Revisiting The Primitives of Transaction Fee Mechanism Design

Aadityan Ganesh

Princeton University

Clayton Thomas

Microsoft Research

Matt Weinberg

Princeton University

Blockchain Basics

- Recall: A blockchain is a sequence of publicly viewable, permanent blocks
- Each block contains up to k *transactions*: payments from one user's "wallet" to another's (for payment-system-only blockchains like Bitcoin)
- Transactions get put into blocks based on a transaction fee mechanism (TFM)

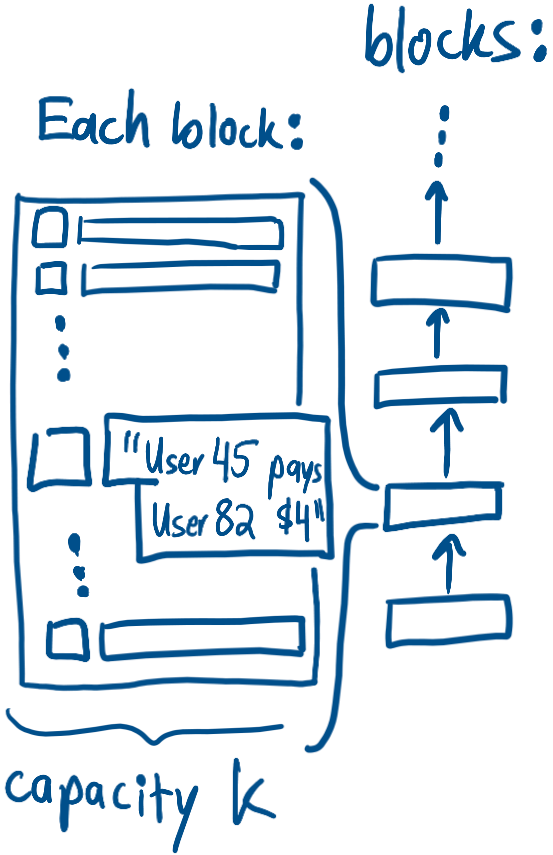

Transaction Fee Mechanisms (TFMs)

- Each block is created by a “miner” (or ‘builder’ in proof-of-stake Ethereum) using a specific, fixed algorithm **B** (“block-building process”)
- **B** implements an auction called the transaction fee mechanism (TFM)
  - Users **bid** to get their transaction included
  - Focus only on this aspect — users place bids and receive the outcome “included” or “not included”
- In contrast to classical auctions:
  - Community designs block-building process **B**
  - But, an **untrusted pseudonymous** miner looks at the bids and submits them to **B**
  - ⇒ Unique concerns (e.g., shill / censored bids)

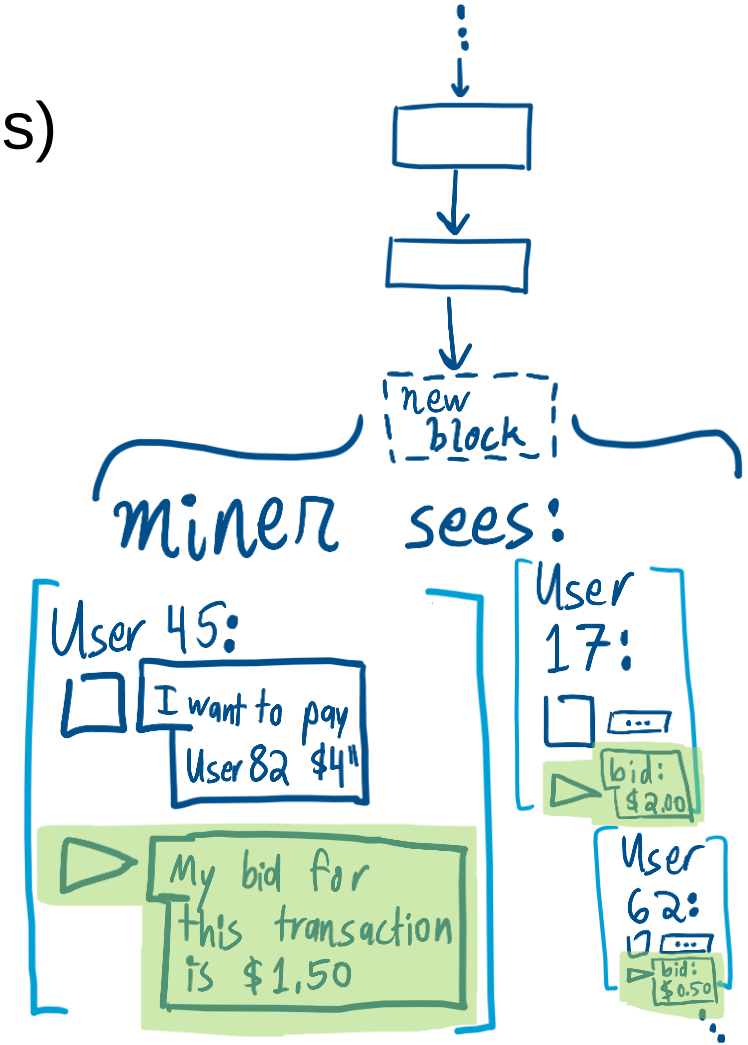

TFM design: Prior works vs. Our Paper

- Observation: **untrusted pseudonymous** miner submits bids to **B**
- Want to know: when can users “just bid their value” without worry?
- Prior work: Observe that miner may not implement the protocol as intended
  [Roughgarden '20, '21] [Chung, Shi '23] [Shi, Chung, Wu '23] [Akbarpour & Li '20]
  - **Concern:** miner lacks **commitment power** – cannot promise to “follow the protocol”
  - Highlights *EIP-1559* (or posted-price with users’ transactions burnt) as a dream TFM
- Our work:
  - **Concern:** miner **exploits their commitment power**;
    if miner “makes a convincing threat”, then can profit.
  - Highlights *Cryptographic Second-Price Auction* (C-2PA)
- This talk:
  - Mostly just the story for EIP-1559 and C-2PA
  - A bit of fancy stuff at the end

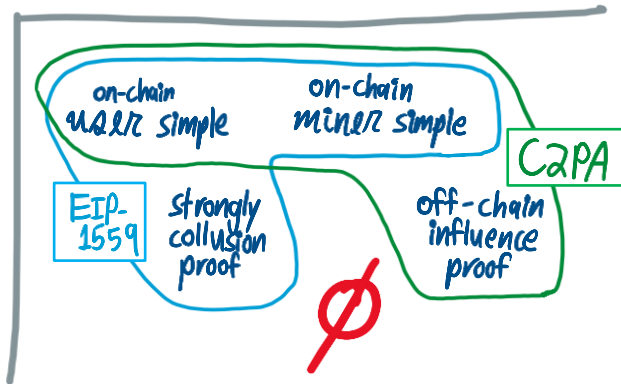

Transaction Fee Mechanisms (TFMs): Examples and “First Attempts”

Original TFM: First-price auction

Definition: first-price auction.

(Essentially the TFM of Bitcoin + pre-2021 Ethereum)

- Miner includes highest bid
- Included users pay their bid;
  this fee is transferred to the miner
  - With capacity k more generally:
    k highest bids included, each paying their bid

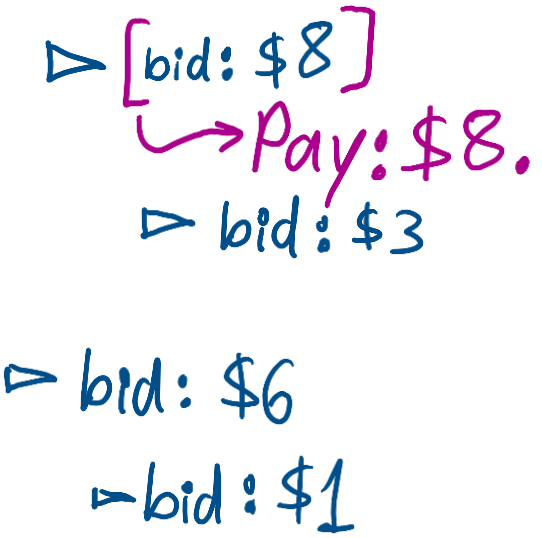

What was wrong with this?

- Users needed a lot of sophistication: had to bid their equilibrium strategies
- (Auto-bidders helped, but bidding was still challenging e.g. due to market volatility.)
- Intuitively, not “simple for users” (formally, not “User Incentive Compatible”, UIC)

Next Attempt: Second-Price Auction (2PA)

What about the solution from classical auction design?

Definition: Second-Price Auction (2PA).

- Miner includes highest bid they see
- Included user pays **the second-highest bid** to the miner
  - With capacity k more generally:
    k highest included while paying (k+1)st bid

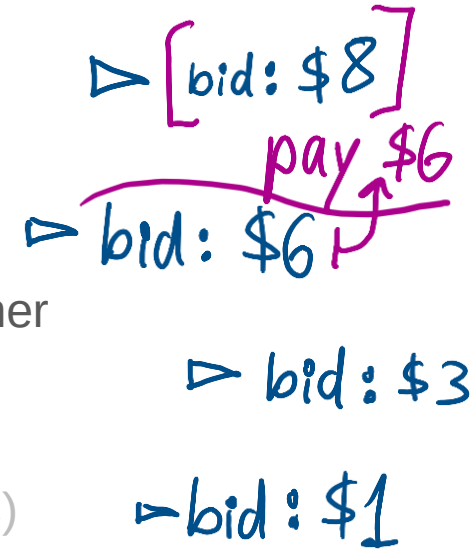

Good thing: (intended) auction is simple for users (UIC)

- Best to submit your value for being included

Bad thing: Miner doesn't want to implement the auction honestly

- E.g., Bids are $4, $12, $6. **Miner submits a bid of $11.99**
- Intuitively, not "simple for miners" (formally, not "Miner Incentive Compatible", MIC, nor credible)

Main Mechanisms #1: EIP-1559

Reformed TFM in practice: EIP-1559

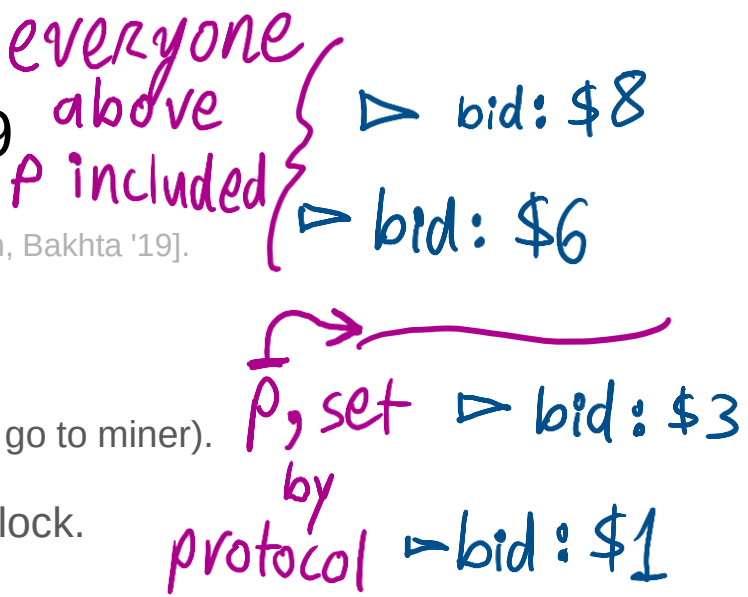

Definition: EIP-1559 [Buterin, Conner, Dudley, Slipper, Norden, Bakhta '19].

- Fixed price p (set by protocol, not by miner).
- (Users can optionally include a tip)
- Miner picks ≤ capacity k users to be included.
  - Every included user **burns p** (pay fee p, but does not go to miner).
  - (Every included transaction pays tips to miner.)
- **Note:** miners also paid a fixed block reward every block.

Lots of good things! [Roughgarden 20]

- **Especially** if (supply) > (demand); for simplicity, we focus on **infinite supply** (k = ∞)
- Straight-forward bidding for users (UIC)
  - Bid if you want it at price p; otherwise don't !
- Miner can't profit from dropping / injecting bids (MMIC)

(We'll cover collusion later – be patient!)

Wrong but not terrible: if finite supply, & > k users are each willing to pay > p…

- Devolves back into first-price auction. MMIC and OCA-proof but not UIC.

EIP-1559 In Our Paper

Definition: EIP-1559 for infinite supply

- Fixed price p (set by protocol, not by miner).
- Miner picks any set of at most k transactions to include.
  - Every included user **burns p**
  - Miner gets **constant reward**

Our paper:

- Miner tweets “users who don’t pay me $5 off-chain do not get included”
- (Or, in full EIP-1559 with tips, “users must **tip** me $5” directly on-chain!)
- Observation: if **any** user is willing to pay $5, and actually “gives in”, then miner will strictly profit from this manipulation

EIP-1559 ⟹ Fist big idea

Definition: EIP-1559 for infinite supply: Posted price with burning
Miner tweets “users who don’t pay me $5 off-chain do not get included”

Our paper:

- Definition: **Off-Chain Influence Proof**: Miner cannot profit by running **any** separate off-chain mechanism to decide how to play in the on-chain TFM (provided the users play in a Bayes-Nash Equilibrium)
- ⇒ EIP-1559 is **not** Off-Chain Influence Proof

Main Mechanisms #2:
Cryptographic Second-Price Auction

Alternative: Cryptographic Second Price Auction (C-2PA)

- Recall second price auction: Highest bid included, pays 2nd highest bid to miner
- Now with *encrypted bids*: Miner can't see the values of bids until auction finished
  - (Technical tools from cryptography: FHE, VDF, MPC, etc)
- Our model: Miner can condition their strategy on *who* submits a bid,
  but has no information about *the value* of the bids
- Avoids credibility issues where miner injects bid just less than the largest bid

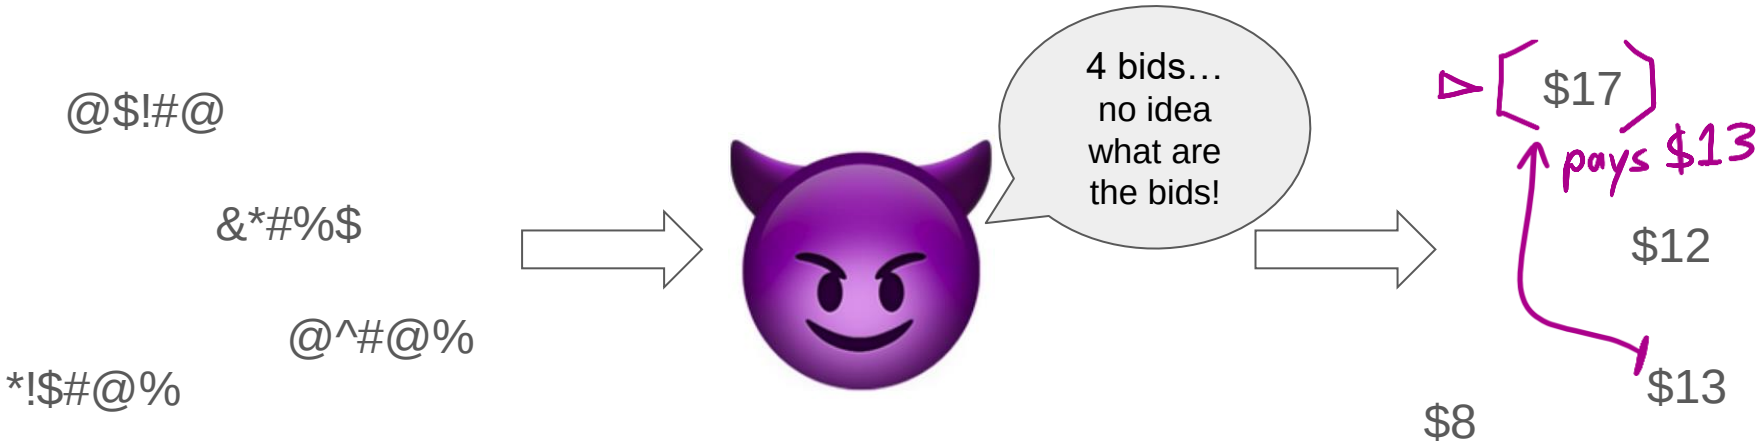

Alternative: Cryptographic Second Price Auction (C-2PA)

Second Price Auction with Encrypted Bids

Possibly bad thing? [Shi, Chung, Wu '23]

- Miner injects fake bid to emulate **reserve price**
  - Classical auction theory ⟹ for all distributions of user values $\mathcal{D}$, optimal auction is “second price auction with reserve r”
  - “Never give the item away for less than r”
  - Observation: Miner can implement this by injecting a fake bid ⟹ As written, C-2PA is not “simple for miners” (MIC)
- However… Is anything really wrong with this?

@$!#@

&*#%$

@^#@%

*!$#@%

miner bids optimal reserve r

@$!#@

&*#%$

@^#@%

C-2PA ⟹ Second Big Idea

Second Price Auction with Encrypted Bids

Miner injects fake bid to emulate **reserve price**

Our paper:

- Allow the miner to set the reserve!
  - For whatever prior $\mathcal{D}$ of user values the miner holds,
    miner's optimal strategy is to simply set the optimal reserve
- Protocols can accept general "advice" from the miner
  - As long as it doesn't harm other good things, like users' incentives
  - ⟹ C-2PA is "simple for users and miners", provided the miner can set reserve
  - Additionally, (since this is the Myersonian optimal auction,)
    C-2PA **is** off-chain influence proof

Formalizations

Sketch of Formal Model and Definitions

- Block-building process **B**. Takes user bids and miner "advice" (e.g., reserve)
- **B** induces "on-chain game" **C**. Users places bids; miner can censor or inject bids
  - Two levels of cryptography: "plaintext" (miner sees the bids); "miner-gatekeeper" (miner sees only **who** submits a bid, not the value of the bid)
- Definition: **On-chain Simplicity** of an equilibrium (s_{miner}, s_{users})
  - On-chain **user** simplicity: Users follow the protocol (i.e., bid their values) and it's DSIC (≈UIC)
  - On-chain **miner** simplicity: Miner follows the protocol (i.e., uses a constant advice without censoring or injecting any bids) and s_{miner} gets max possible revenue given s_{users} (≈ MMIC)
- **C** induces "off-chain game" **D**. Miner commits to an off-chain mechanism $\mathcal{M}_{off}$; users report to $\mathcal{M}_{off}$ to determine play in **C**.
- Definition: **Off-Chain Influence Proofness** of on-chain equilibrium $\sigma = (s_m, s_u)$
  - In **any** $\mathcal{M}_{off}$ and in **any** "user-Bayes-Nash-Equilibrium", miner revenue is not higher than in σ
  - $\Longrightarrow$ Miner cannot profit above σ in **any** user-equilibrium consistent with **B**

Formal Theorems

- **Theorem:** For all distributions of user values $\mathcal{D}$, EIP-1559 (under the truth-telling equilibrium) is on-chain (miner and user) simple, but **not** off-chain influence proof.
- **Theorem:** For all distributions of user values $\mathcal{D}$, C2PA (under the truth-telling equilibrium, and when the miner sets the Myersonian-optimal reserve price for $\mathcal{D}$) is on-chain (miner and user) simple, **and also off-chain influence proof**.

Collusion Resistance: Tradeoffs + Remarks

Our desiderata and results

Our desiderata and results

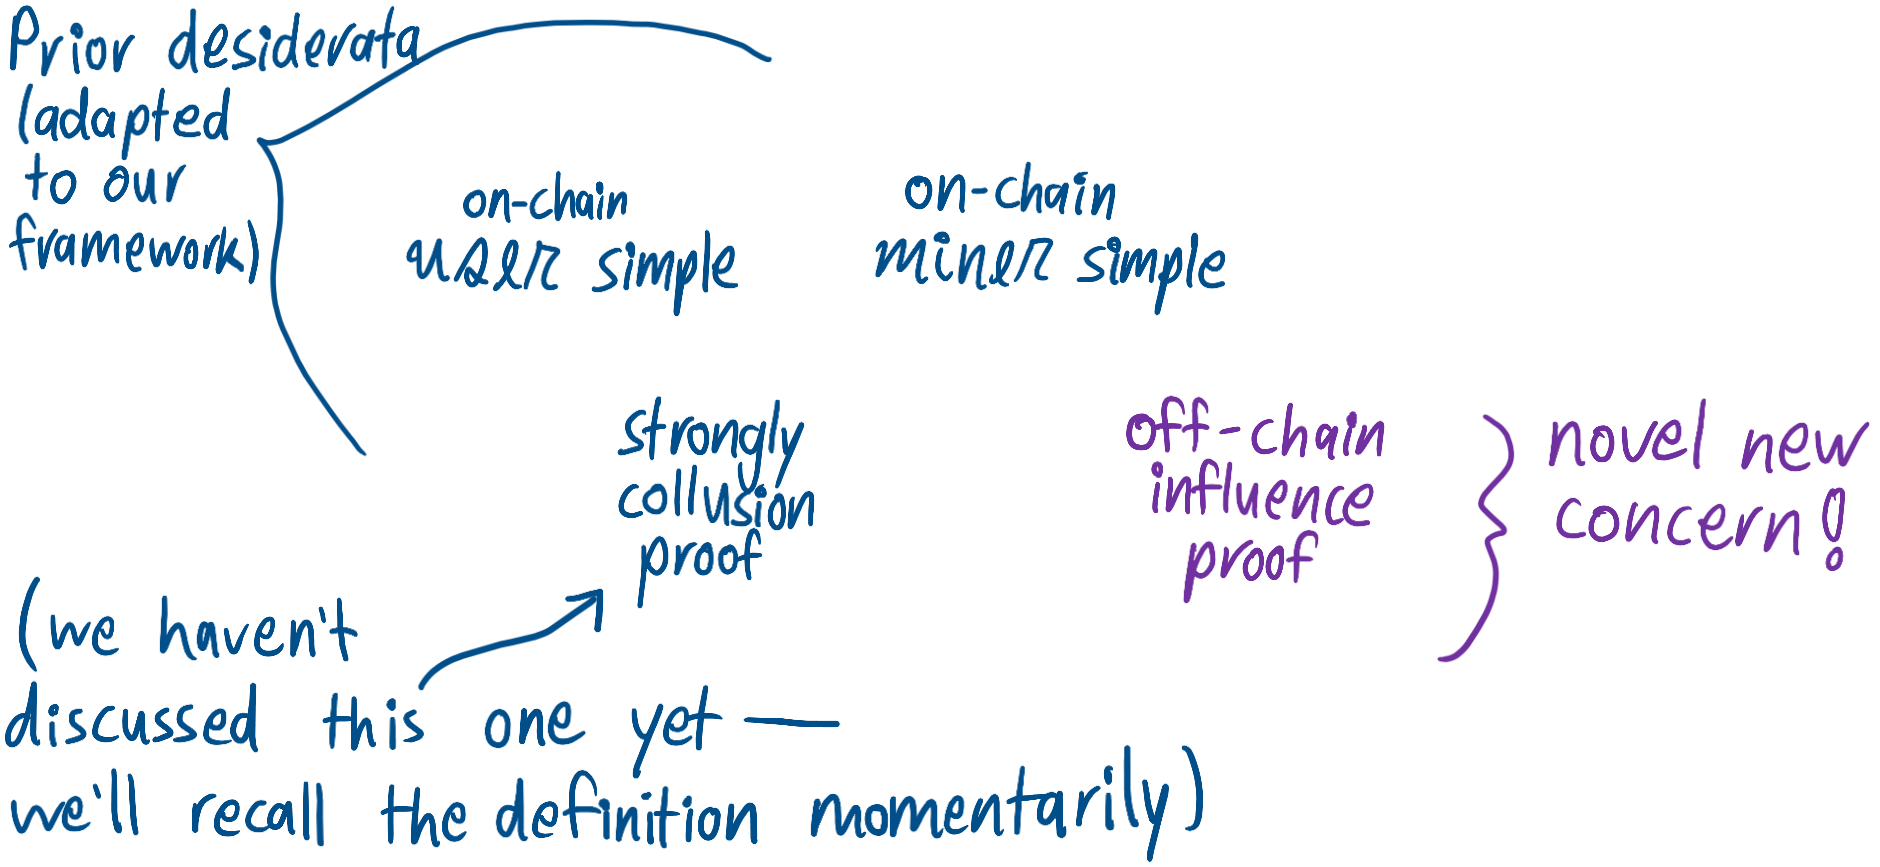

Our desiderata and results: Two mechanisms satisfy different properties

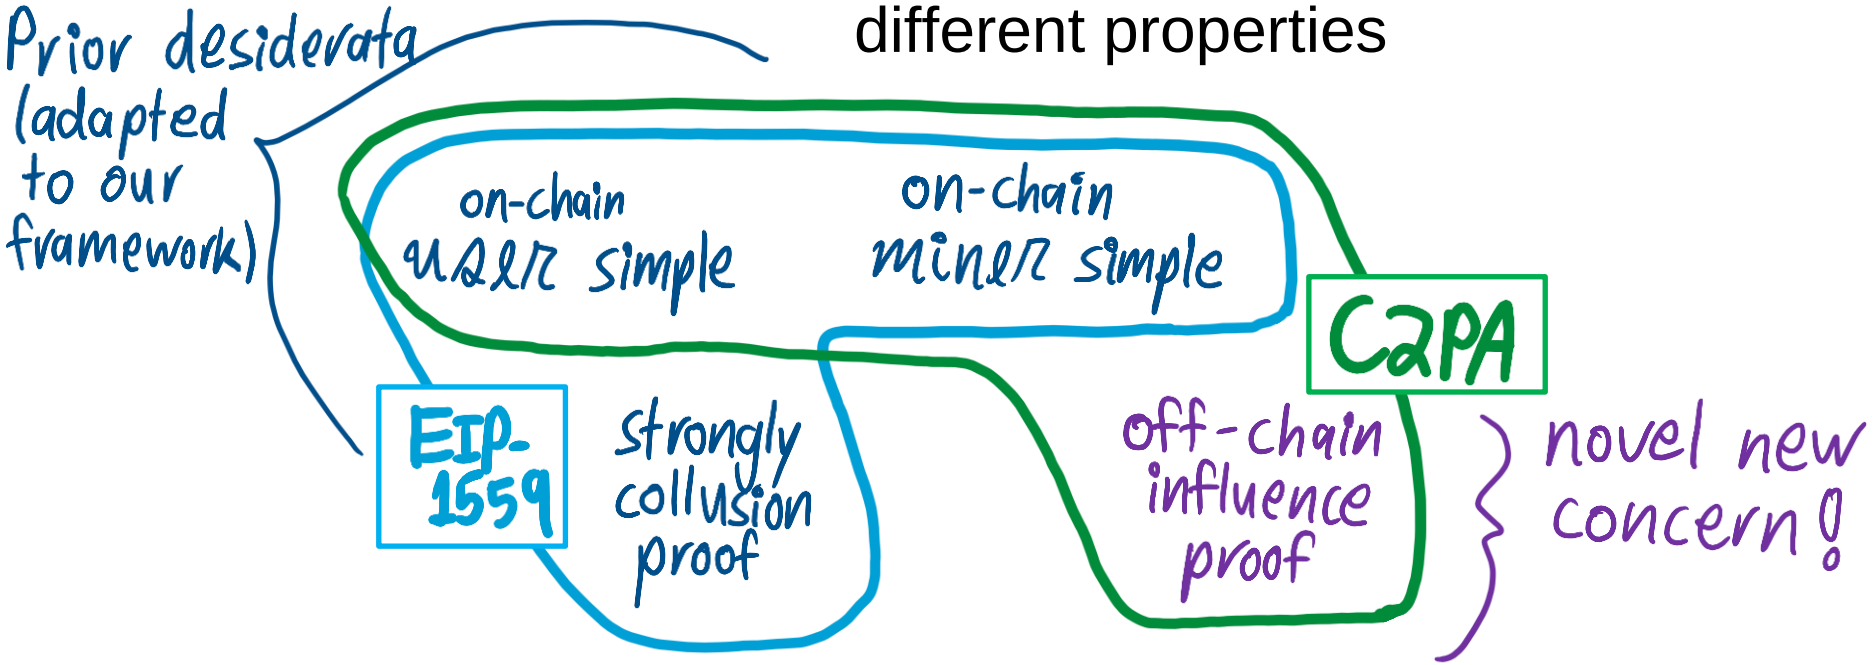

Our desiderata and results: There's a formal tradeoff between these properties

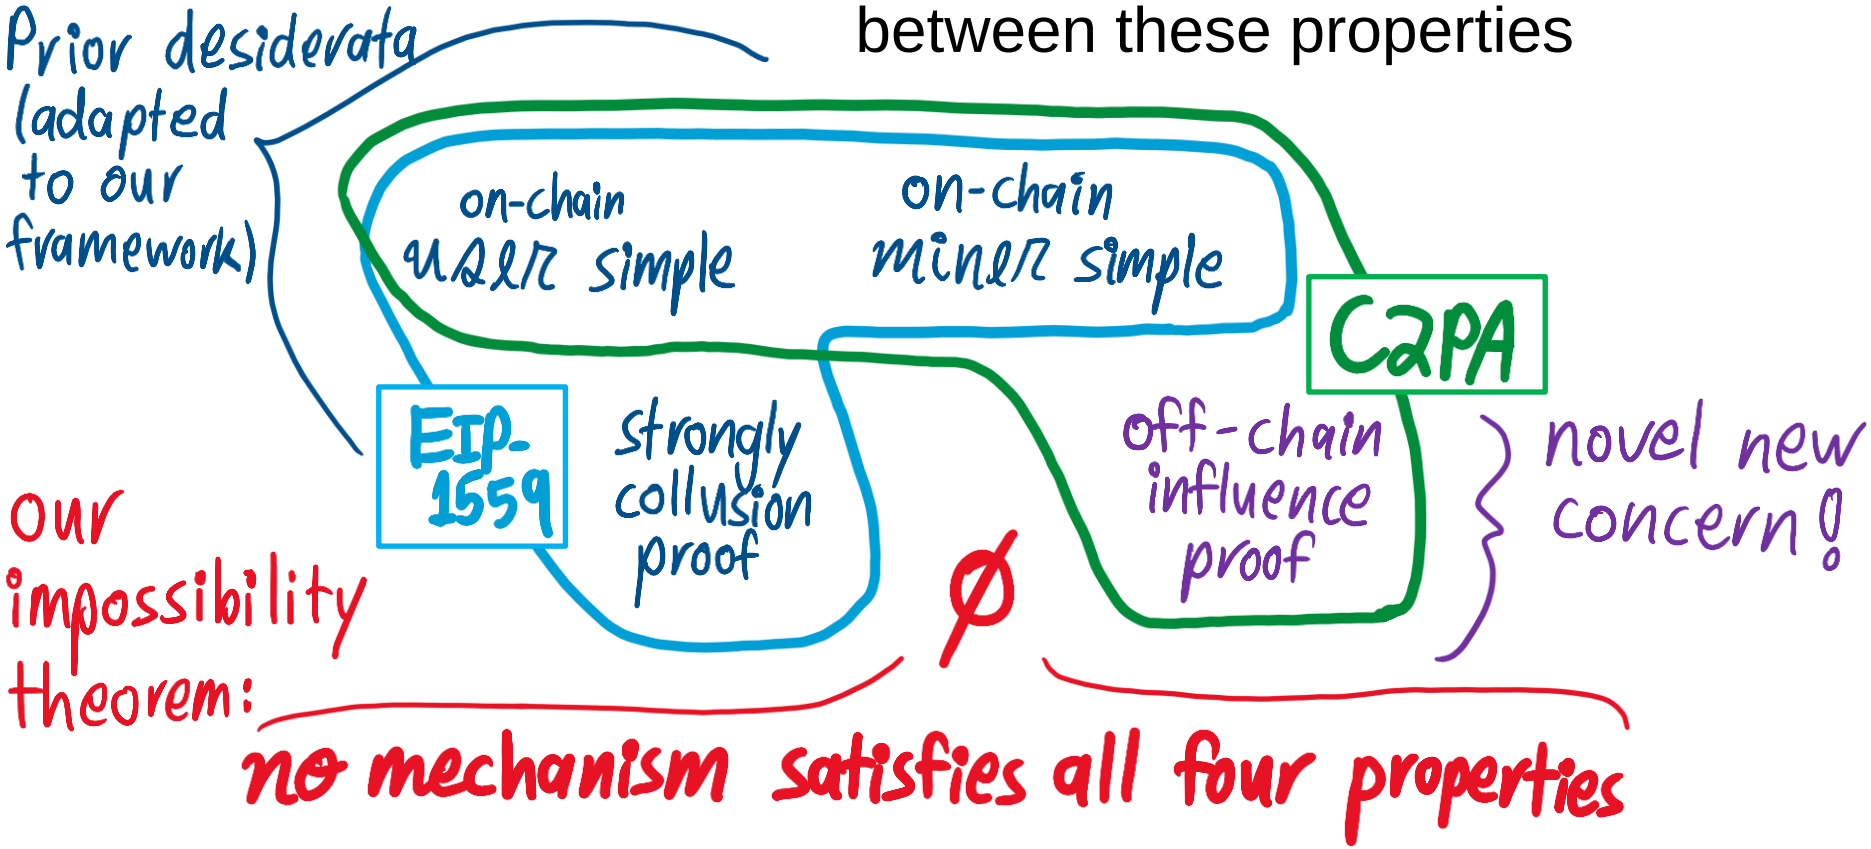

Formal Versions

- Definition: **Strong Collusion Proofness** of an equilibrium $\sigma = (s_{miner}, s_{users})$
  - For every user i and every value v_i, and for every $(\widetilde{s_m}, \widetilde{s_i})$, the sum of user i's utility and the miner's profit **cannot** be higher in $(\widetilde{s_m}, \widetilde{s_i}, s_{users\,-i})$ than in σ (≈ 1-SCP [Chung, Shi '23])
- **Theorem:** No nontrivial TFM (i.e., one that allocates with positive probability) satisfies all four of: on-chain (user and miner) simplicity, off-chain influence proofness, and strong collusion proofness.

How much collusion resistance does C2PA lose?

- By adopting C2PA, you must give up strong collusion proofness
- As an aside, we prove two observations regarding collusion in C2PA

1. Requires “true profit sharing”
    - If colluders agree on a specific “profit sharing contract”, and the user best-responds to this contract, then revenue equivalence already applies ⟹ the miner cannot profit
2. When two “truly profit sharing” agents (e.g. miner submits a transaction) collude optimally, *other users’ best responses is still truth-telling*

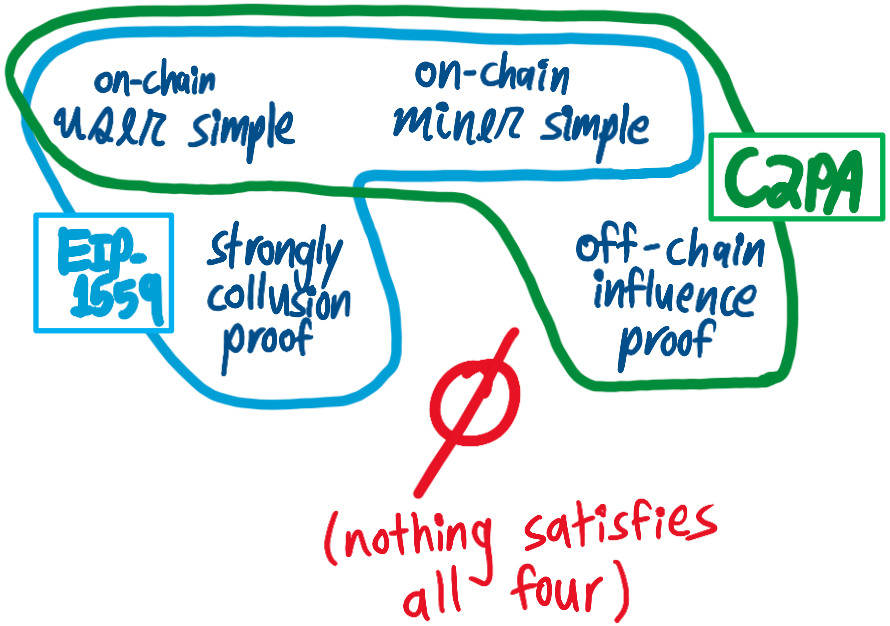

Conclusion

Takeaway: New desiderata ⇒ different desirable mechanisms

- Off-Chain Influence Proof: account for miner running an off-chain
  - ⇒ EIP-1559 may not be a “dream” (even with unlimited supply)
  - Still: Attack has never been observed in practice
  - Still: EIP-1559 has advantages like easy cryptography & predictability
- Might as well allow input from miner, e.g. setting a reserve
  - ⇒ Cryptographic Second Price Auctions worth (re)considering