Discussion of H. Halaburda, Z. He and J. Li's

An economic model of consensus on distributed ledgers

Cyril Monnet (Uni Bern and Study Center Gerzensee)

Question

- How to reach consensus in distributed ledgers…
- When some agents are honest and **utility maximizing**…
- … while others seek to jeopardise the whole system (**Byzantines**)?

Definition

- **Consensus** is achieved when **all** honest agents “commit” to a block (add the block to their local blockchain)

Simple set up

- A continuum (n-f) of honest agents
- A continuum f of Byzantine agents
- A randomly selected leader suggests a block (message)
- Agents who received the block can also sends that block to others
- **Given the number of messages received**, should an honest agent commit?
- If an honest agent commits and all others do, this agent gets R
- If an honest agent commits, and some do not, this agent gets -c
- If an honest agent does not commit, this agent gets 0

Simple(r) set up

- A ~~continuum~~ finite number (n-f) of honest agents
- A ~~continuum~~ finite number f of Byzantine agents
- A randomly selected leader suggests a block (message)
- ~~Agents who received the block can also sends that block to others~~
- **Given the number of messages received**, should an honest agent commit?
- If an honest agent commits and all others do, this agent gets R
- If an honest agent commits, and some do not, this agent gets -c
- If an honest agent does not commit, this agent gets 0

Simple(r) set up

Honest Leader

- Honest leader does not know who is honest or Byzantine
- To reach consensus, **same message** should reach **all** honest agents
- Does not matter if Byzantine agents receive it as well

Simple(r) set up

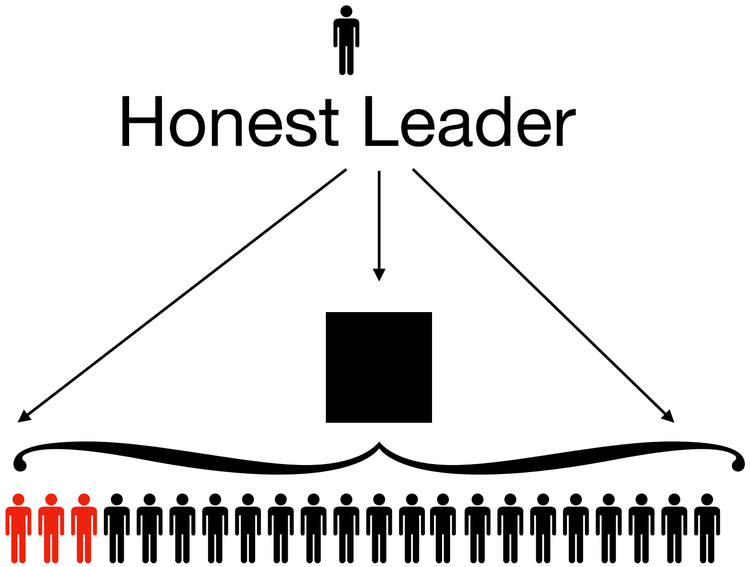

Simple(r) set up

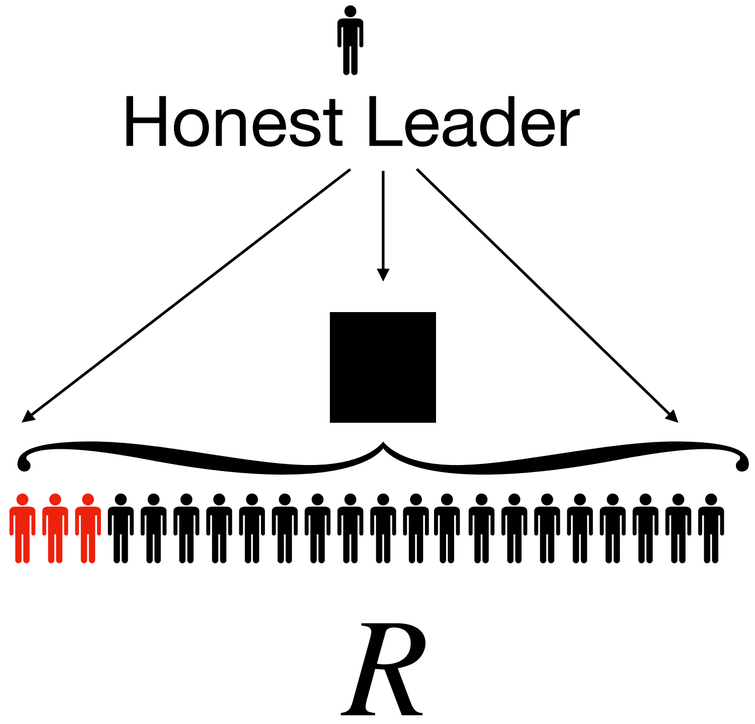

- If honest agents know the leader is honest: 2 equilibrium
  - they commit and get R
  - they don’t commit and get 0 (gridlock)

Simple(r) set up

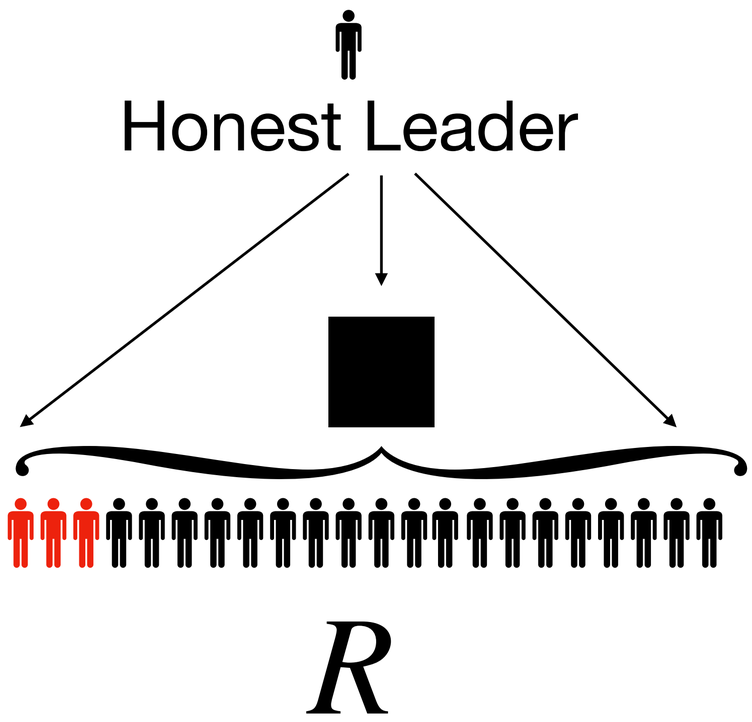

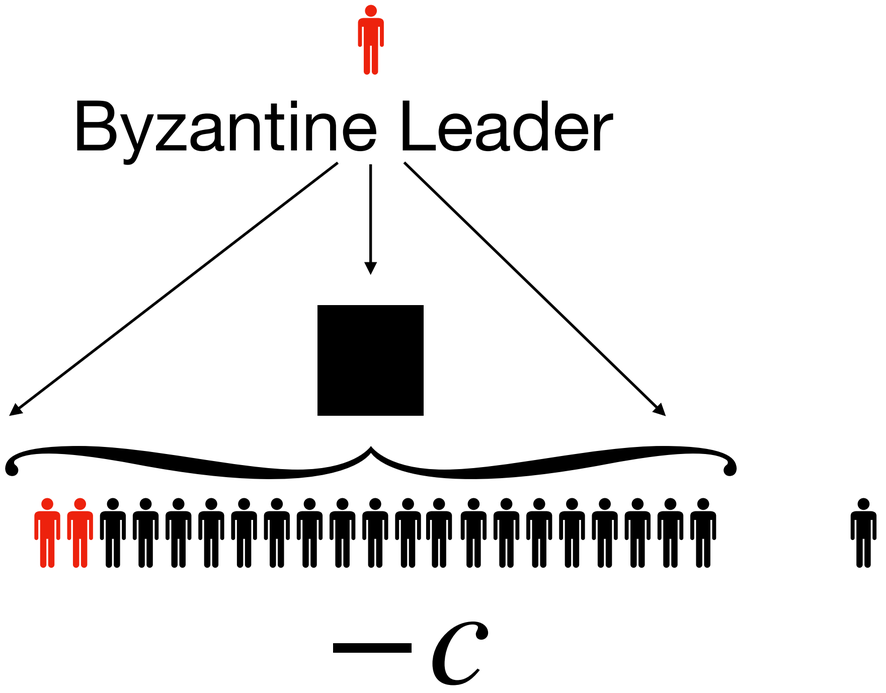

- A Byzantine leader maximizes damage by communicating the block to all but one.
- Honest agents who receive message cannot tell if the leader is Byzantine

Simple(r) set up

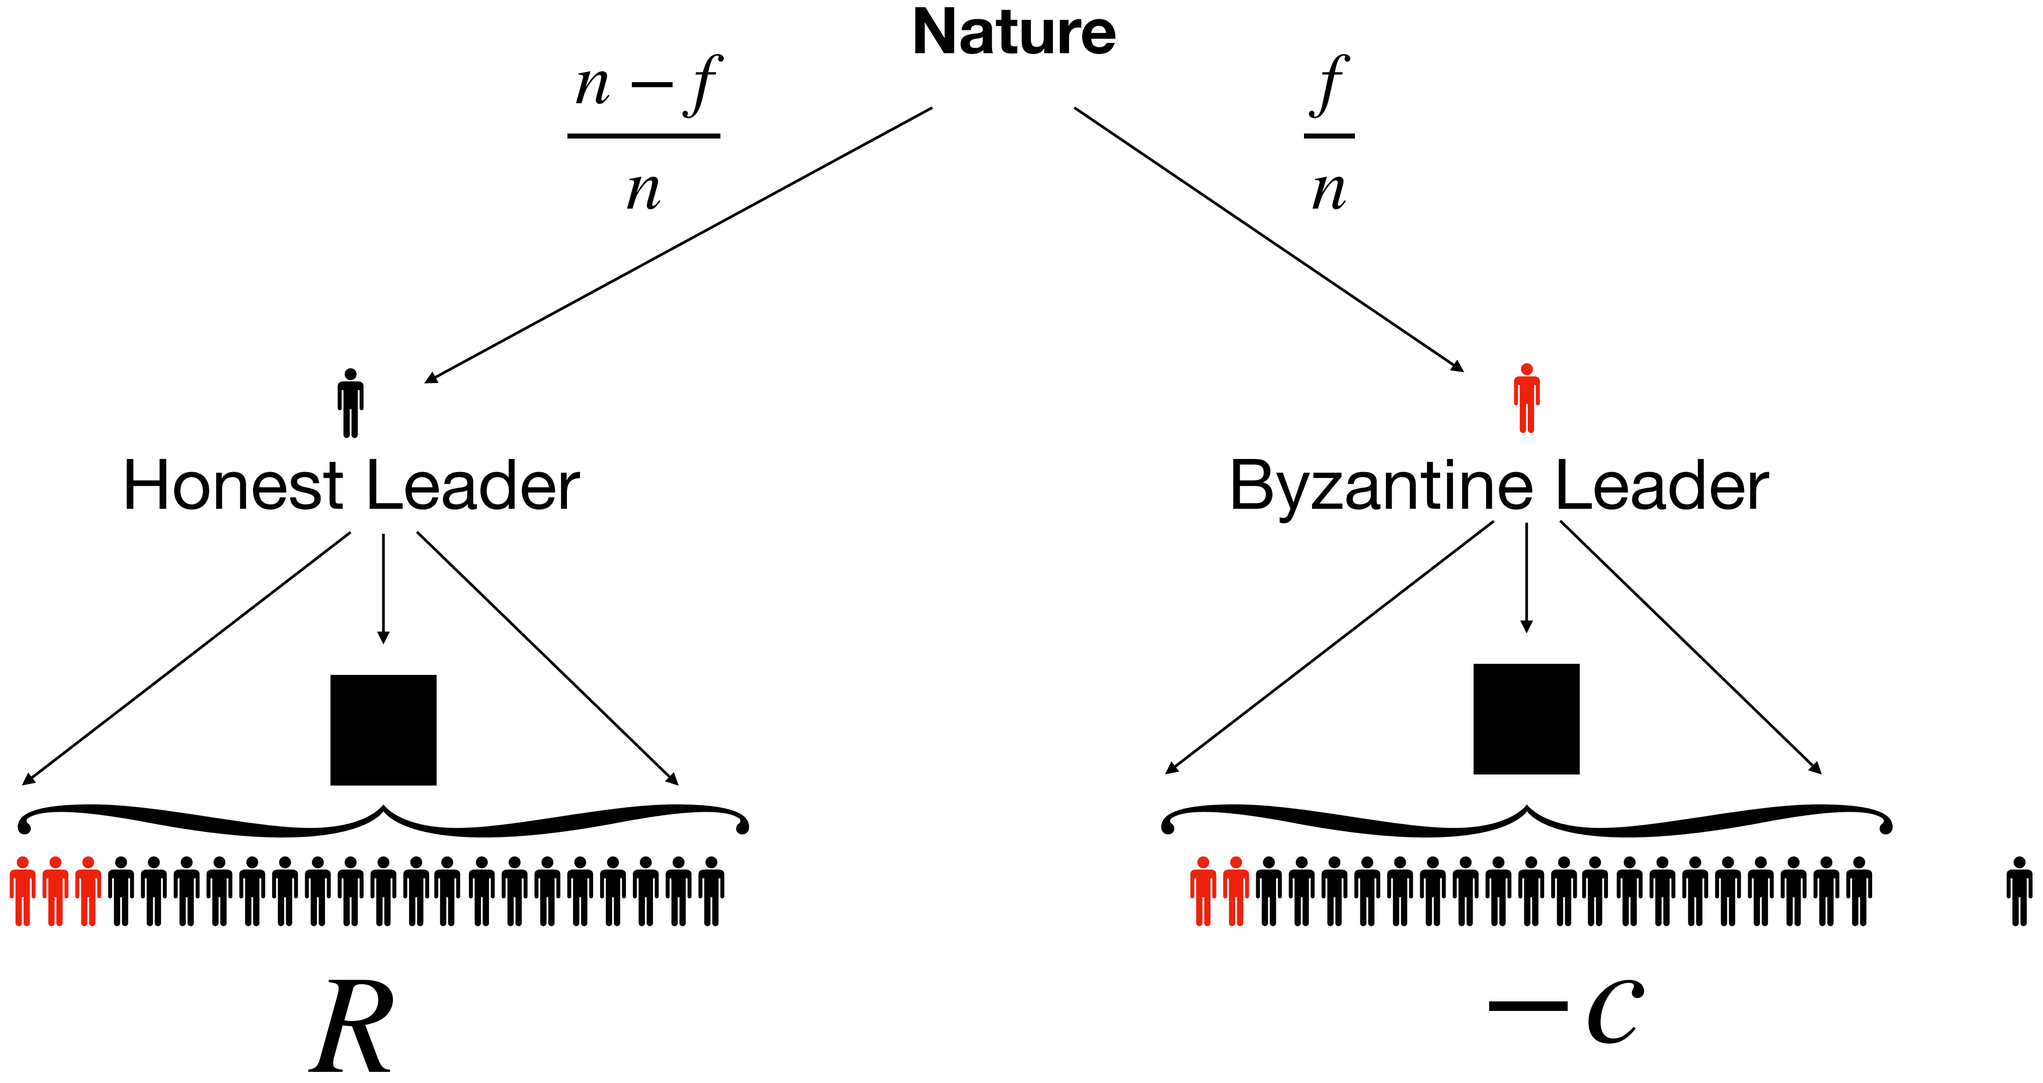

Simple(r) set up

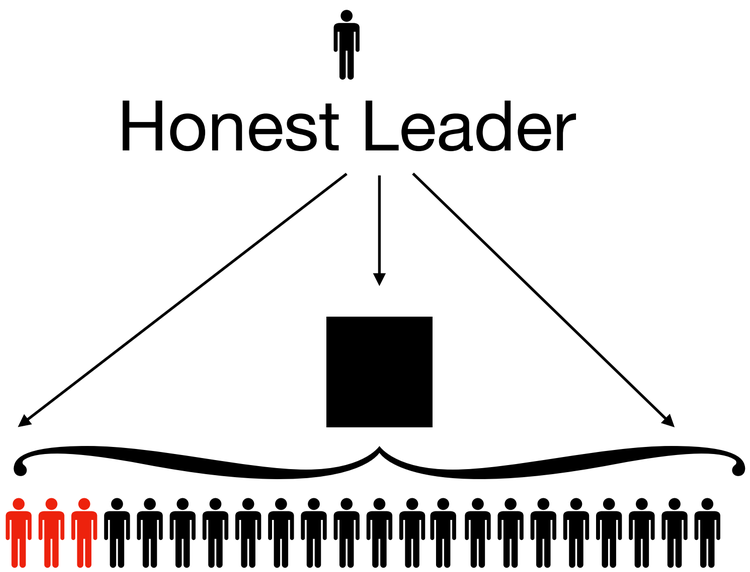

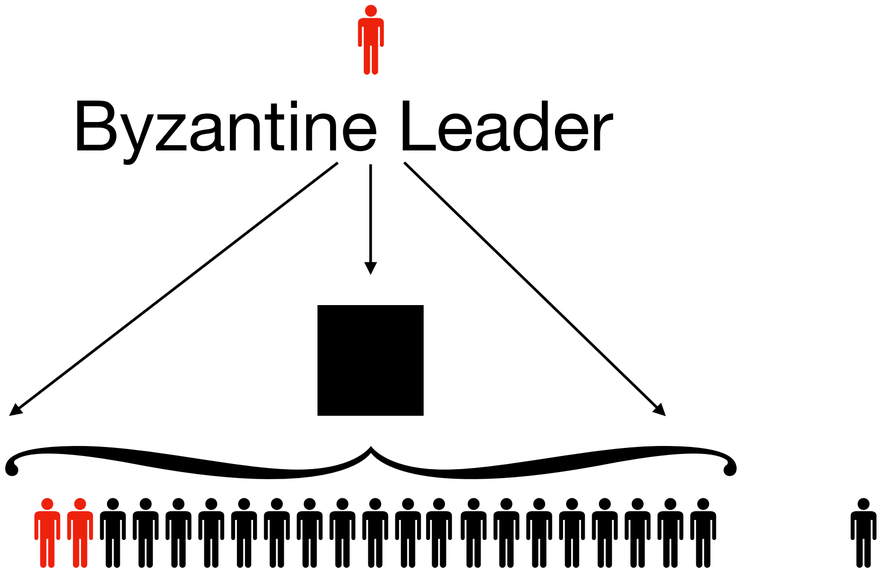

- Honest agents commit iff $\frac{n-f}{n}R - \frac{f}{n}c \geq 0$
- This is also their expected payoff

Adding a layer of communication can help

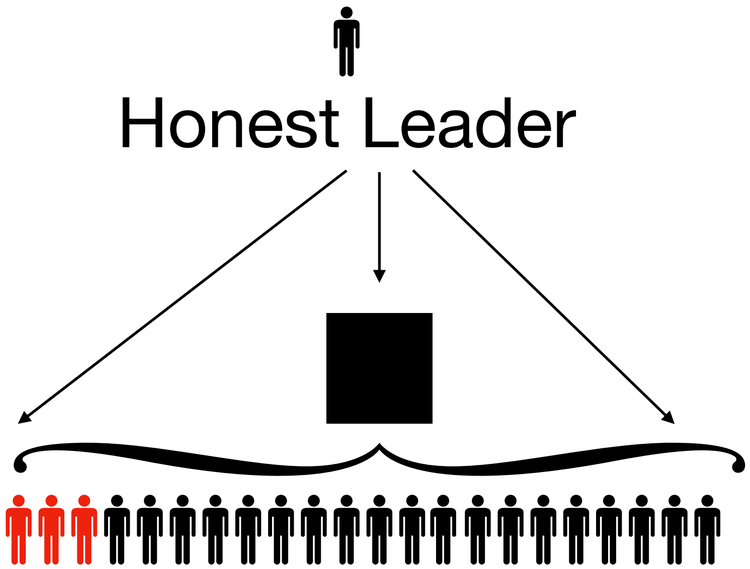

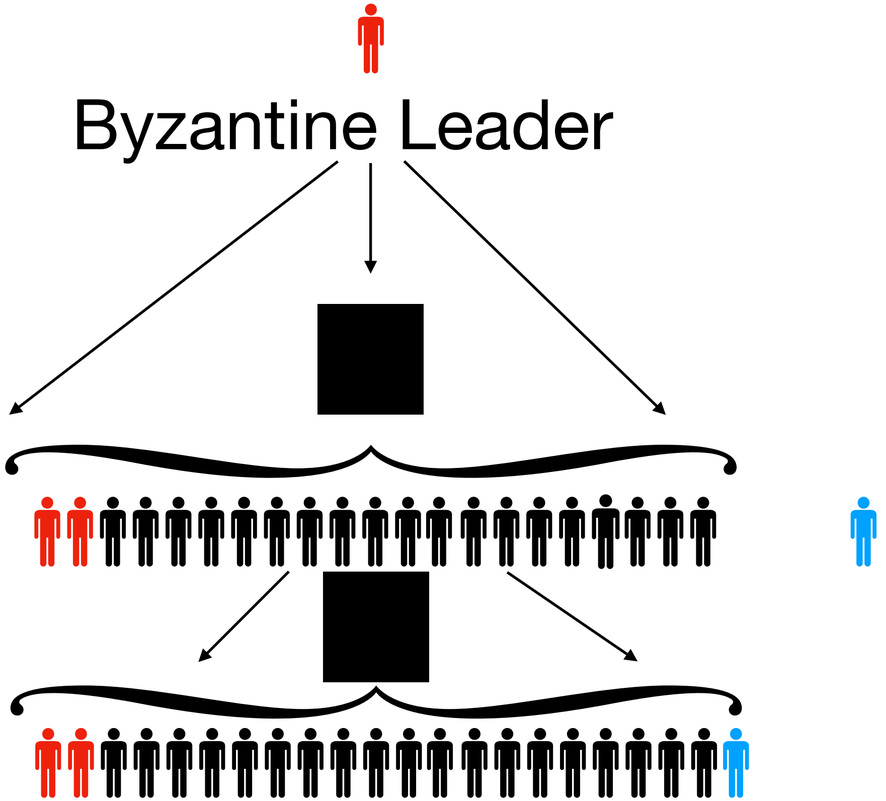

- Honest agents always commit
- The best strategy for a Byzantine leader is to submit no block
- Expected payoff : $\frac{n-f}{n}R$

Multiple messages would not affect the result

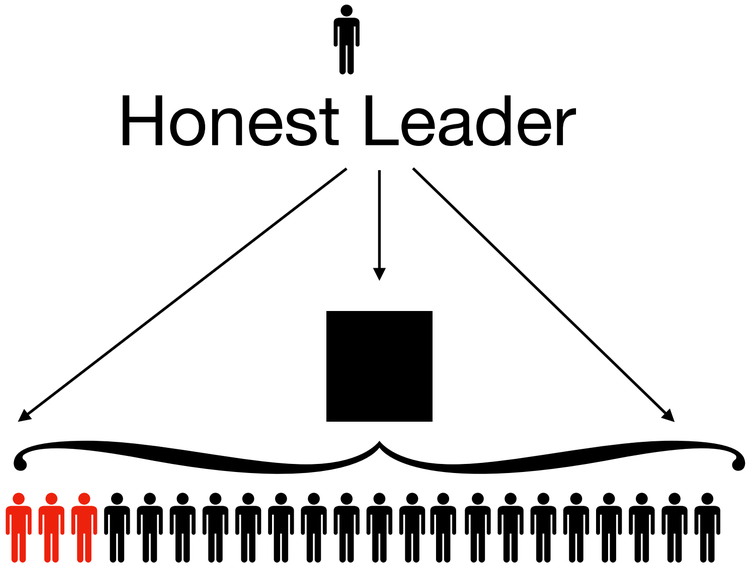

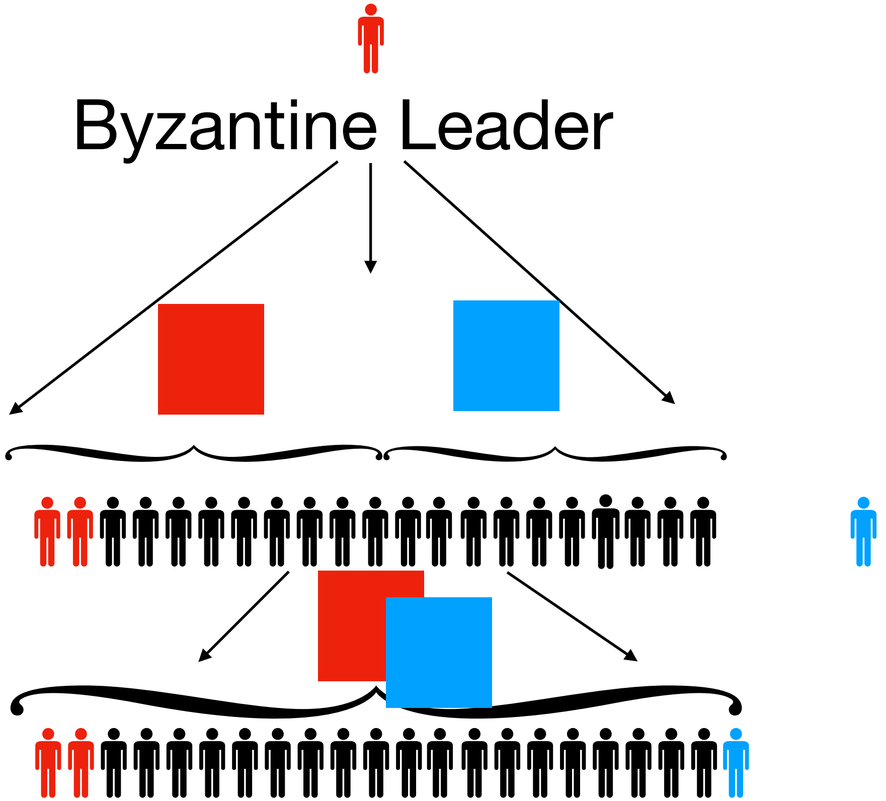

- As soon as honest agents receive two different blocks, they do not commit
- The best strategy for a Byzantine leader is to submit no block
- Expected payoff : $\frac{n-f}{n}R$

But failure to deliver message would

- Suppose a message reaches an honest agent only with probably π
- Proba it reaches all honest agents is $\pi^{(n-f)}$

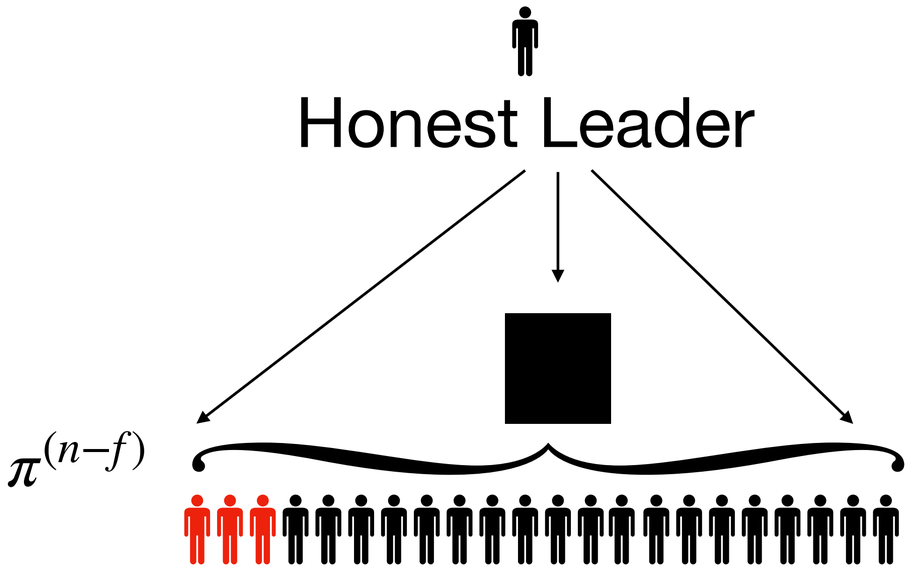

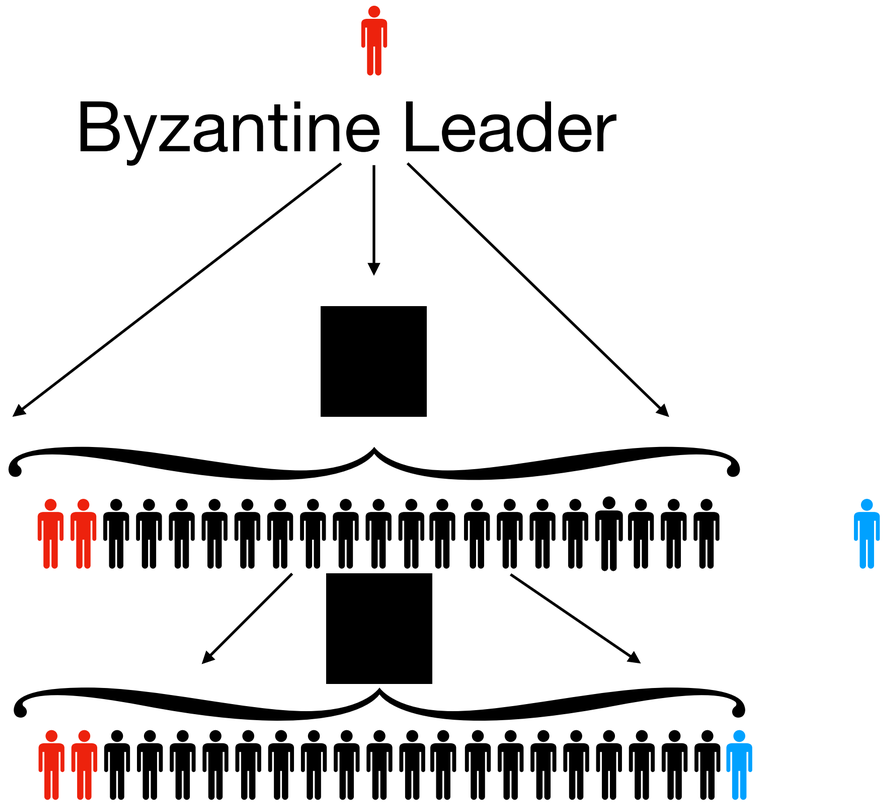

But failure to deliver message would

- The Byzantine leader now can choose h to inflict maximum damage — sends to 1 honest agent
- The probability all honest agents become informed is $\pi(\pi^{n-f-1})$

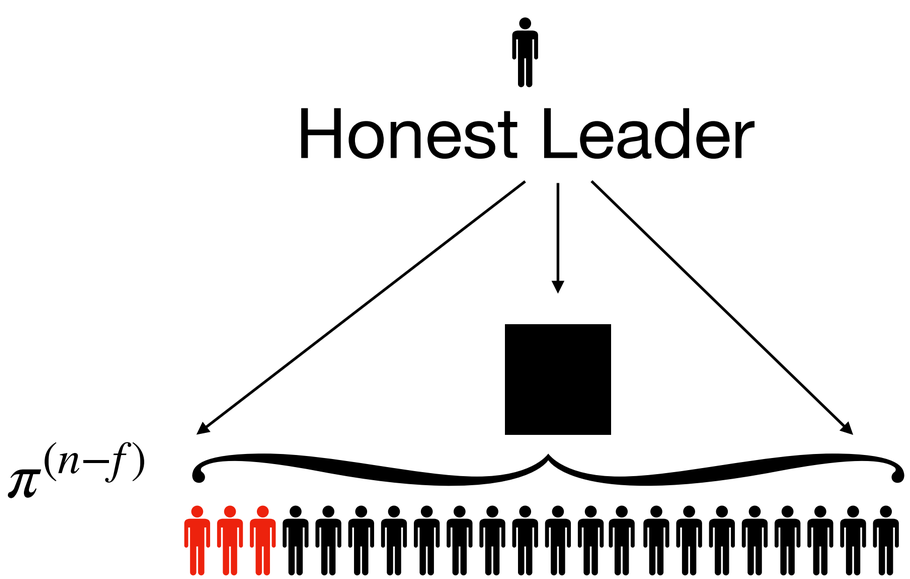

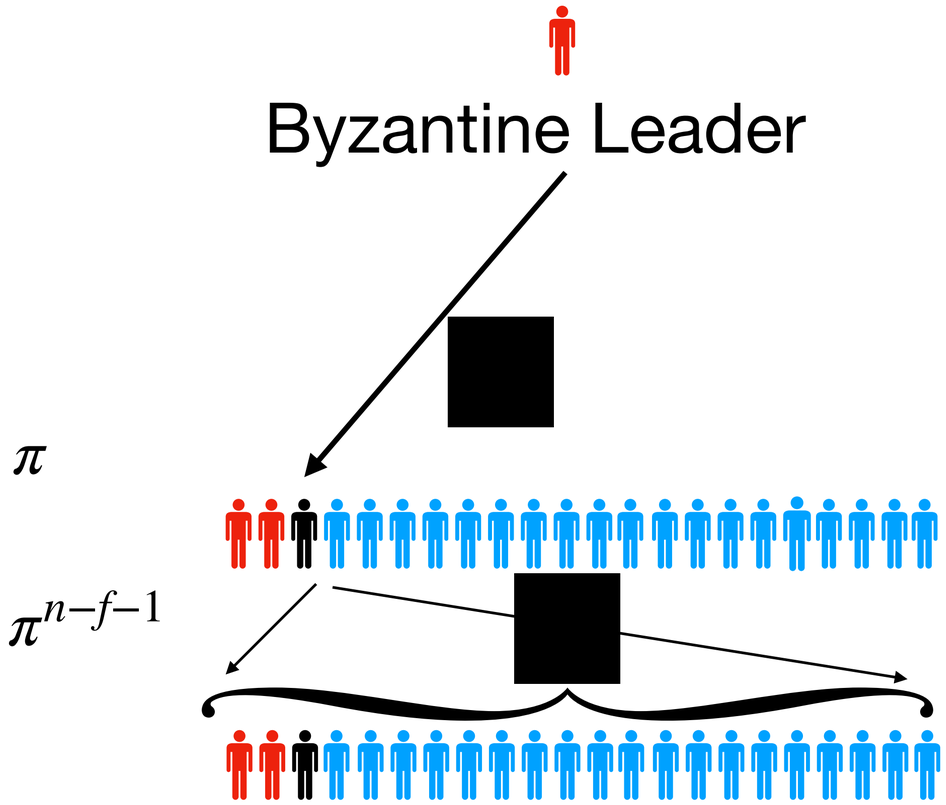

But failure to deliver message would

- The Byzantine leader now can choose h to inflict maximum damage — choose to send to 1 honest agent
- The probability all honest agents become informed about block is $\pi(\pi^{n-f-1})$
- Honest agents commit whenever $\left[\frac{n-f}{n}+\frac{f}{n}\right]\pi^{n-f}R-\left[\frac{n-f}{n}+\frac{f}{n}\right]\left[1-\pi^{n-f}\right]c\geq 0$
- Holding the fraction of honest agents constant, an increase in number of agents would make it more difficult to achieve consensus (R is a convex function of n-f!)

Key Takeaways

- With rational honest agents, there always **exists a gridlock equilibrium** : If honest believe one other will not commit it is optimal not to commit

- (Layers of) **communication** helps consensus (!)

- “More” **distribution** (a higher number of agents) implies **more communication** which makes it **more difficult** to achieve consensus

- Consensus requires **higher rewards as faults** become increasingly likely (convex!)

Final remarks

- I laud the authors for characterising all(!!) (symmetric) equilibria
  - Nice proof using iterated deletion of dominated strategies
  - But could this be simplified by determining the objective of the honest leader?

- Also, honest agents maximise their payoff under the worst case scenario —> helps reduce the set of equilibrium strategies
  - but what is the objective of the Byzantine agents (achieve maximum damage?)

Final remarks

- “Anything goes”-consensus
  - But the message better be correct: consensus on the wrong block jeopardizes the whole system

- Achieving consensus on the truth is hard —> requires verification and validation
  - Garratt and Monnet (2022)
  - Also Amoussou-Guenou et al. (2021), Auer et al. (2021)
  - This paper can help, e.g. learning if the leader is B or not through # messages

Last slide

- Would adding communication rounds help?
- Can the mechanism allow “near” consensus? (if 99% of honest agents agree, is that enough?)
- Must read paper on consensus on distributed ledgers!