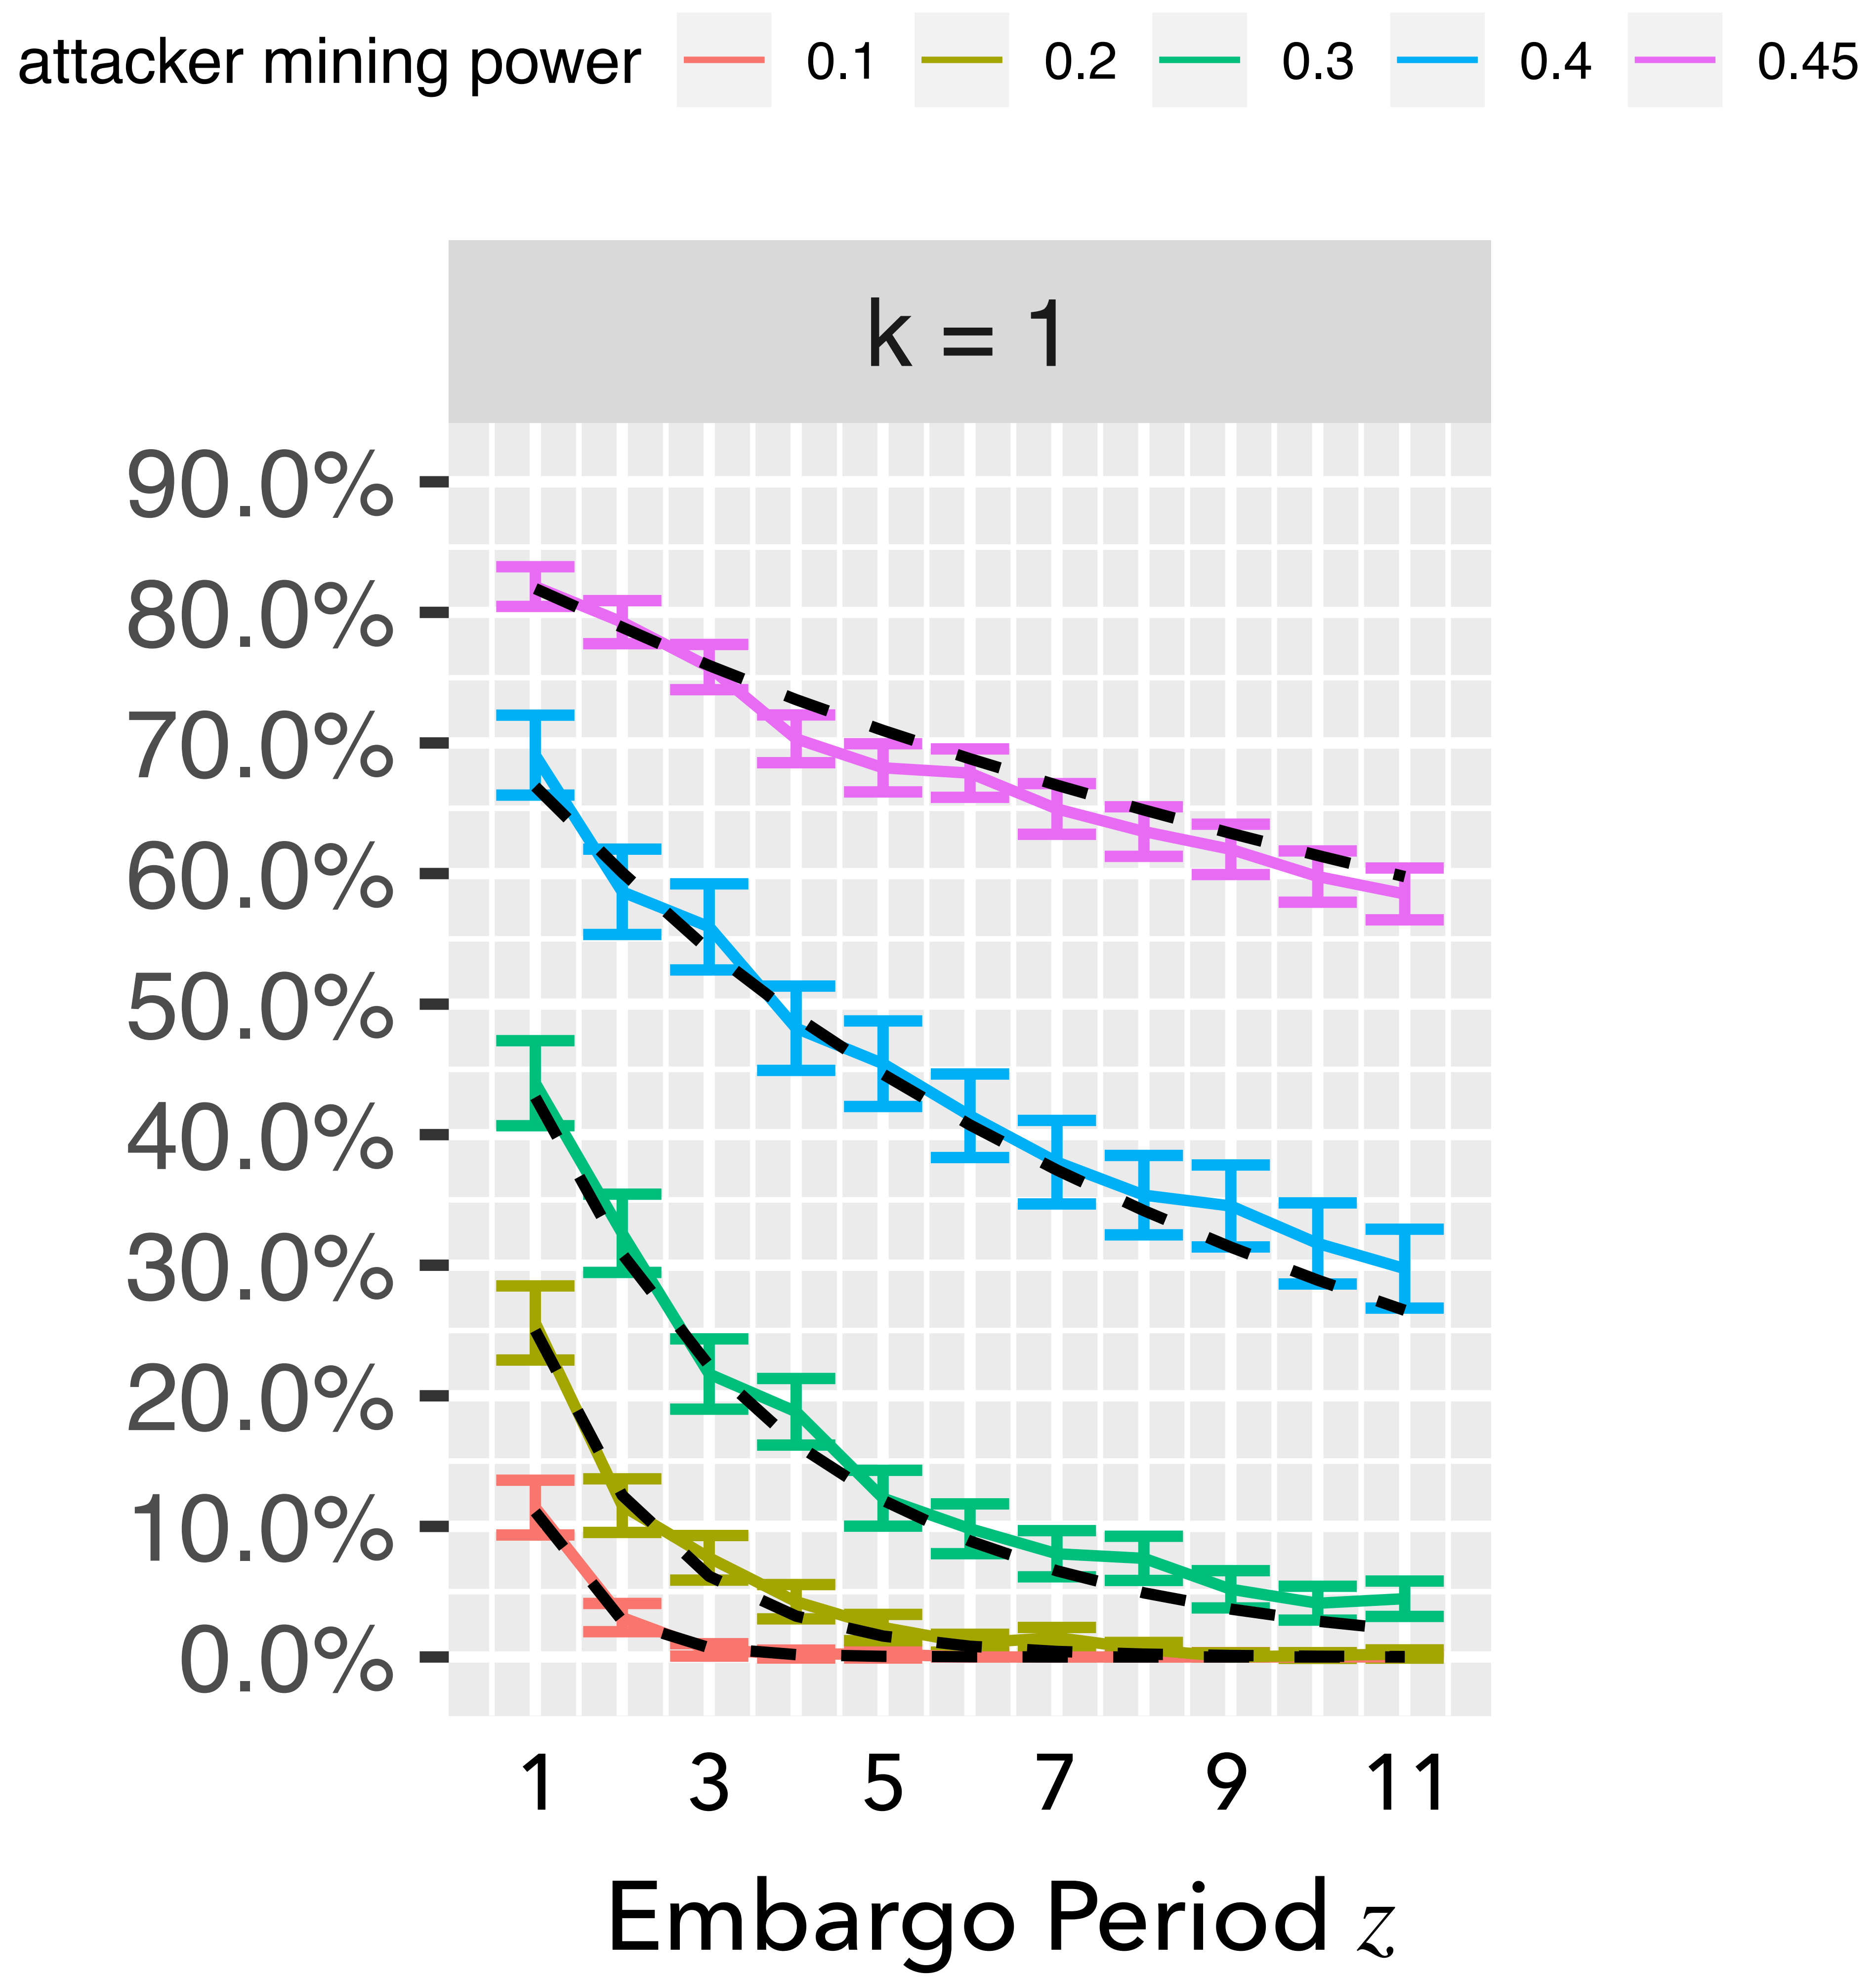

Bobtail Protocol Details

- Assemble a block containing transactions
- Hash header as usual to generate "proofs"
- Disseminate proofs that are "low enough" to neighbors
- Maintain queue of lowest k proofs
- Assemble k proofs whose mean is below t
- Each proof miner receives reward

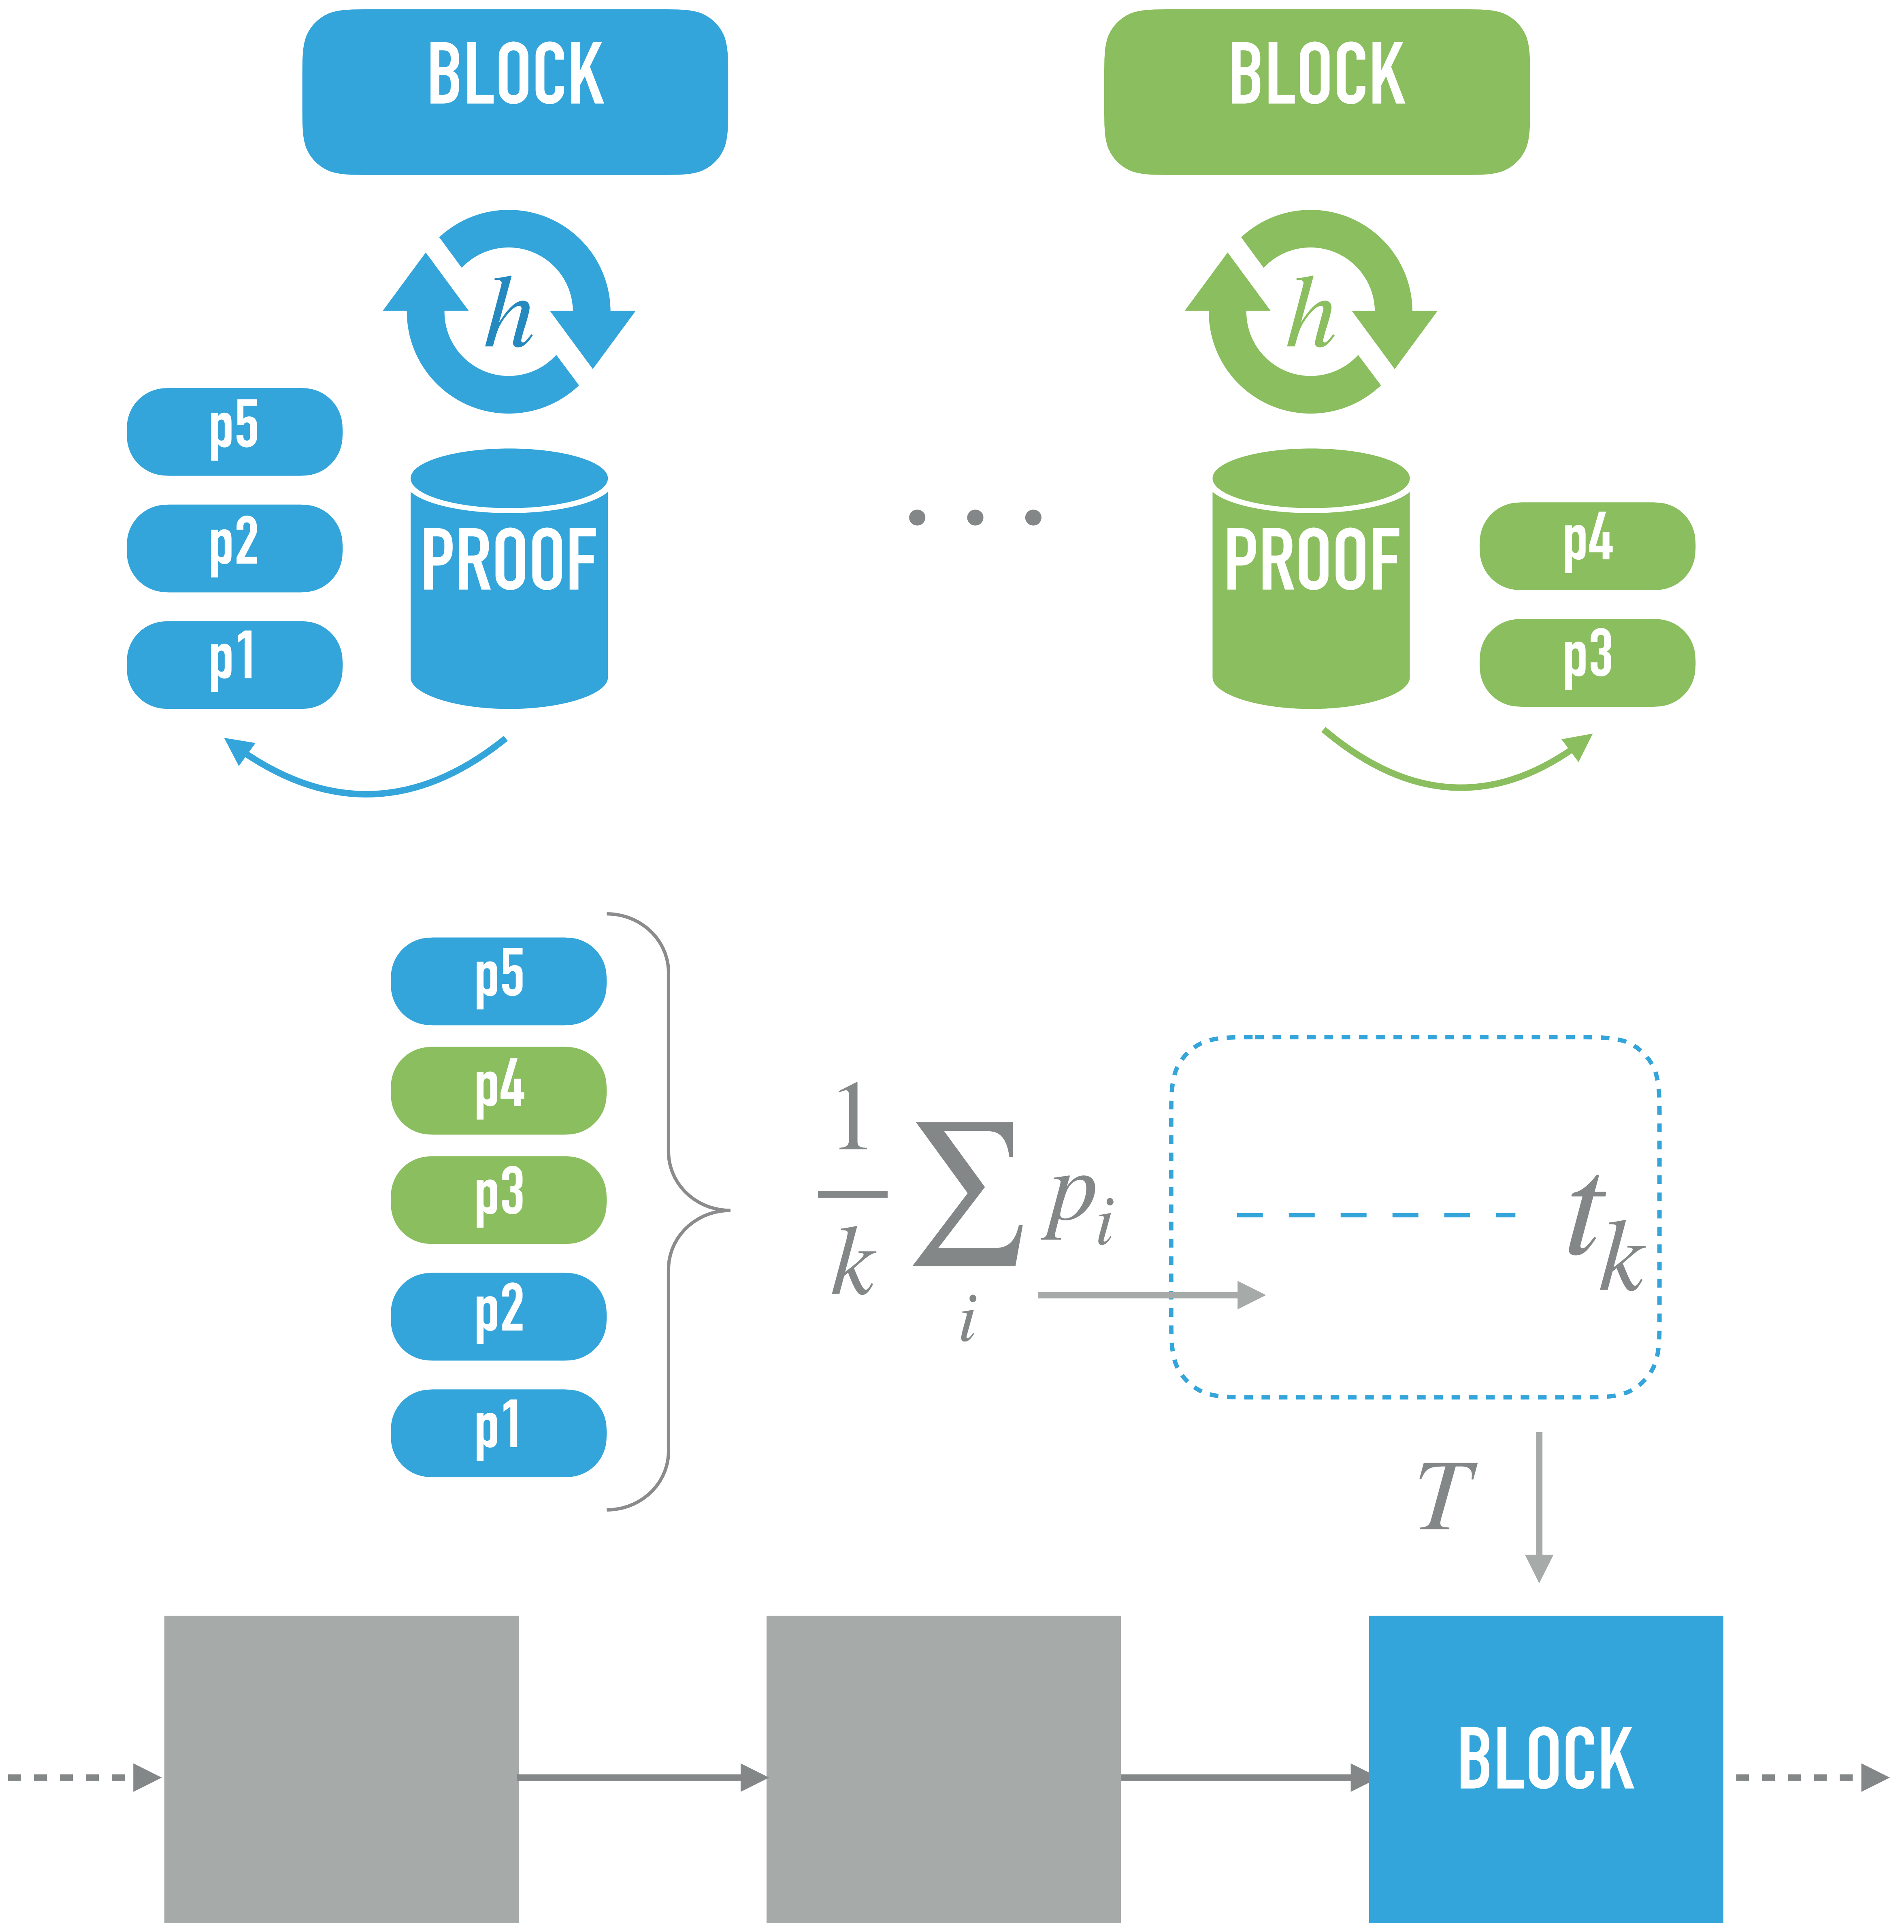

New Lottery: Bobtail

- Miners draw numbers until the average of any 2 cross threshold 5

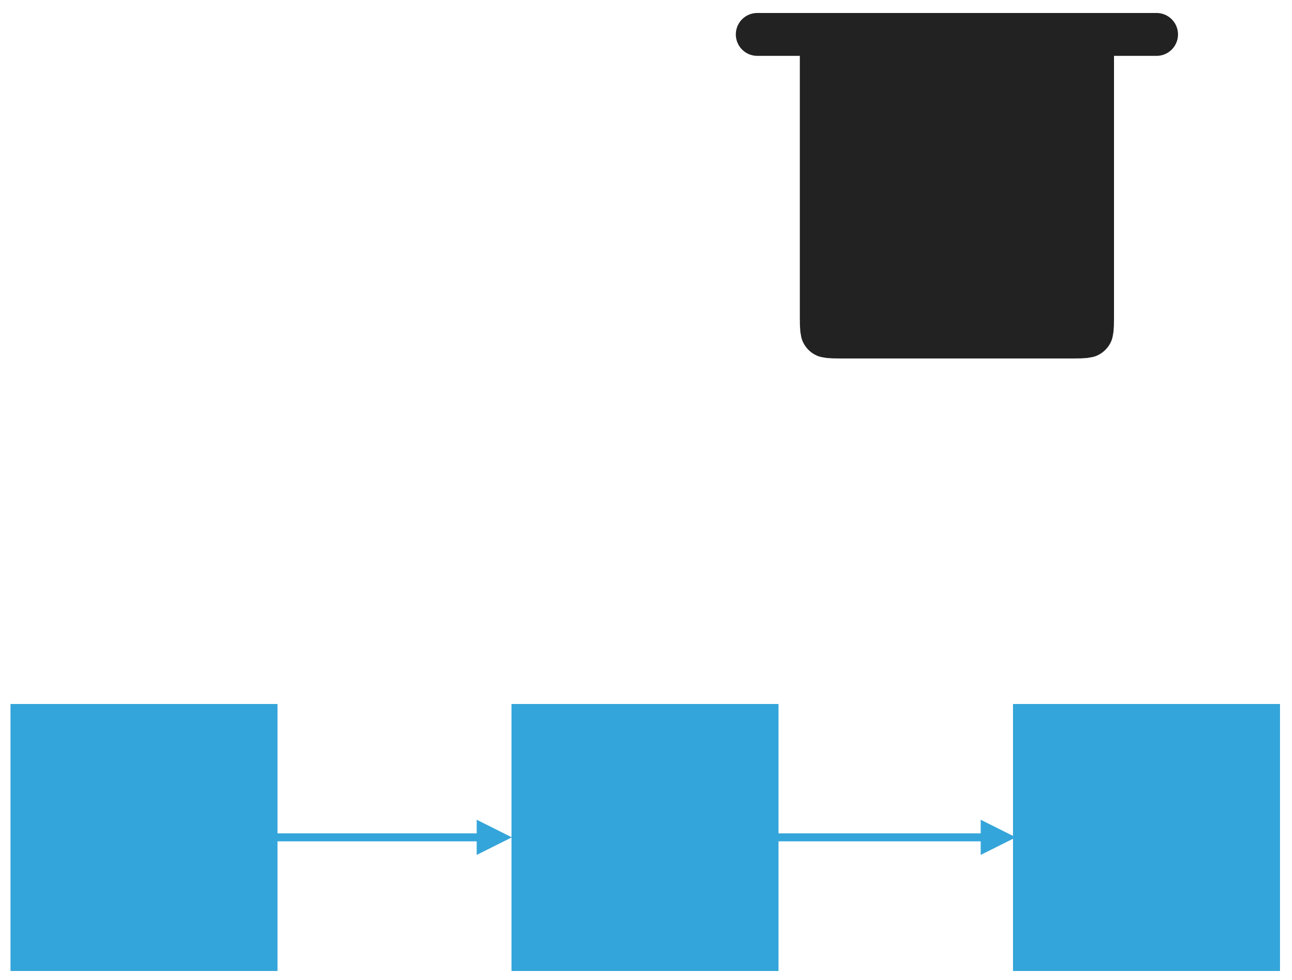

New Lottery: Bobtail

- Miners draw numbers until the average of any 2 cross threshold 5
- Each draw still "costs" a hash

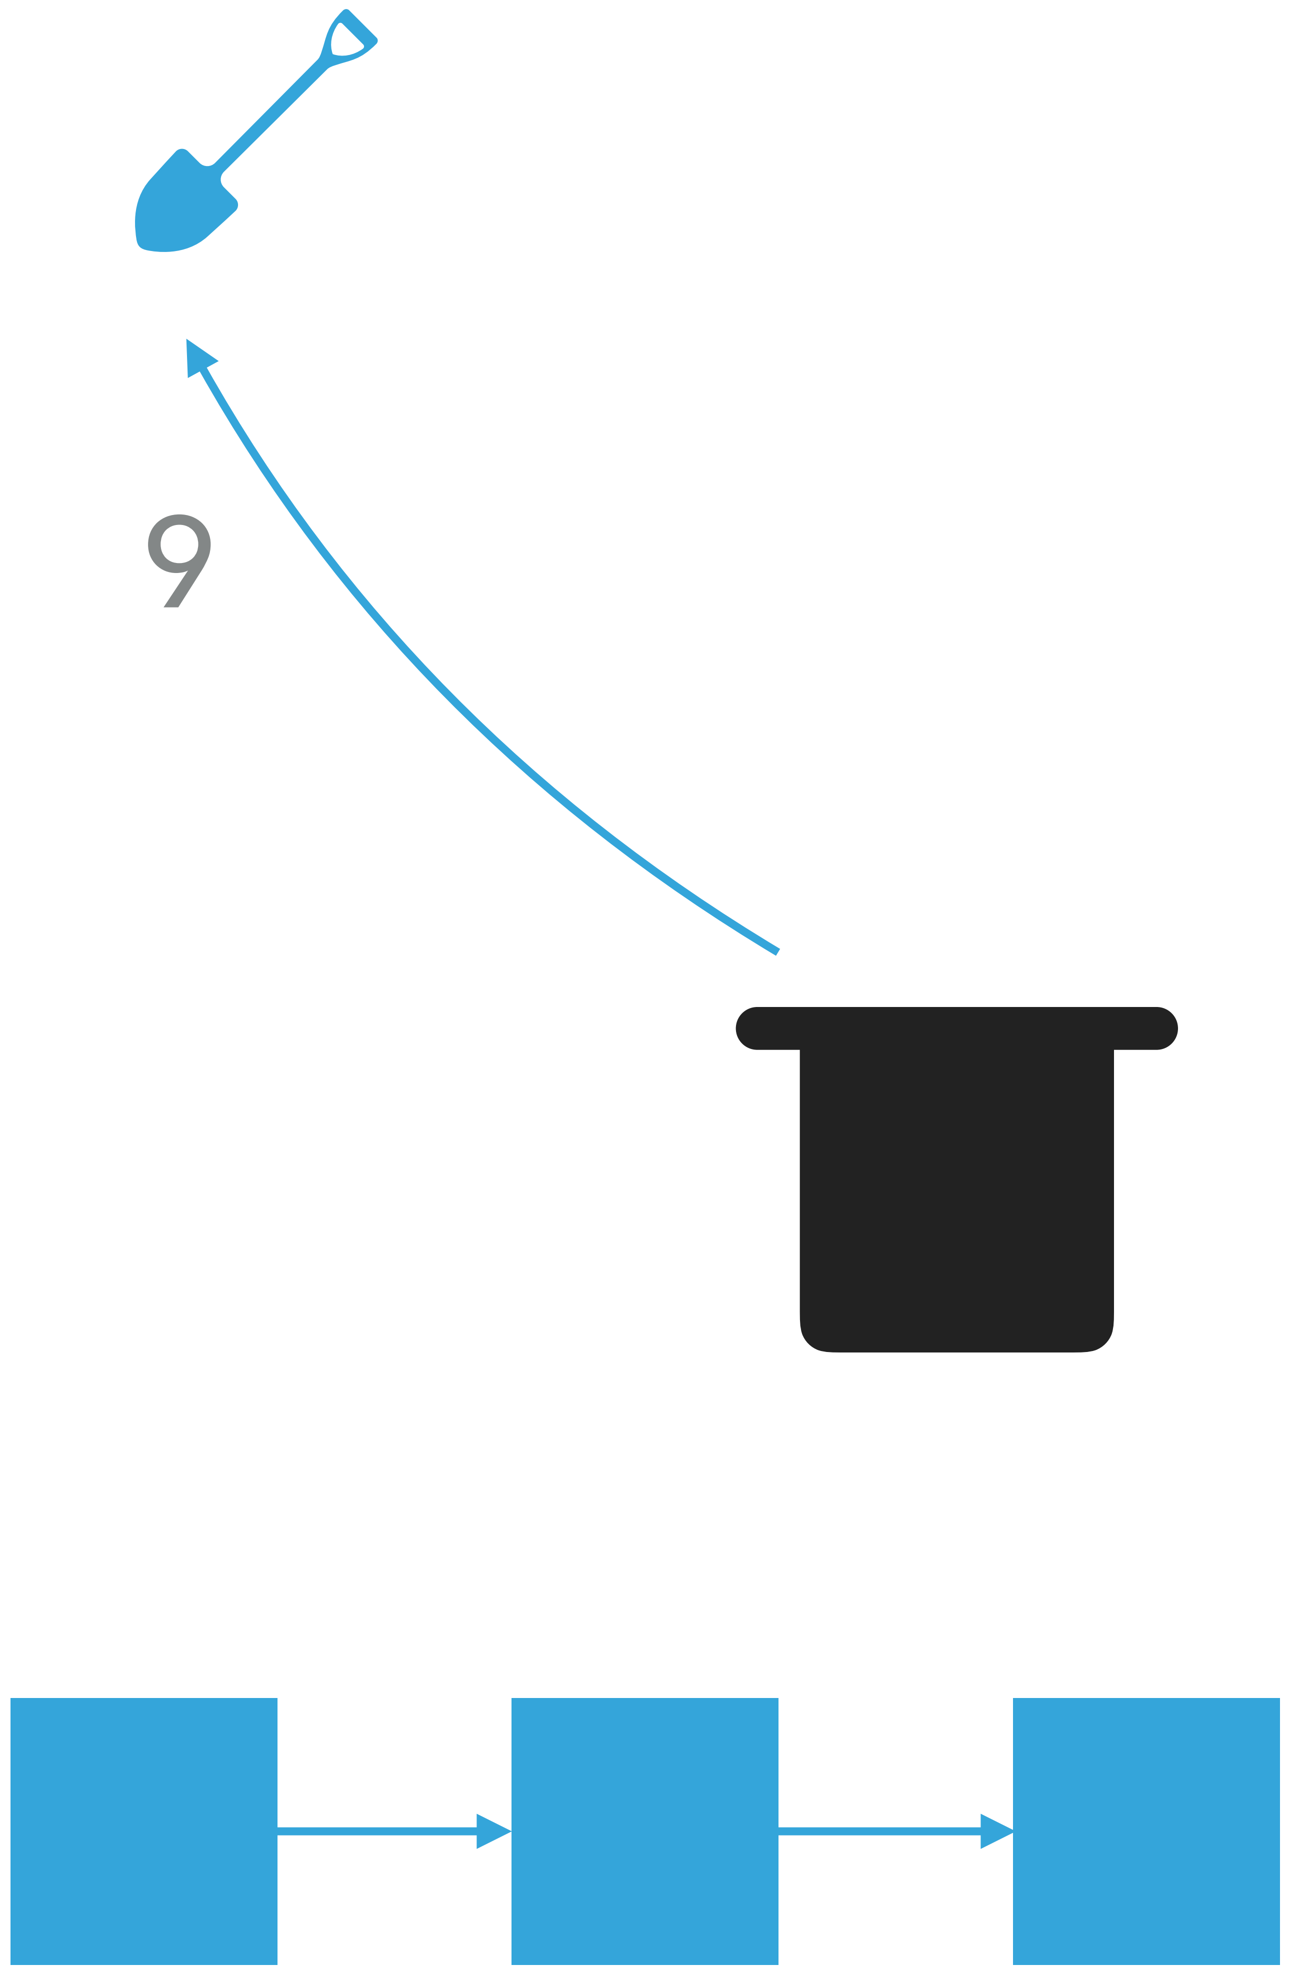

New Lottery: Bobtail

- Miners draw numbers until the average of any 2 cross threshold 5
- Each draw still "costs" a hash

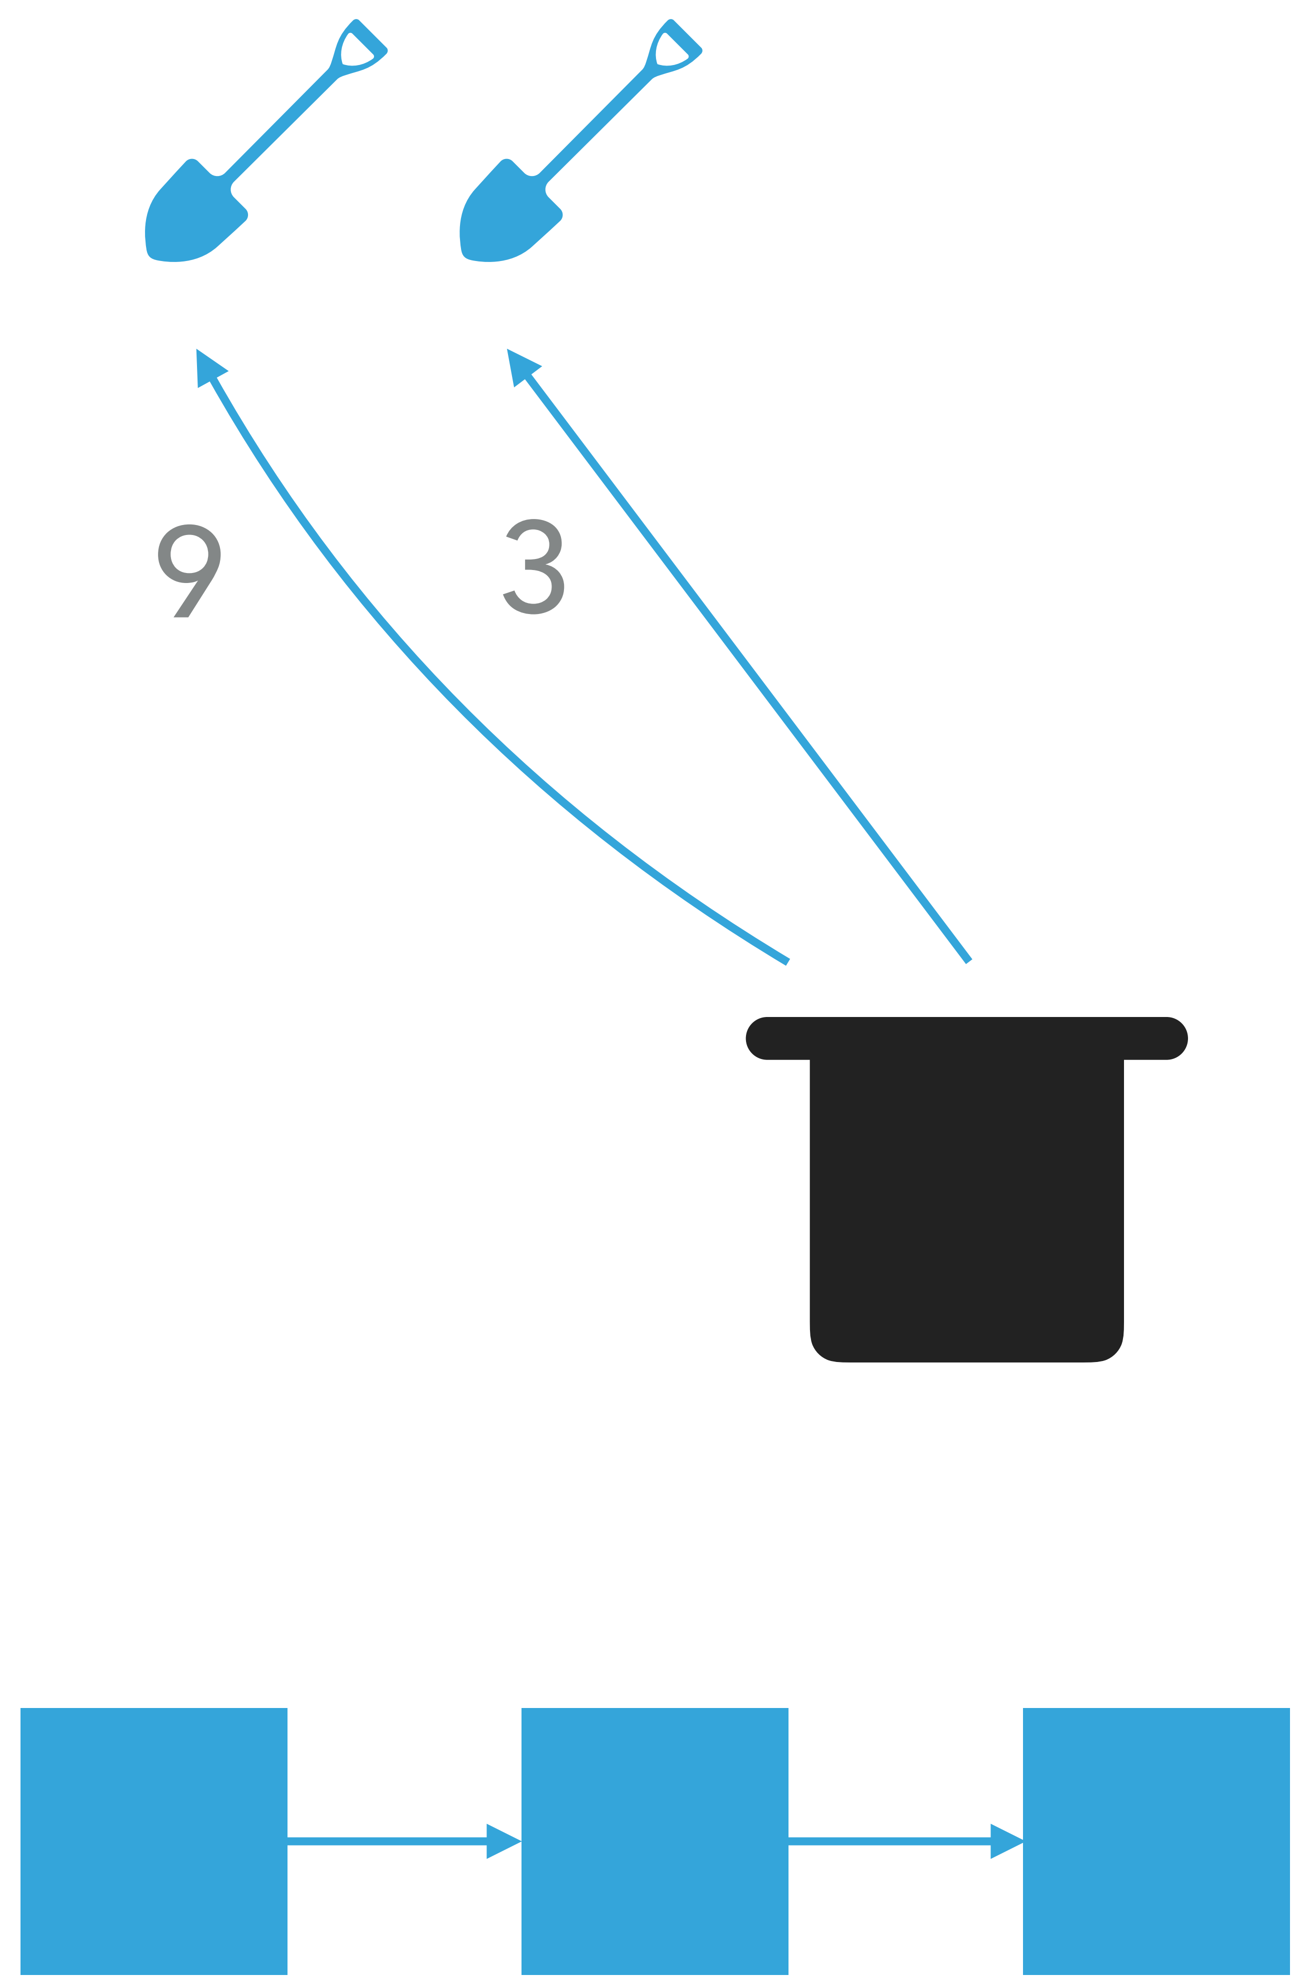

New Lottery: Bobtail

- Miners draw numbers until the average of any 2 cross threshold 5
- Each draw still "costs" a hash

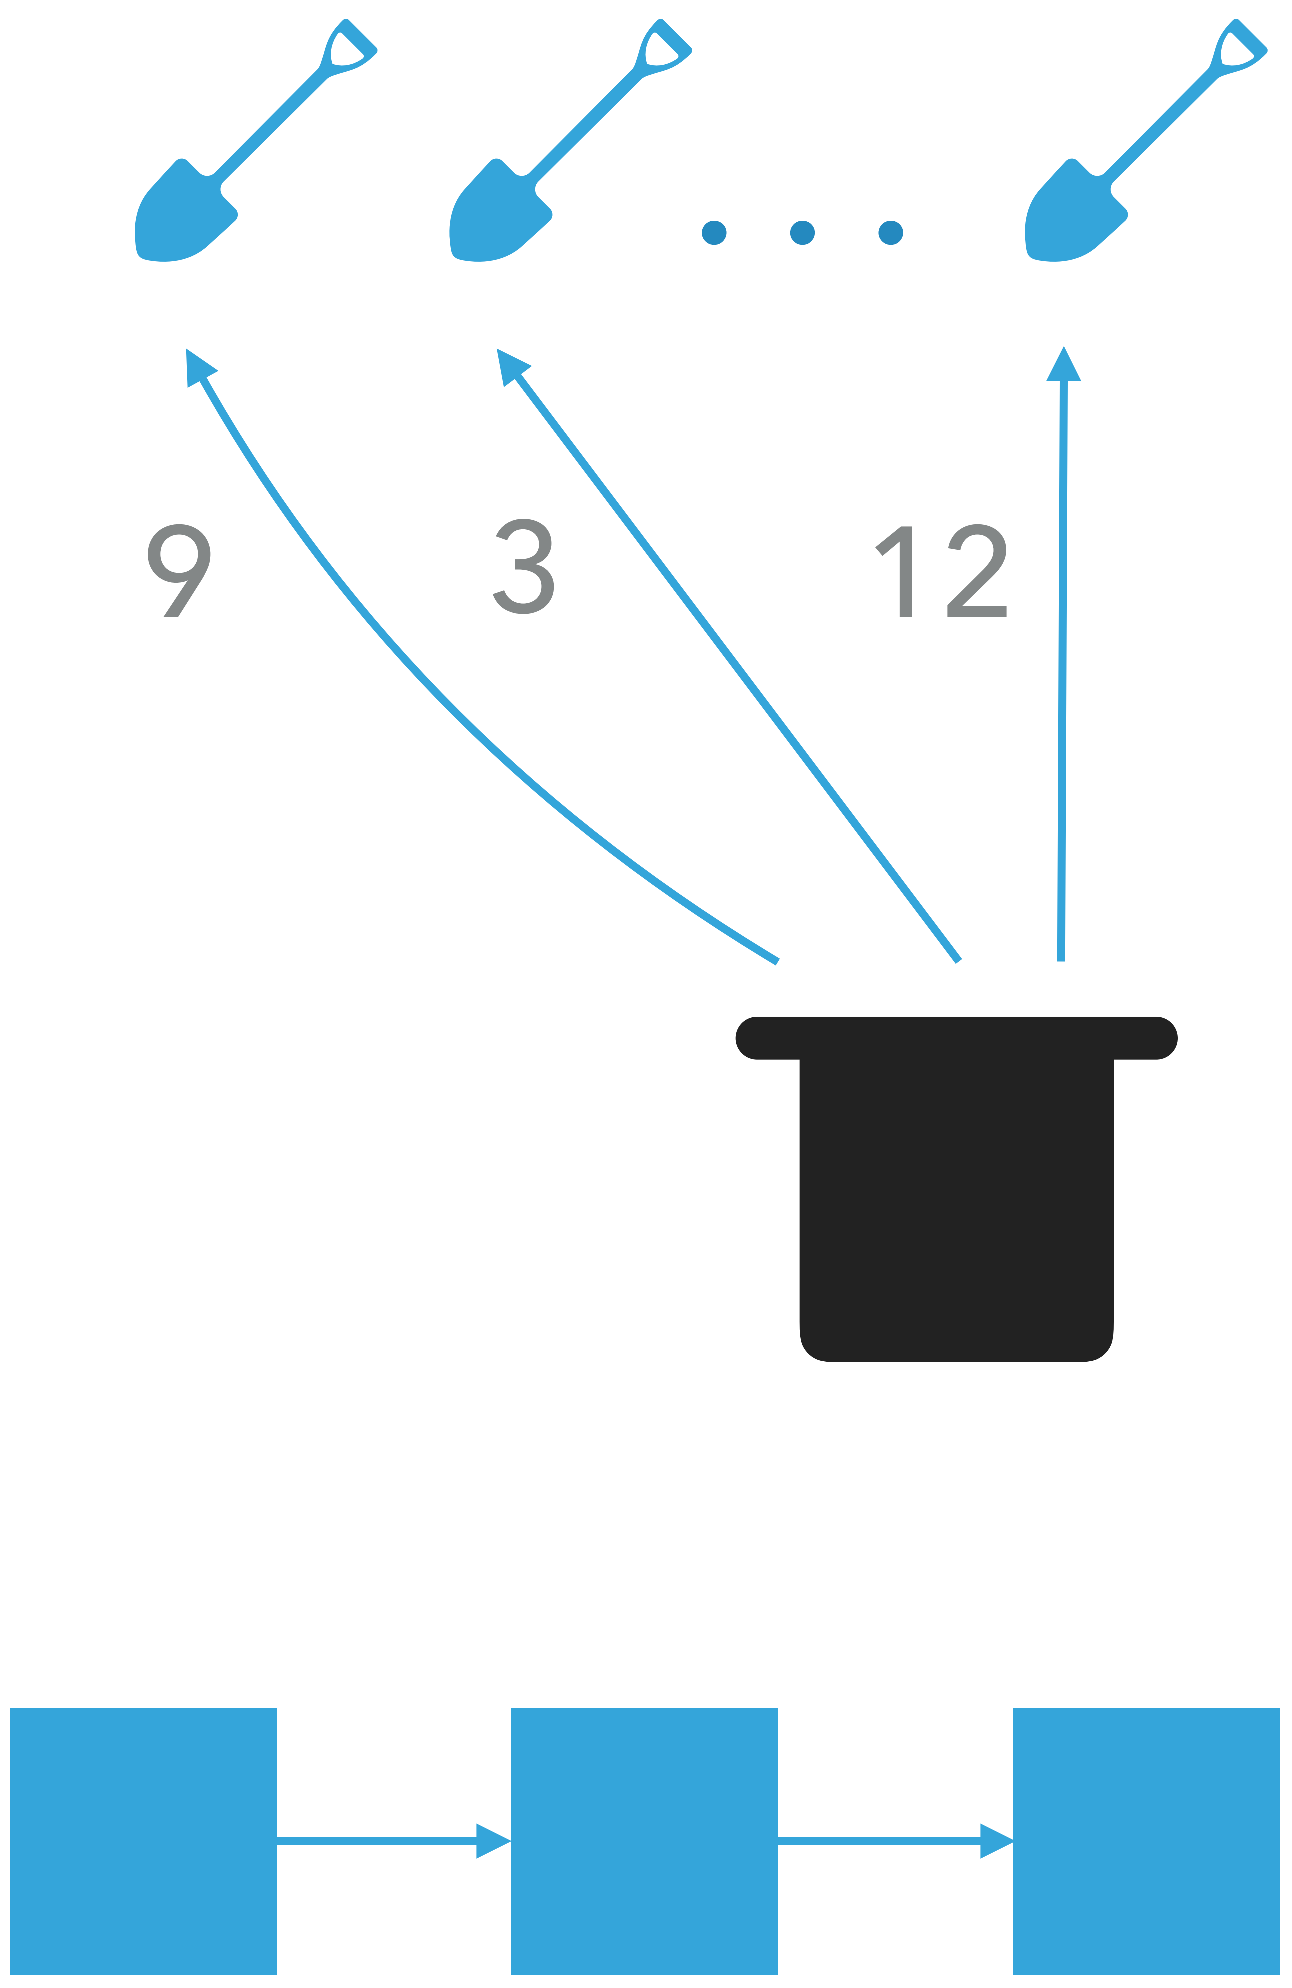

New Lottery: Bobtail

- Miners draw numbers until the average of any 2 cross threshold 5
- Each draw still “costs” a hash
- First 2 to cross threshold win
- Winners receive a reward and lowest proposes a block

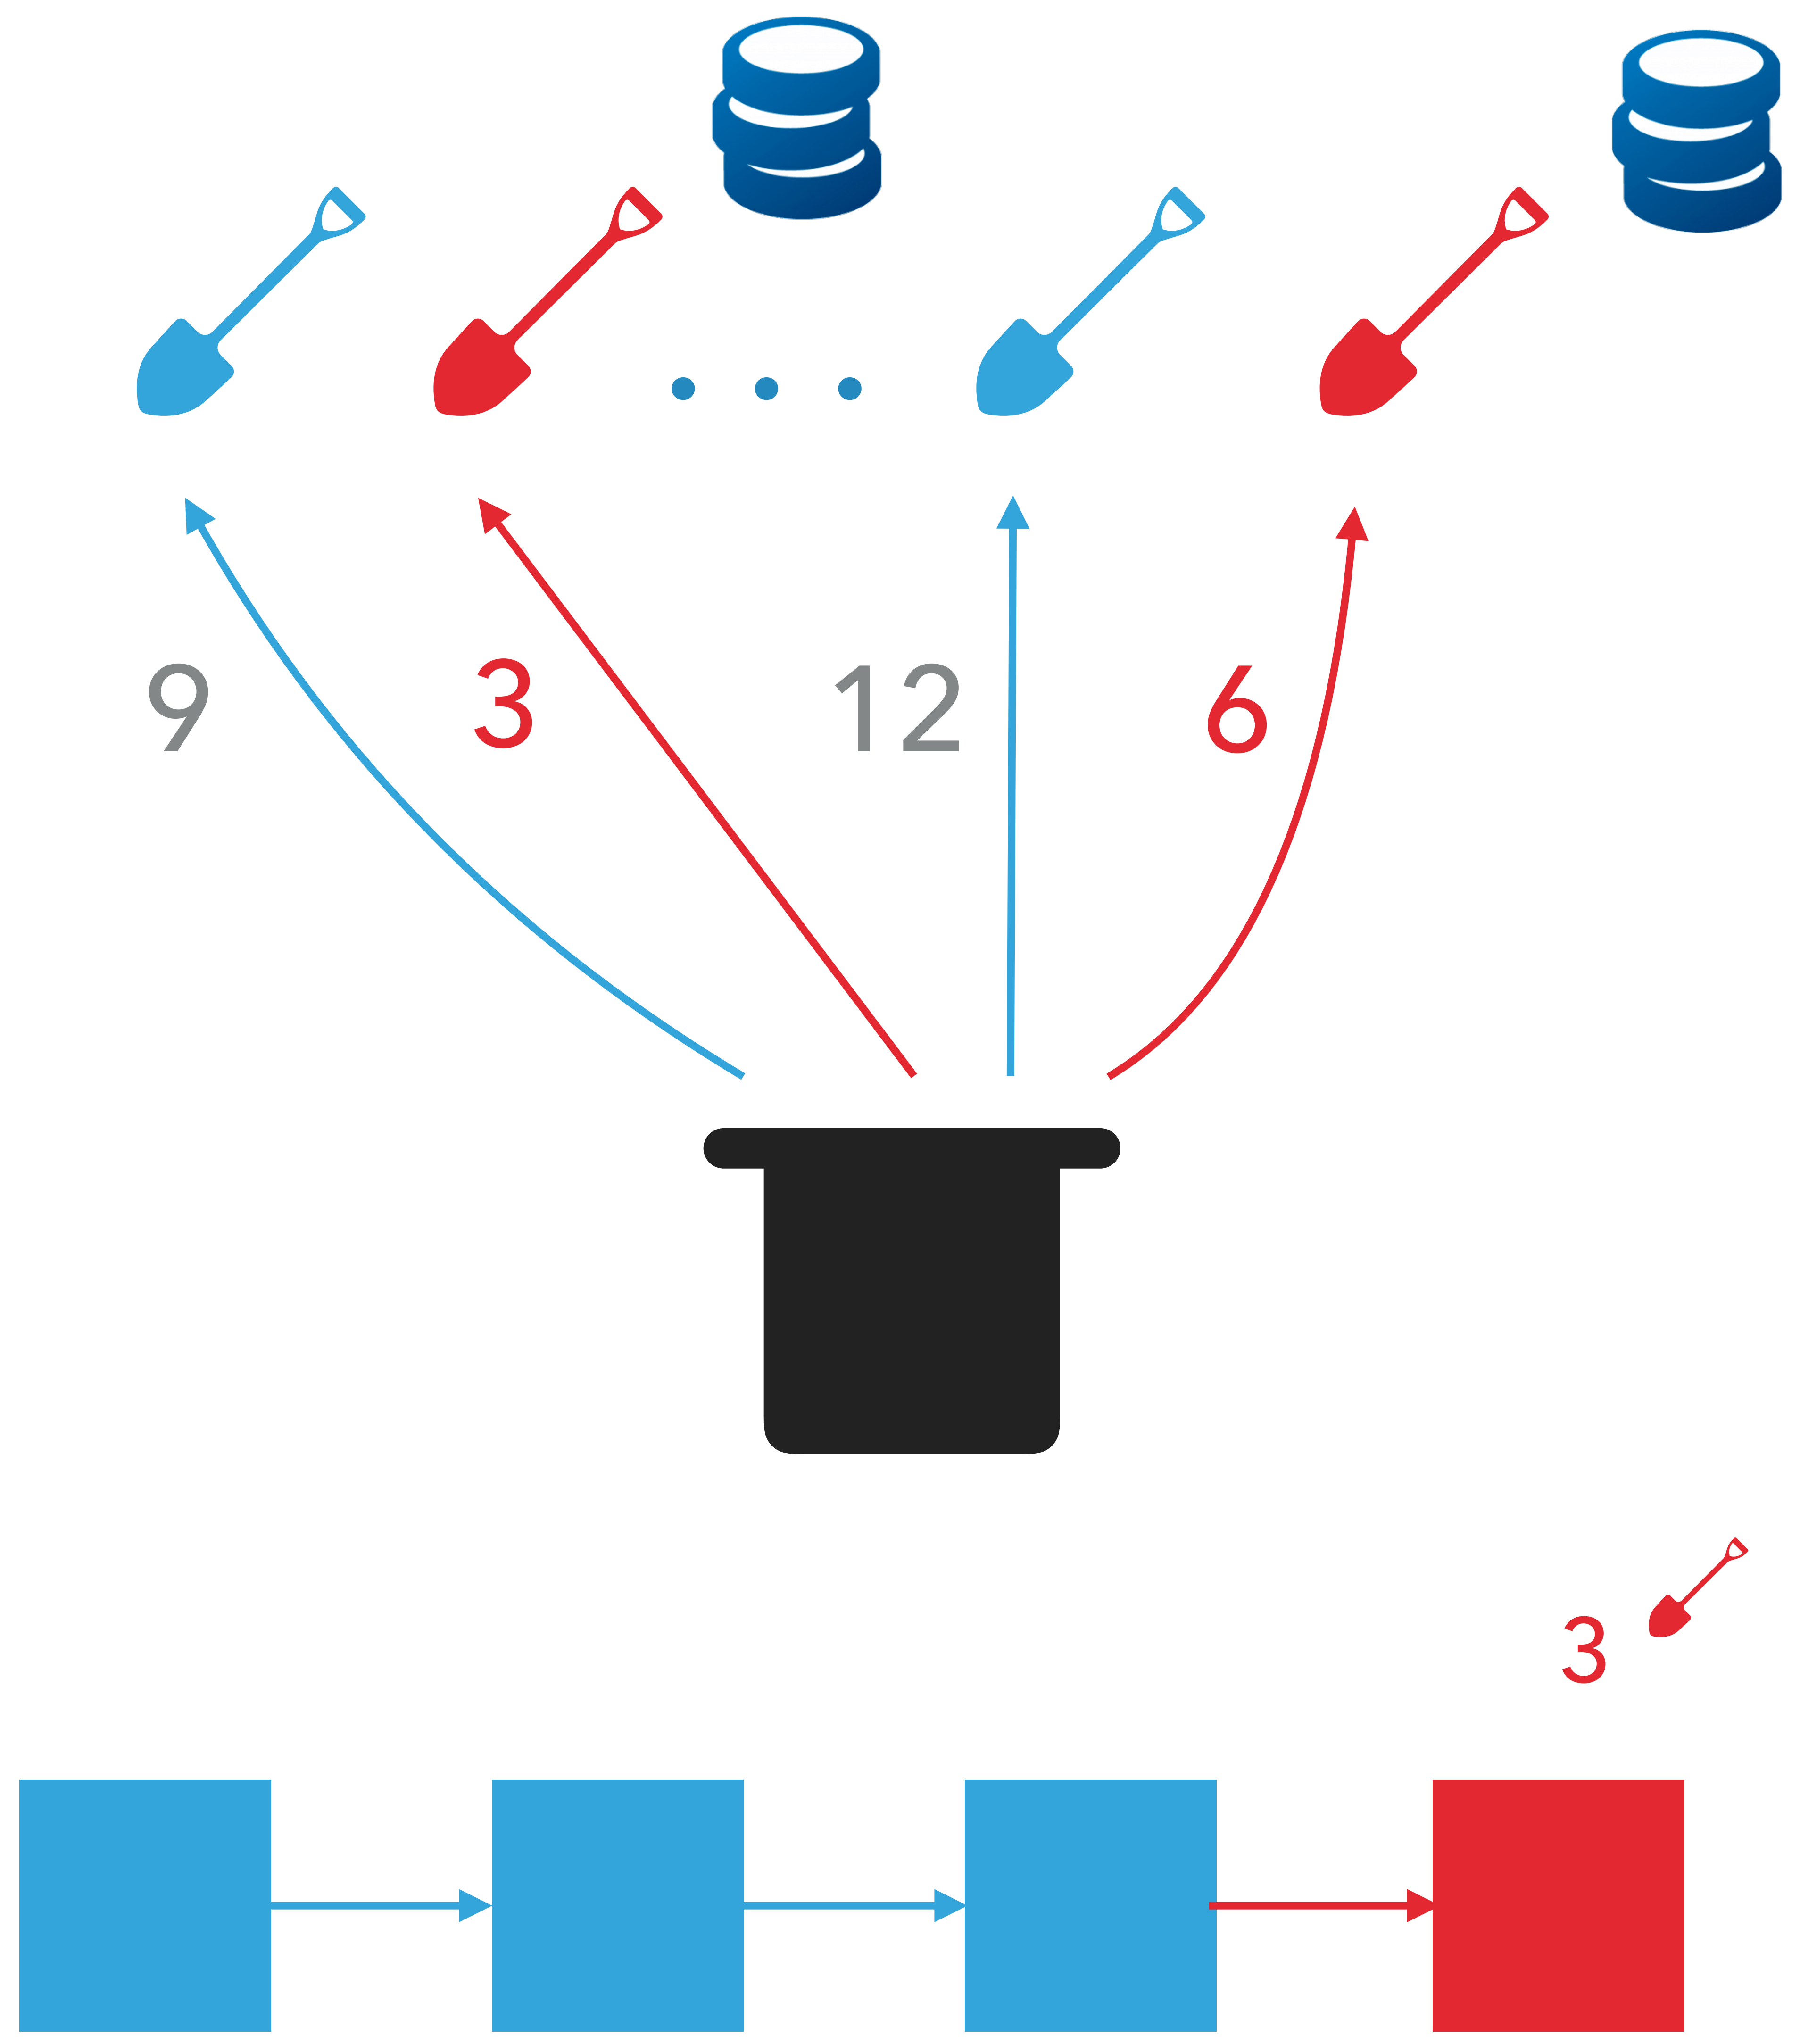

Impact on Doublespend Attack Efficacy

- Status quo (Bitcoin)
  - 20% attacker succeeds approximately 5% of the time after 6 confirmations
- Bobtail with k=20
  - 20% attacker succeeds less than 1% of the time with just 2 confirmations

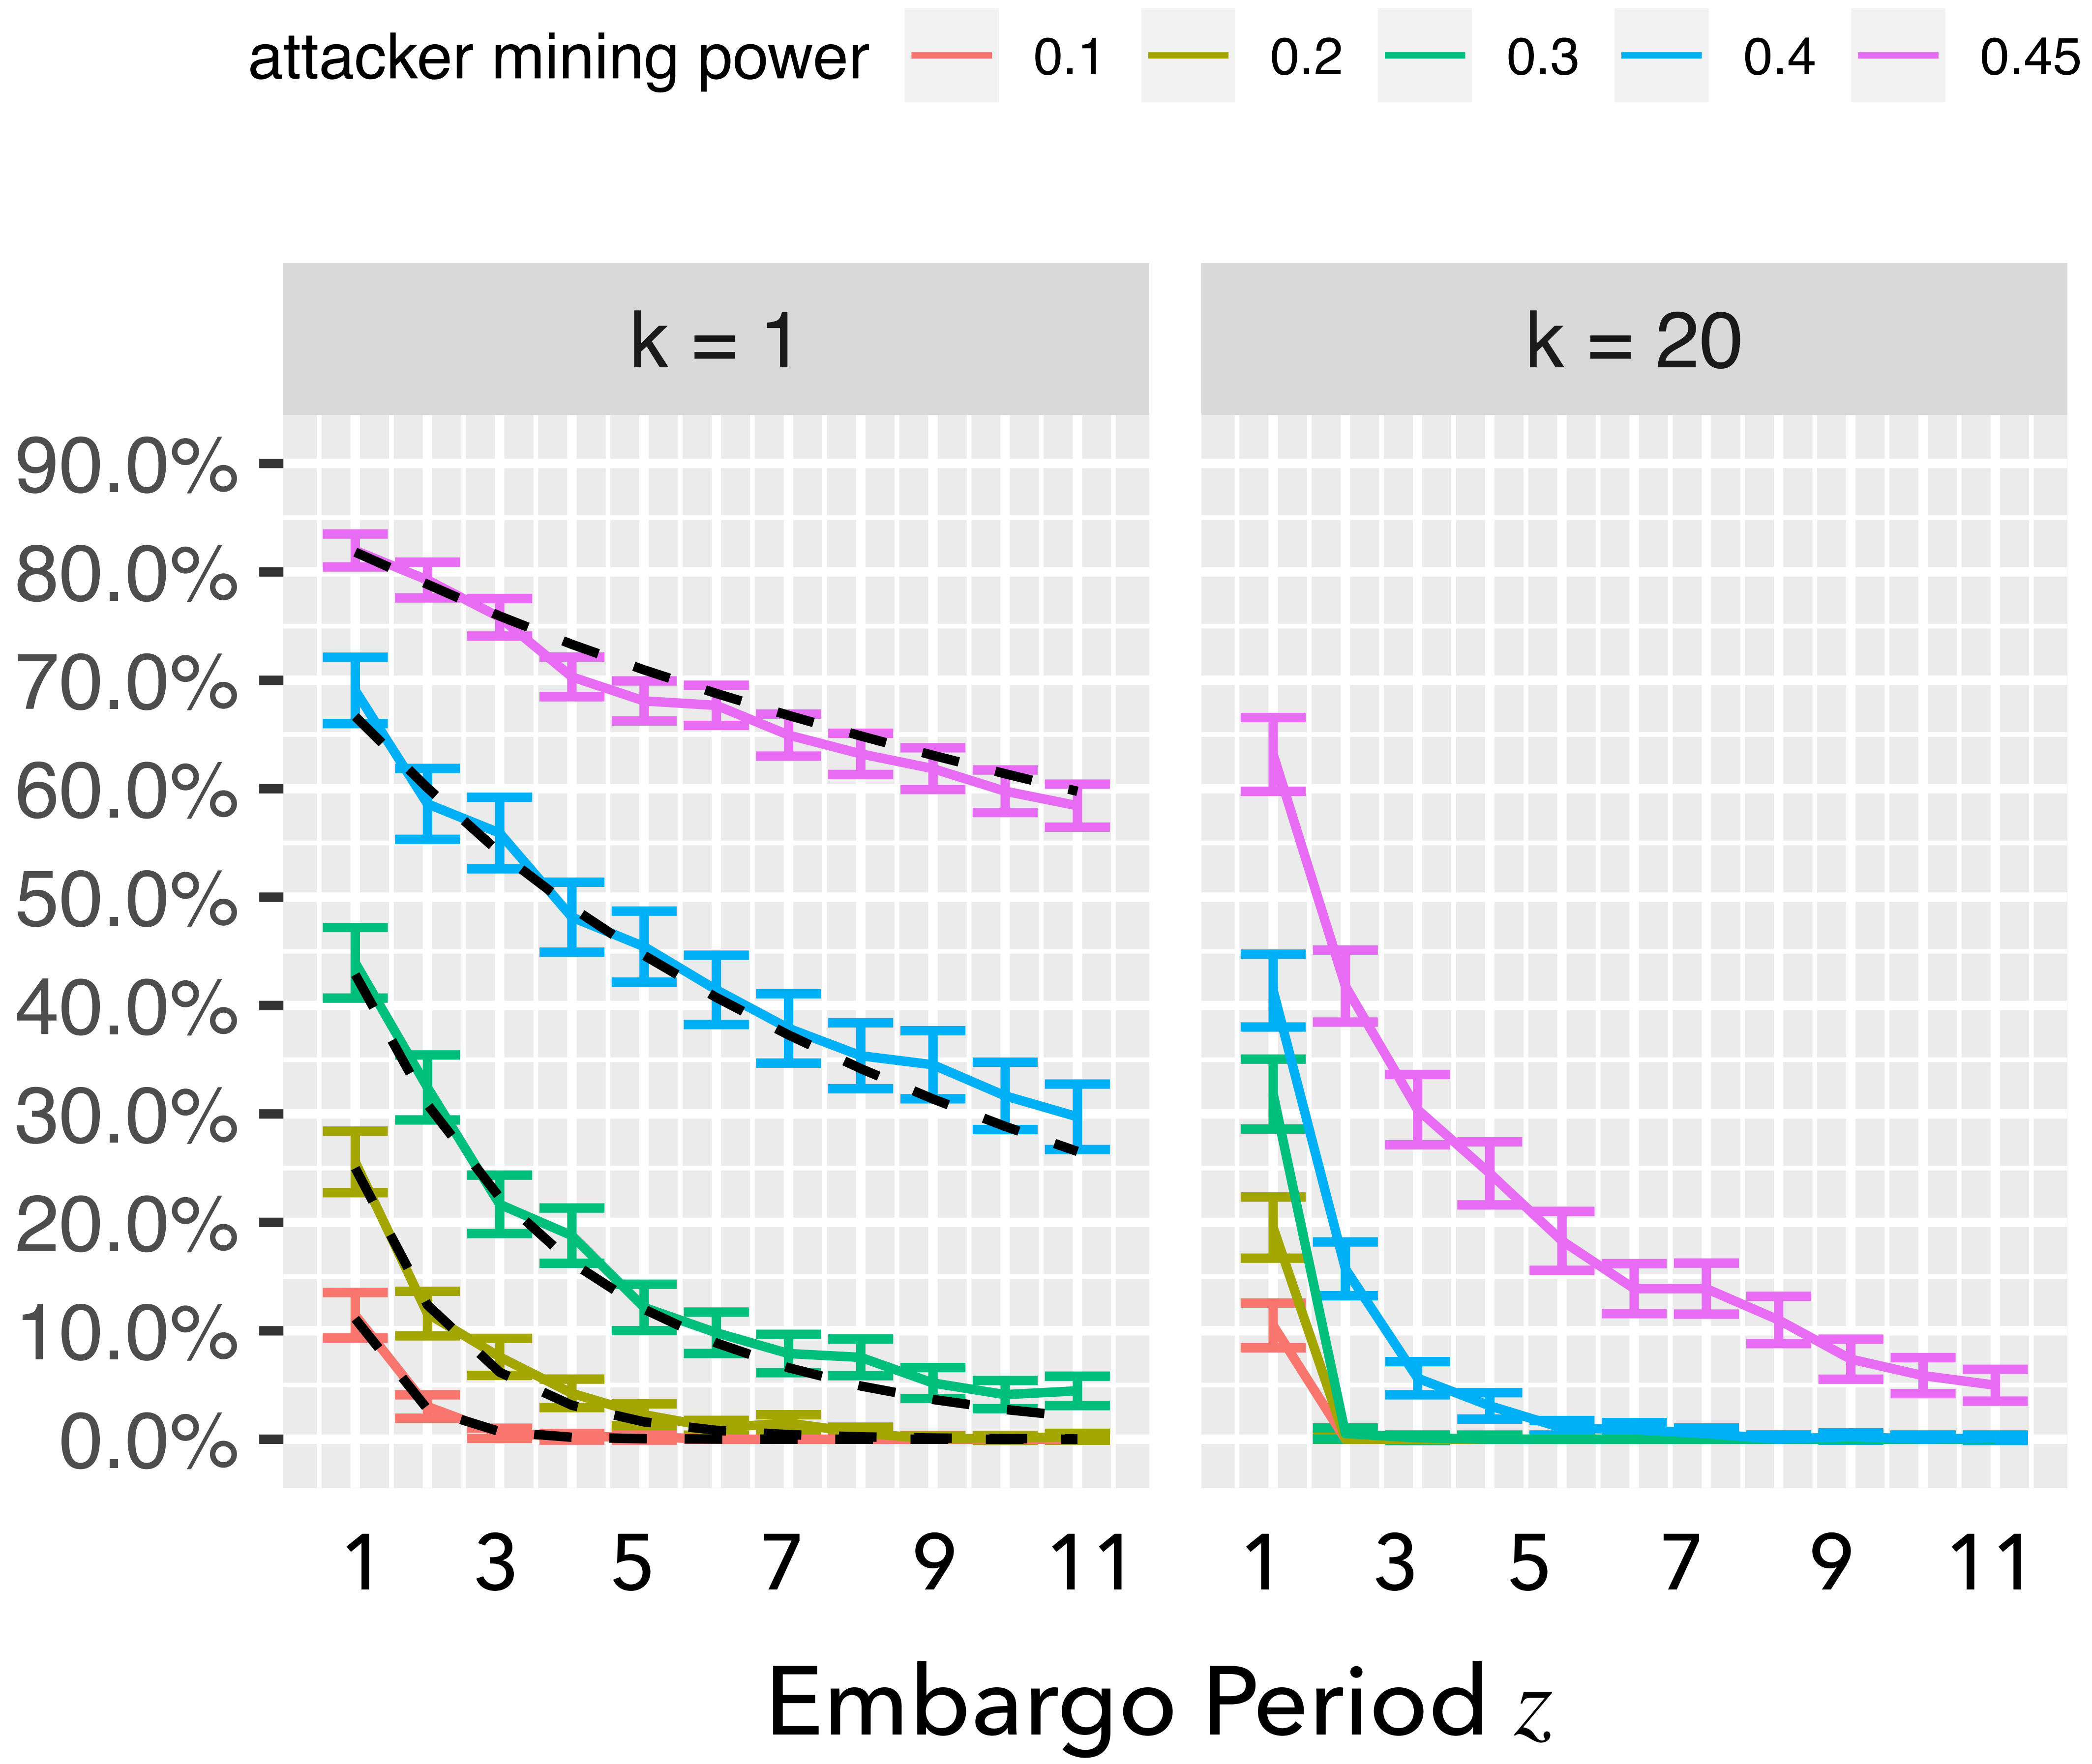

Relative Statistics

- Mining time with Bobtail for fixed target t:
  - Expected value increases by $\dfrac{k+1}{2}$
  - Variance increases by $\dfrac{(k+1)(2k+1)}{6k}$
- When expected times are aligned:
  - $t_k = \dfrac{k+1}{2}t$
  - Relative variance $O(1/k)$

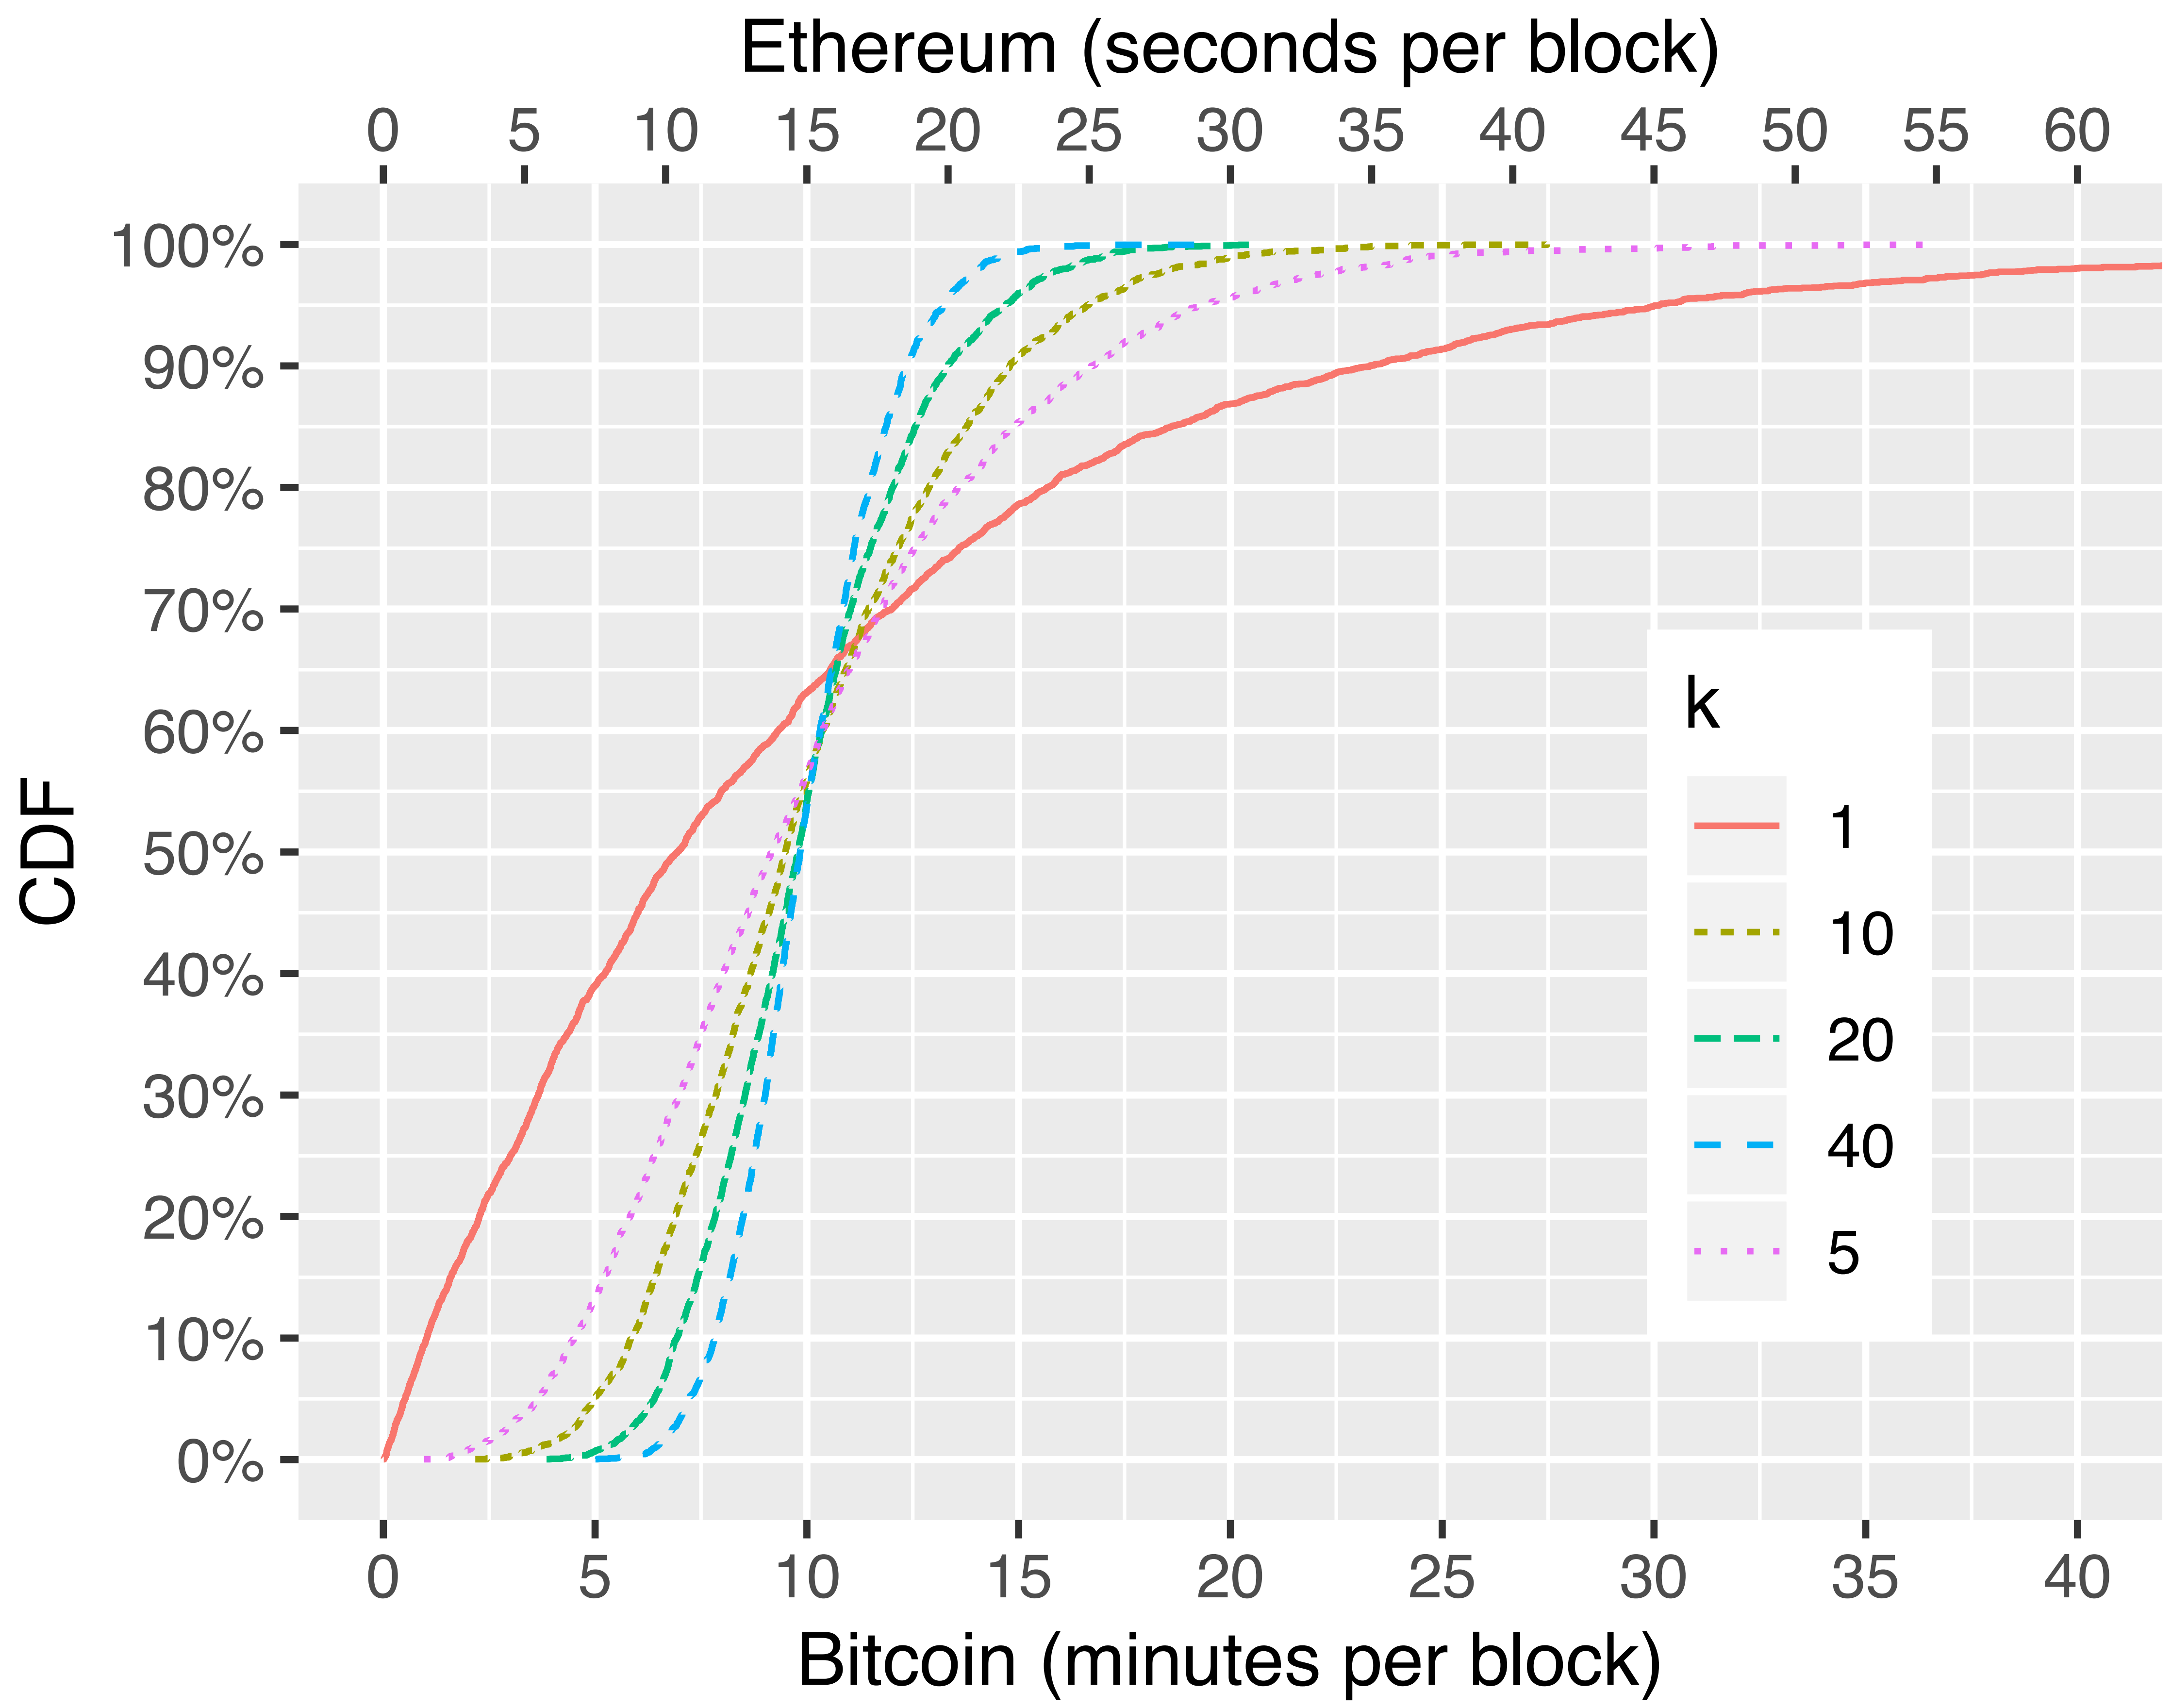

What is the Cost?

- Size of meta data increases by $k \cdot 160$B

What is the Cost?

- Size of meta data increases by $k \cdot 160\text{B}$
- **Increased network overhead**
  - Mitigated by not sending proofs in the "tail"
  - Graphene can be used to reduce redundancy

Gamma shape k

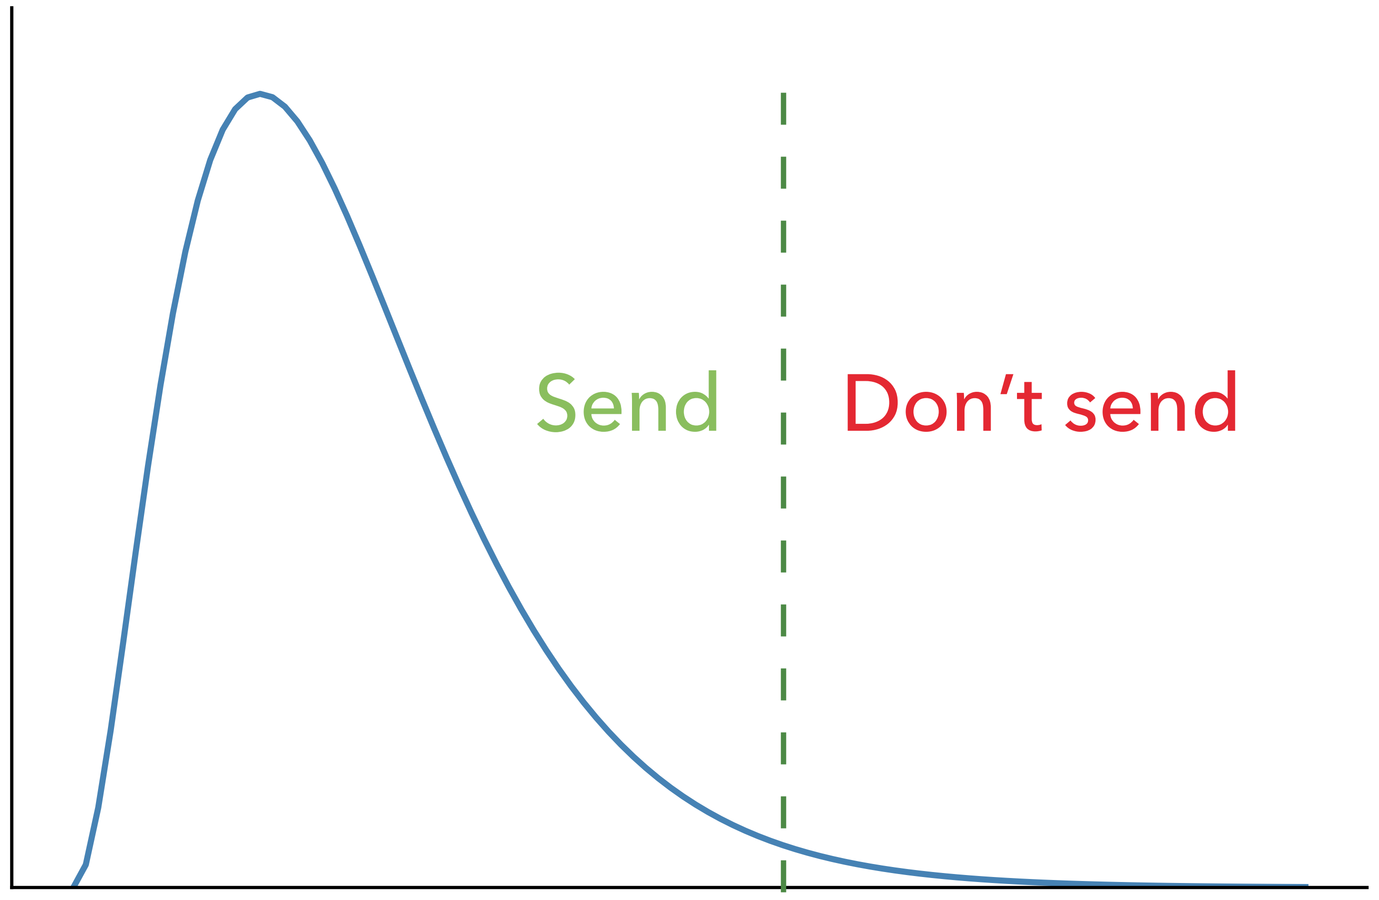

What is the Cost?

- Size of meta data increases by $k \cdot 160\text{B}$
- Increased network overhead
- **New attacks must be considered**
  - Proof withholding
  - Denial-of-Service (DoS)

Summary

- Mining process is akin to a lottery
- We can skew statistics in favor of honest majority
- This greatly mitigates fundamental attacks
  - Doublespend susceptibility reduced by orders of magnitude
- Primary cost is increased network and block overhead