On the impossibility of surviving (iterated) deletion of weakly dominated strategies in rational MPC.

 TCC 2023 ia.cr/2022/1762

PADERBORN UNIVERSITY

Johannes Blömer, Henrik Bröcher

ZK Lab

THE UNIVERSITY of EDINBURGH

Jan Bobolz

Traditional MPC

Secure mooltiparty computation

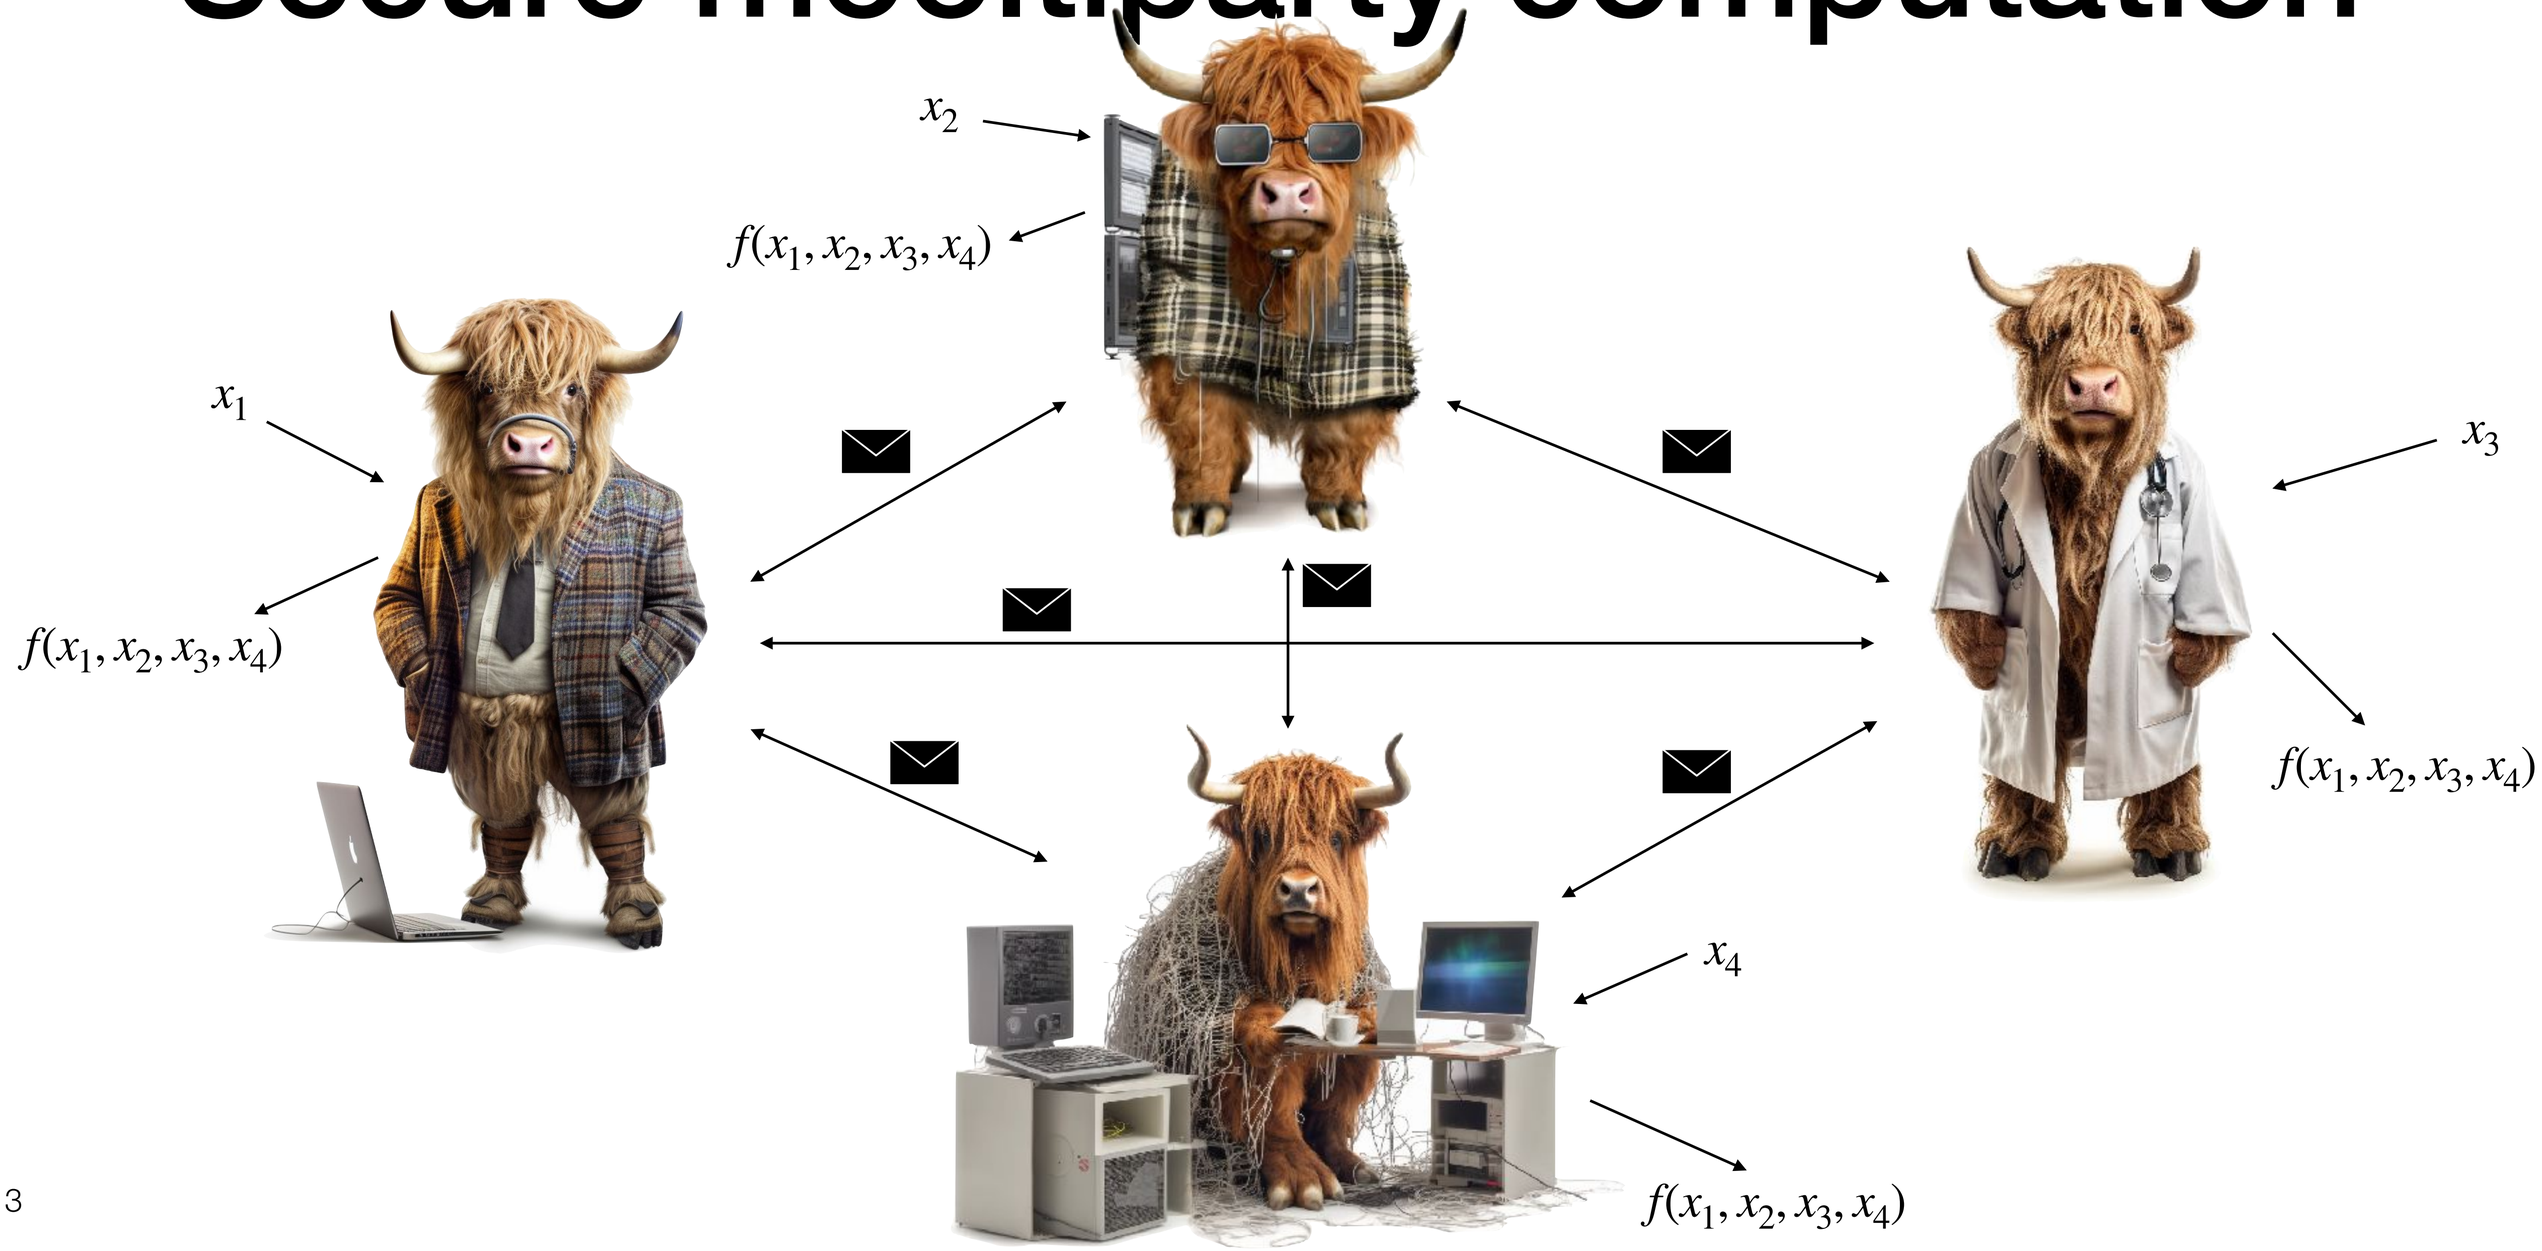

Secure mooltiparty computation

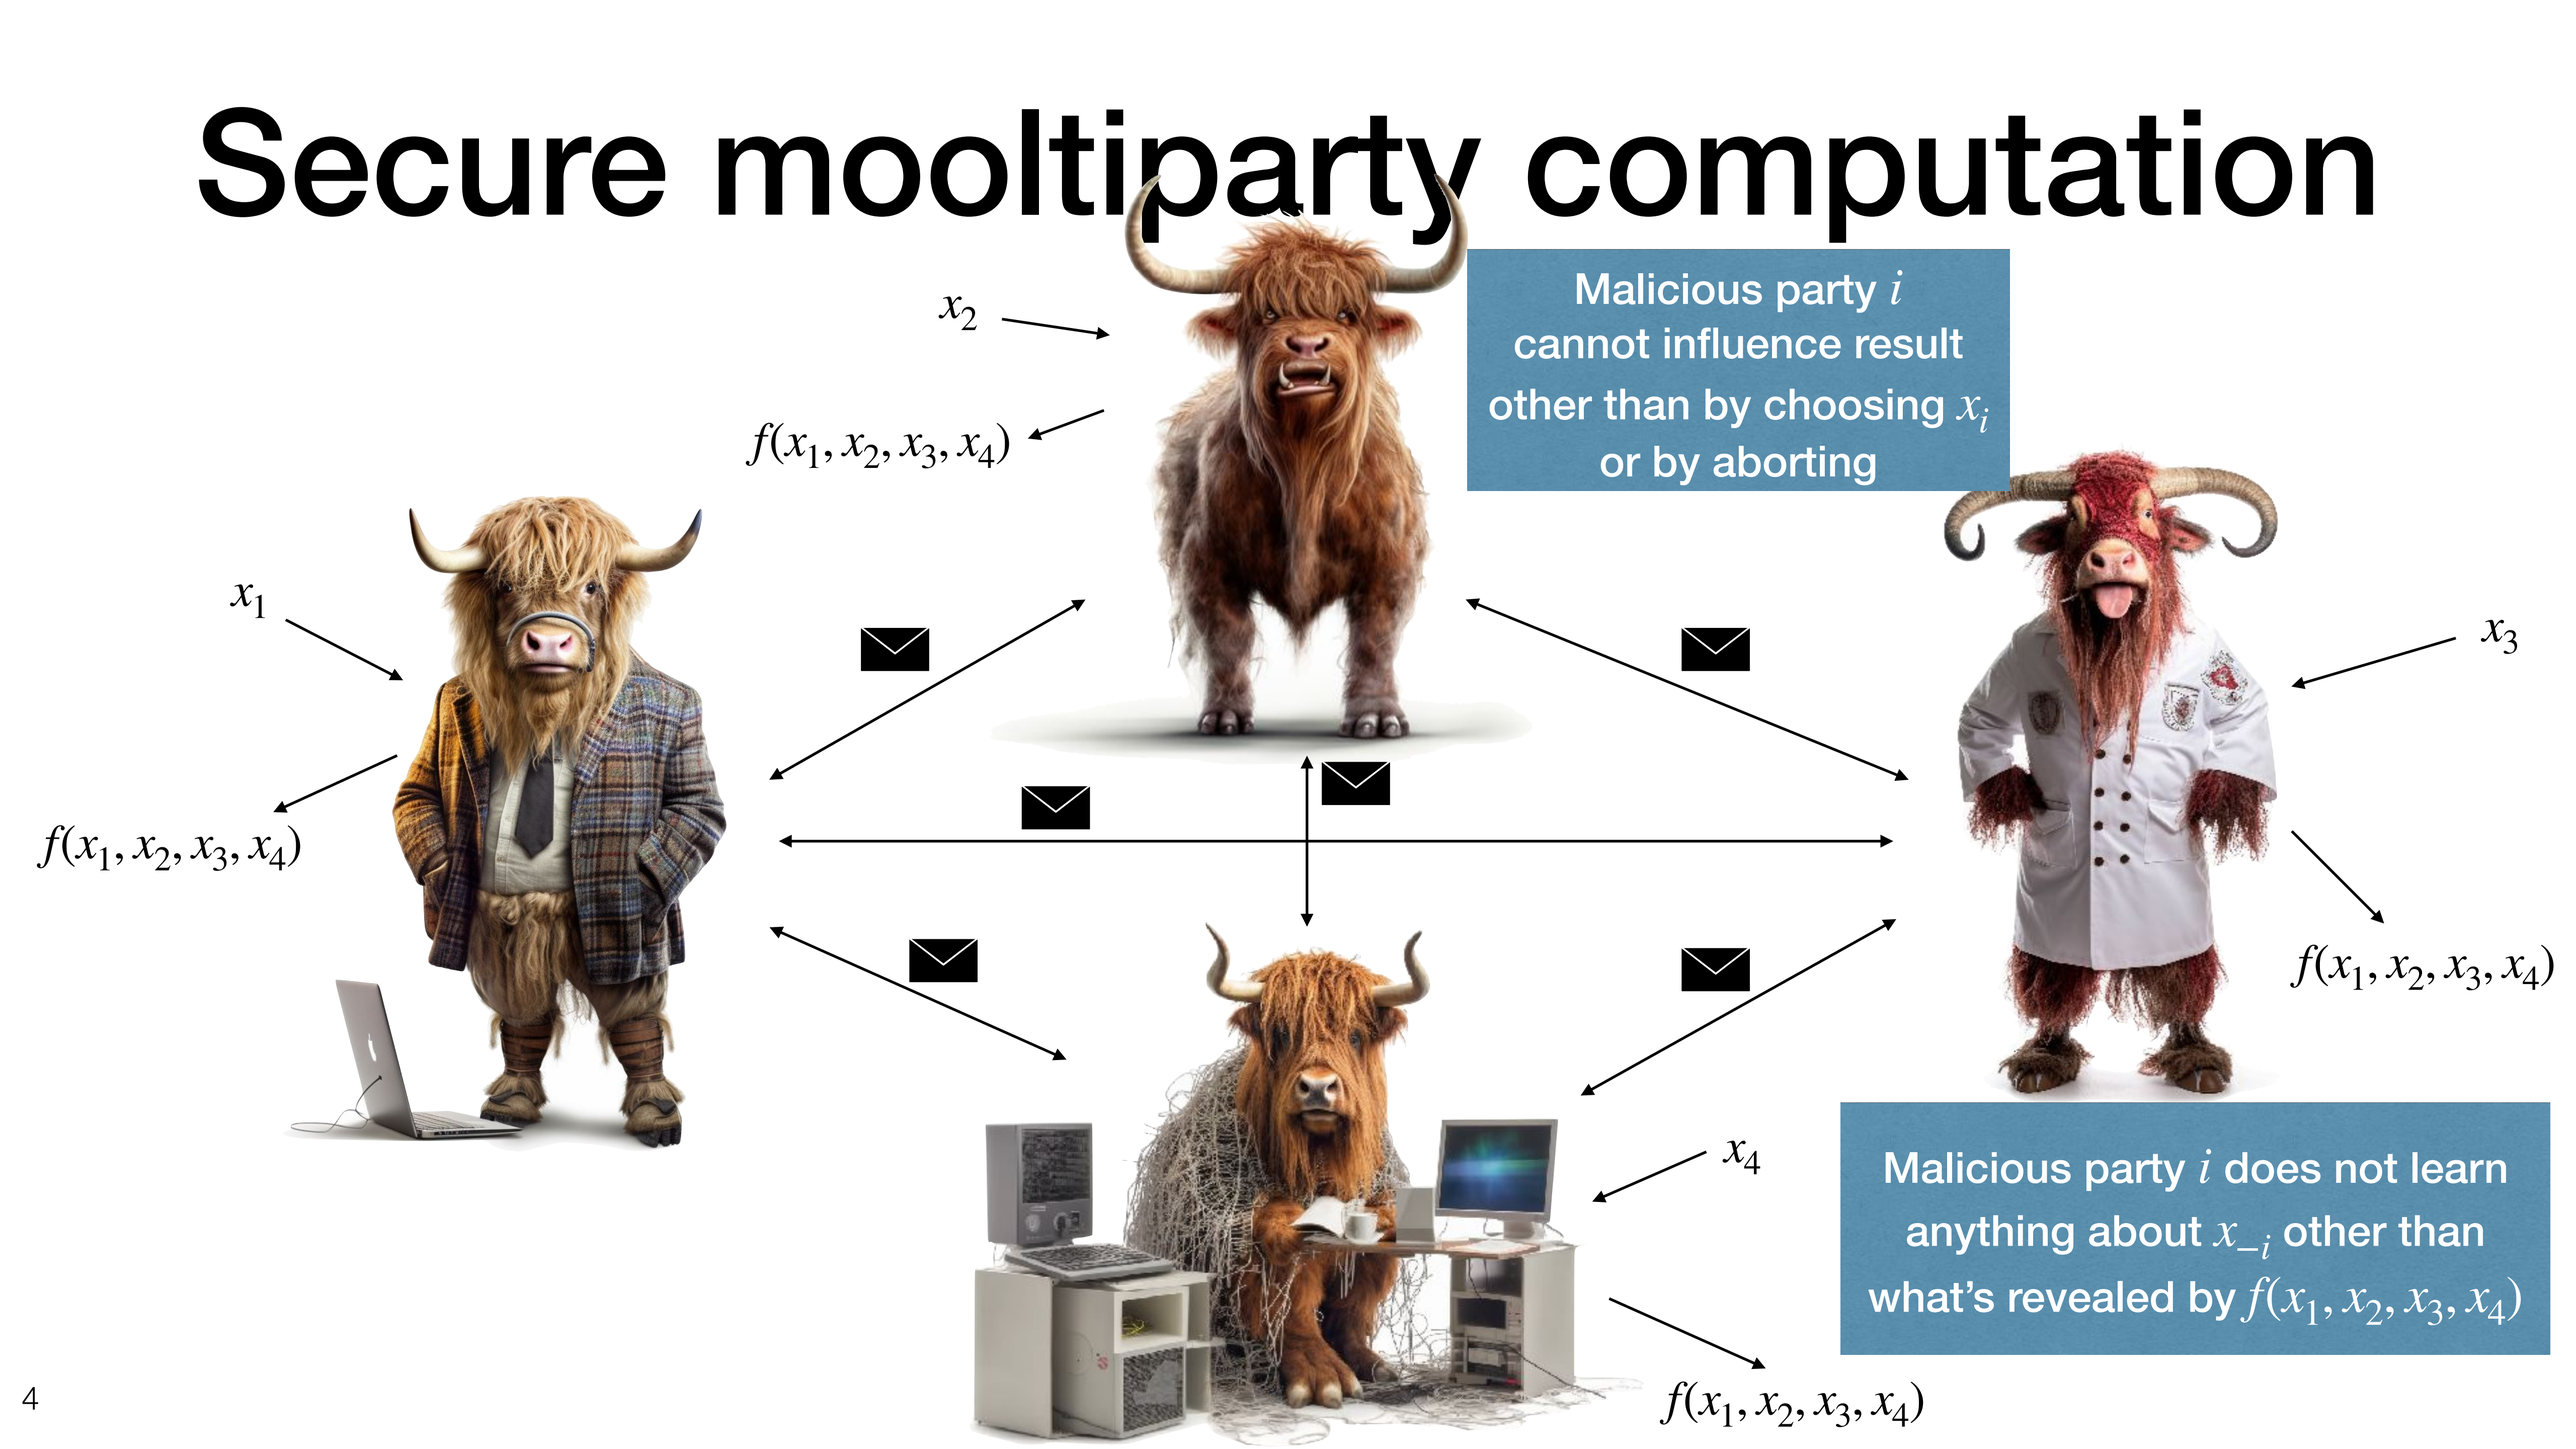

Criticism on standard MPC definition

- 💪 Too strong: protects against *irrational* attacks

  ⇒ Fairness impossible (2 party)

- 🤷‍♂️ Too weak: assumes that at least one party *altruistically* follows the protocol

  - What if *nobody* is purely *honest*?

Rational MPC

Rational MPC

- n **players**.
- Each player i chooses a **strategy** s_i.
- (A security parameter λ is chosen).
- Then the **game** commences.

s_i **is a ppt interactive Turing machine**

1. **Choose random input** $(x_1, \ldots, x_n) \leftarrow I(\lambda)$
2. **Run all the** $s_i(x_i)$**, interacting with each other (sync rounds)**
3. **Each** s_i **outputs some** y_i

- Each player i tries to maximize expected **utility**
  $u_i(\lambda, s_1, \ldots, s_n) := \mathrm{E}_{\lambda,s}[u_i'(x_1, \ldots, x_n, y_1, \ldots, y_n)]$

Utility depends on input/output

Mechanisms

- n **players**.
- Each player i chooses a **strategy** s_i.
- (A security parameter λ is chosen).
- Then the **game** commences.

A *mechanism* $(s_1, \ldots, s_n)$ **is a strategy prescription (algorithm) for the players.**

Intuitive security expectation:
"it's rational for player i to play s_i"

- Each player i tries to maximize expected **utility**

$$u_i(\lambda, s_1, \ldots, s_n) := \mathrm{E}_{\lambda,s}[u_i'(x_1, \ldots, x_n, y_1, \ldots, y_n)]$$

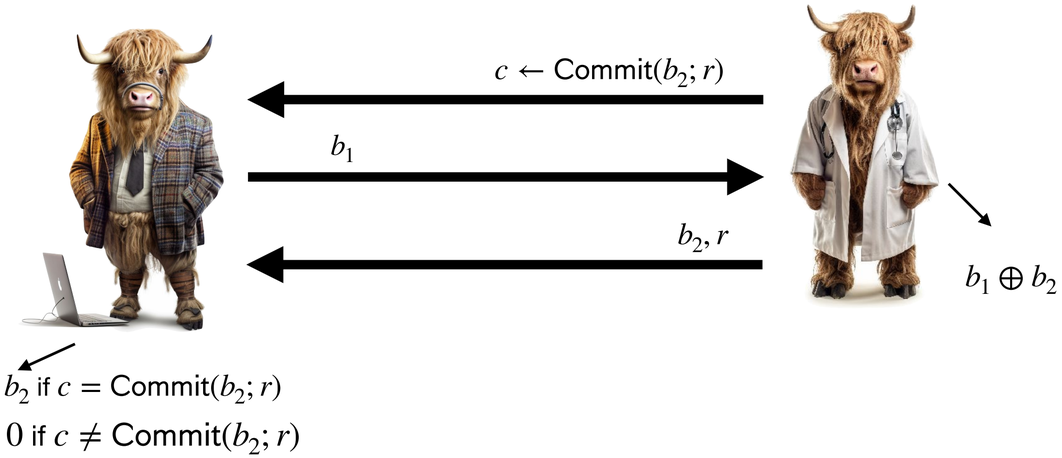

Nash equilibria

- When is a mechanism *good*?
  If no player has an incentive to deviate from their prescribed strategy.

A strategy profile $s = (s_1, \ldots, s_n)$ is a *Nash equilibrium* if no player gains (non-negligibly) by unilaterally deviating:

$$\forall i \ \forall s_i' \ \exists negl \ \forall \lambda : \ \ u_i(\lambda, s_i, s_{-i}) \geq u_i(\lambda, s_i', s_{-i}) - negl(\lambda)$$

Building rational MPC protocols

Blueprint for rational MPC

- Step 1: Make sure function/utilities are such that lying about input is irrational.
- Step 2: Use GMW-style maliciously secure MPC to compute the function value.
  At the end, everyone holds a secret-share of the result + proofs of correctness.
- Step 3: Use rational secret-reconstruction protocol so that everyone learns the result.

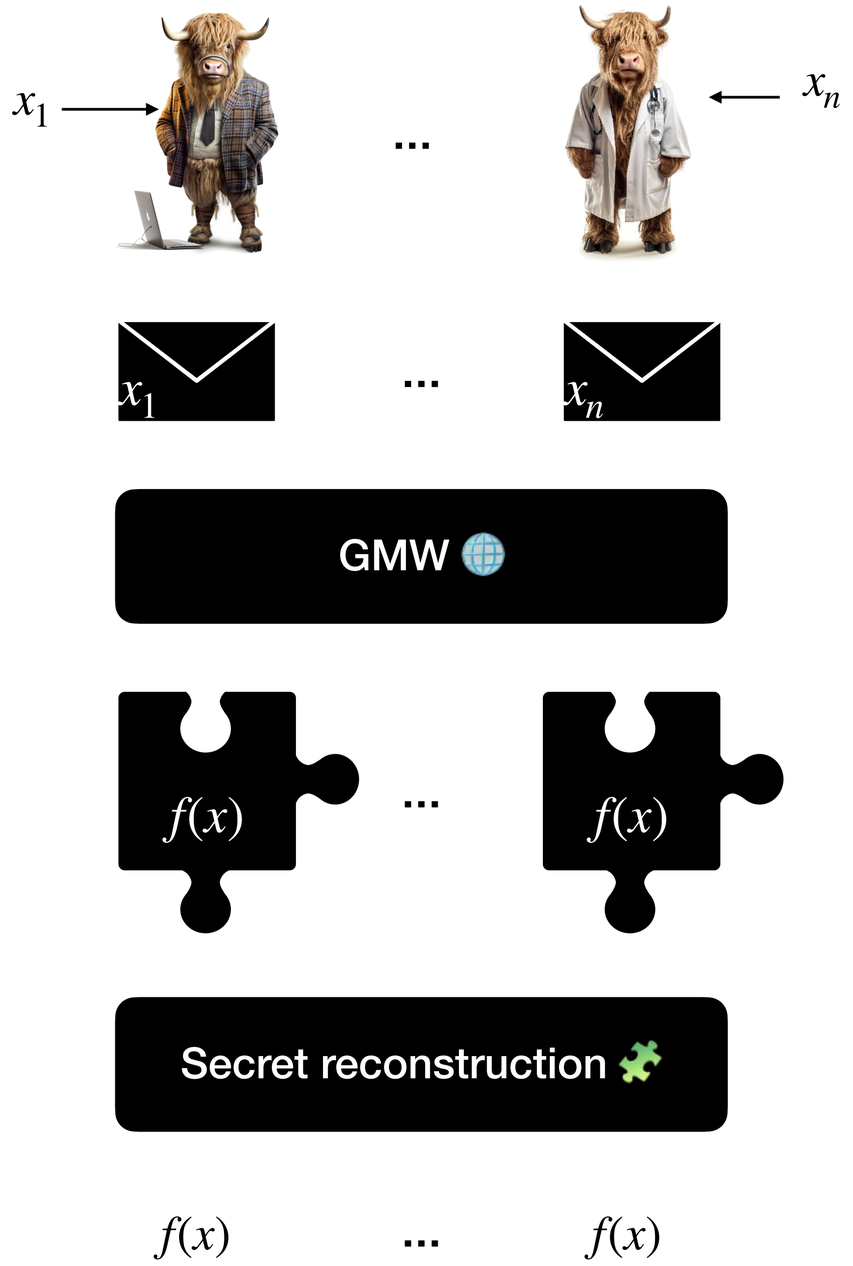

Secret reconstruction

The secret reconstruction scenario

Secret reconstruction:
MPC with $f(x_1, \ldots, x_n) = \text{ReconstructSecret}(x_1, \ldots, x_n)$

- n **players**.
- Each player i chooses a **Turing machine** s_i.
- (A security parameter λ is chosen).
- Then the **game** commences.
  1. Dealer chooses **random secret** y^* (known distribution).
  2. Dealer shares y^* with **authenticated secret-sharing**
  3. Inputs $x_i = (pk, share_i, \sigma_i)$.
  4. Run all the $s_i(pk, share_i, \sigma_i)$, interacting with each other (sync rounds).
  5. Each s_i outputs some y_i, hoping $y_i = y^*$.
- Each player i tries to maximize expected **utility**
  $u_i(\lambda, s_1, \ldots, s_n) := \mathrm{E}_{\lambda,s}[u'_i(x_1, \ldots, x_n, y_1, \ldots, y_n)]$

Utilities: any, as long as:

(1) *correctness*: i prefers $y_i = y^*$ over $y'_i \neq y^*$.

$$u'_i(\ldots, y_i, \ldots) \geq u'_i(\ldots, y'_i, \ldots) + p(\lambda)$$

(2) *exclusivity*: i prefers j *not* to output right result.

$$u'_i(\ldots, y'_j, \ldots) \geq u'_i(\ldots, y_j, \ldots) + p(\lambda)$$

Trivial protocol for secret-reconstruction

$(n-1)$-out-of-n secret sharing.

Protocol s_i: just broadcast input $(share_i, \sigma_i)$, receive all broadcasts, check signatures, if enough valid shares: reconstruct y.

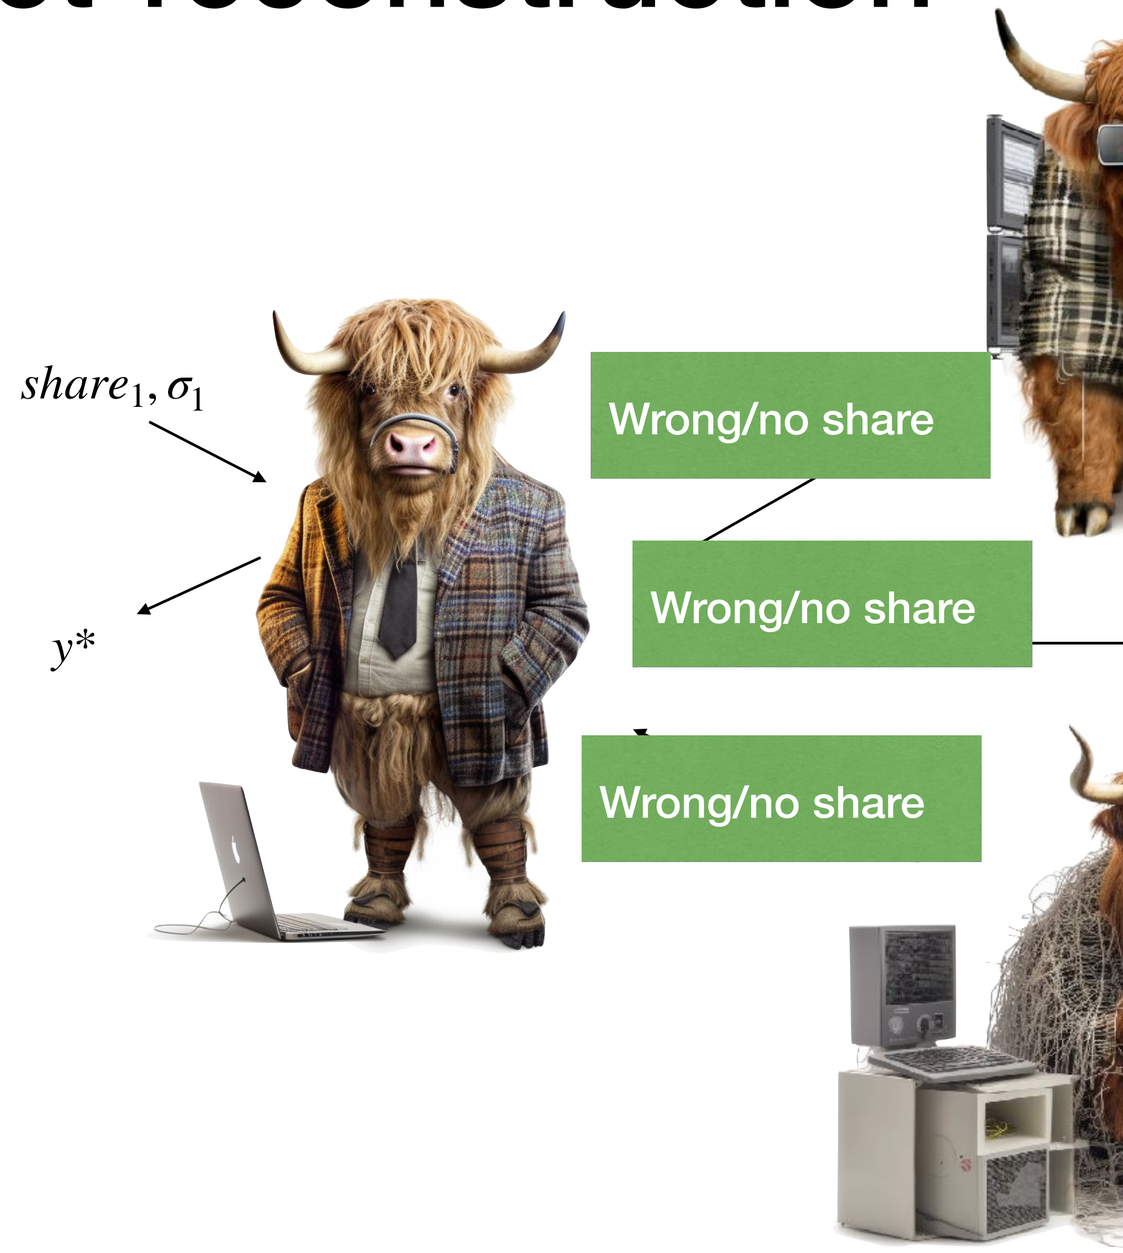

✅ Nash Equilibrium

(1) *correctness*: i prefers outputting y^*.
(2) *exclusivity*: i prefers j *not* to output y^*.

Trivial protocol for secret-reconstruction

n-out-of-n secret sharing.

Protocol s_i: just broadcast input $(share_i, \sigma_i)$, receive all broadcasts, check signatures, reconstruct y from shares, output y.

❌ Nash Equilibrium

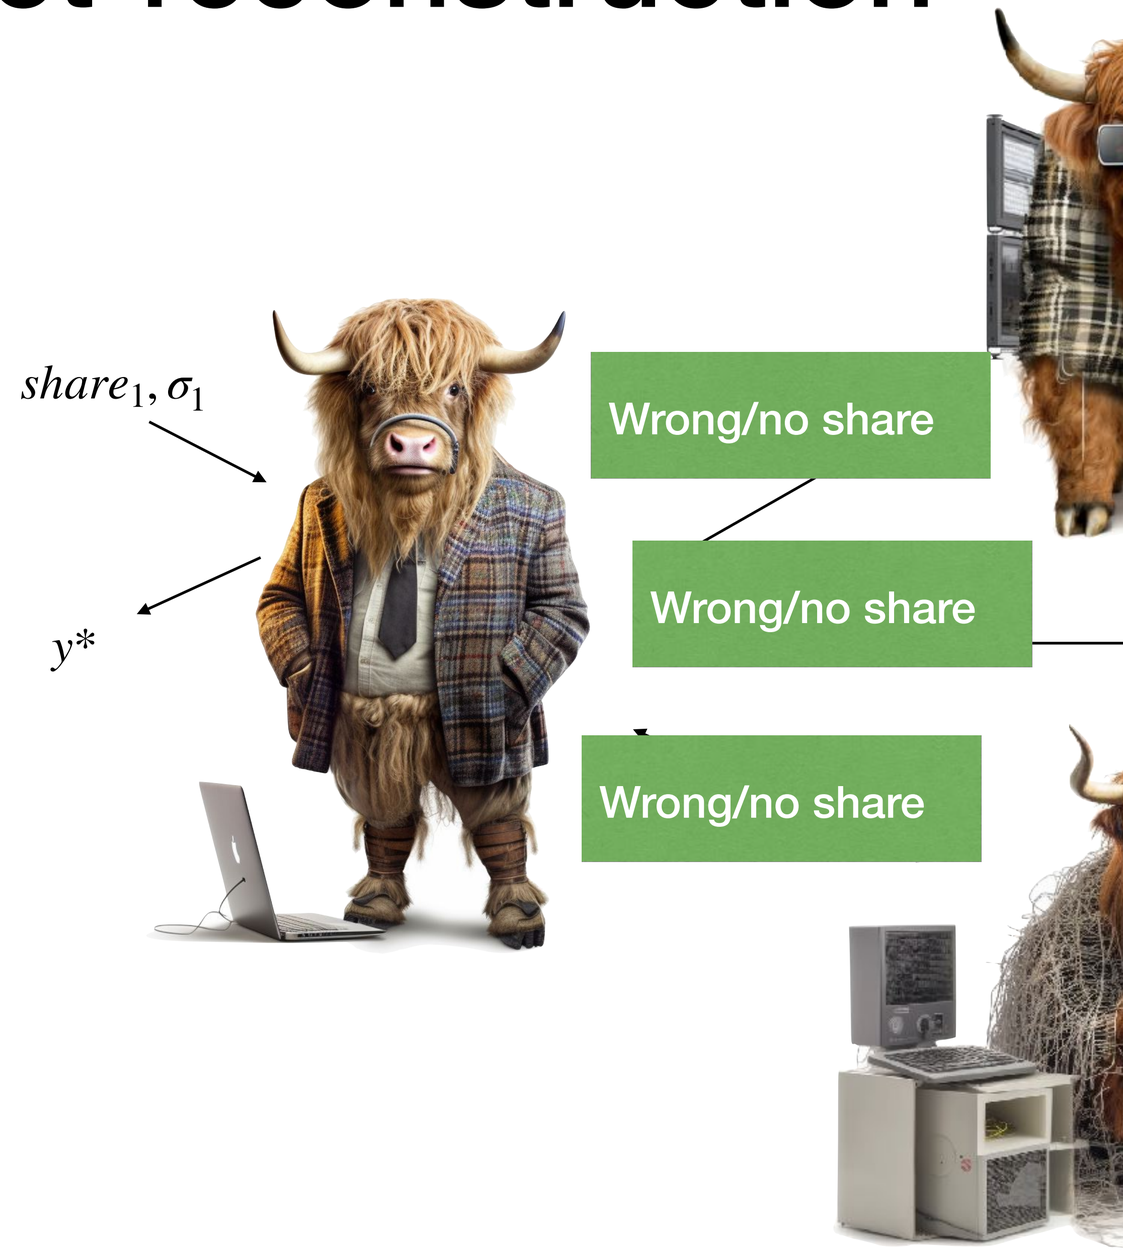

(1) *correctness*: i prefers outputting $y*$.
(2) *exclusivity*: i prefers j *not* to output $y*$.

Secret reconstruction for $n = 2$

Can we make n out of n work?

[Gordon, Katz 2006] and, concurrently, [Abraham, Dolev, Gonen, Halpern 2006]

n out of n reconstruction

If we have a fixed last round in which everyone learns the secret: bad.
So we hide which round is *the one*.

Deviation is punished:
don't broadcast in ⊥ round:
don't learn secret ever.

- Run MPC that will
  - With probability β: produce (auth'd) sharing of y^*
  - With probability $1-\beta$: produce (auth'd) sharing of ⊥
- Everyone broadcasts the new sharing (hopefully)
- If someone doesn't broadcast, abort everything 🔥
- If reconstructed secret is ⊥, repeat. Otherwise, output y^*.

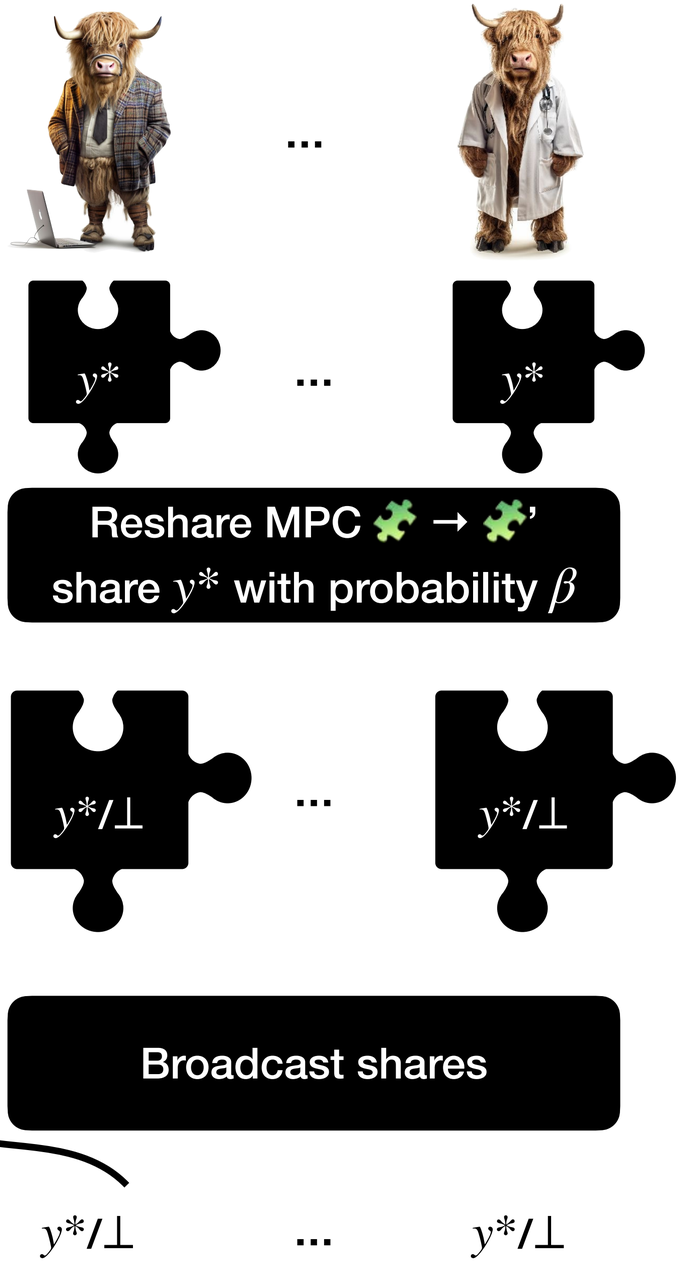

[Gordon, Katz 2006] and, concurrently, [Abraham, Dolev, Gonen, Halpern 2006]

n out of n reconstruction

- The other party definitely sends the shares every time.
- When not broadcasting, don't know whether I'll get y^* or $\perp$.
  - If y^*: I win! Expected utility for deviating: $\beta \cdot U^+$
  - If $\perp$: I'll never learn secret b/c other party stops 🔥.
    Expected utility for deviating: $(1-\beta) \cdot U^-$

Deviate: utility $\beta \cdot U^+ + (1-\beta) \cdot U^-$

Honest: utility U

$U^+ > U > U^-$

Set β **such that** $\beta \cdot U^+ + (1-\beta) \cdot U^- < U$

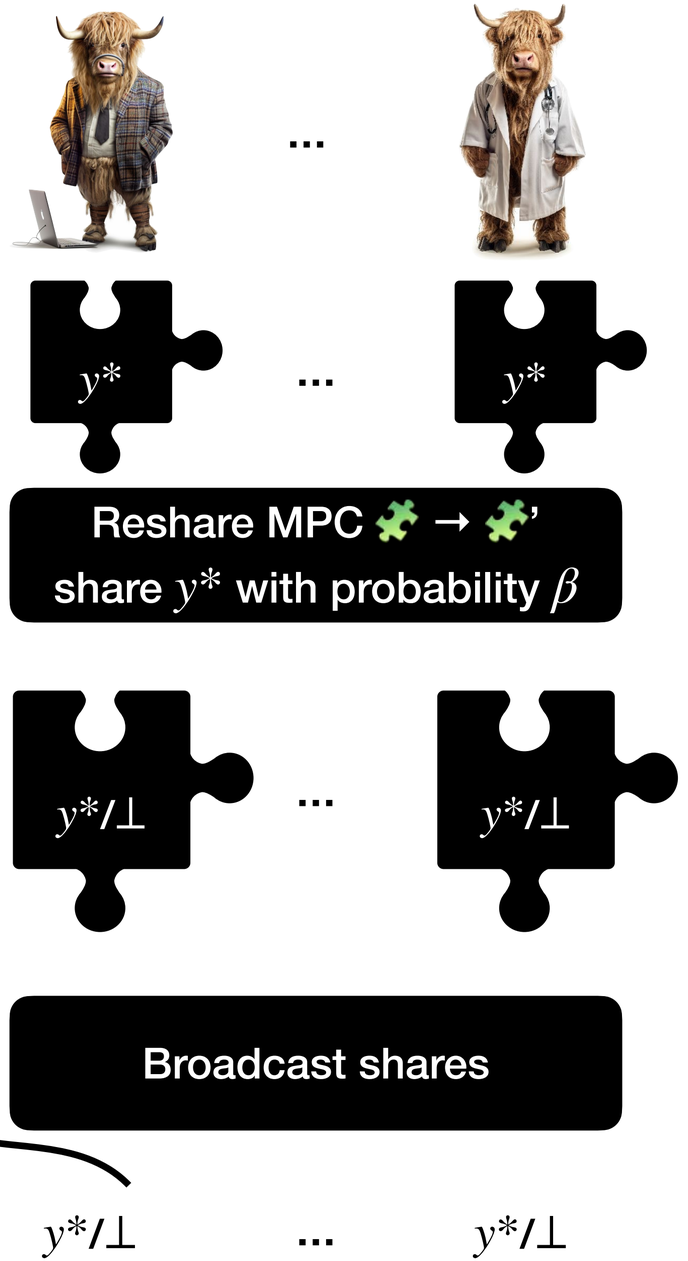

Weakly dominated strategies

a story with twists and turns

Weakly dominated strategies

A strategy s_{better} ***weakly dominates*** s_{bad} for player i if:

Sometimes better: $\exists s_{-i} : \; u_i(\lambda, s_{\text{better}}, s_{-i}) > u_i(\lambda, s_{\text{bad}}, s_{-i}) + p(\lambda)$

Never worse: $\forall s_{-i} : \; u_i(\lambda, s_{\text{better}}, s_{-i}) \geq u_i(\lambda, s_{\text{bad}}, s_{-i}) - negl(\lambda)$

Idea: why would I play s_{bad} if s_{better} could be better in some scenarios and is never really worse?

⇒ additional criterion for "is this mechanism *good*?":

Nobody is expected to play weakly dominated strategy.

The spirit of weak domination

A strategy s_{better} ***weakly dominates*** s_{bad} **for player** i **if:**

Sometimes better: $\exists s_{-i}: \; u_i(\lambda, s_{\text{better}}, s_{-i}) > u_i(\lambda, s_{\text{bad}}, s_{-i}) + p(\lambda)$

Never worse: $\forall s_{-i}: \; u_i(\lambda, s_{\text{better}}, s_{-i}) \geq u_i(\lambda, s_{\text{bad}}, s_{-i}) - negl(\lambda)$

✅ Nash equilibrium

❌ Weak domination

$(n-1)$-out-of-n secret sharing.

Protocol:

- Do nothing in round 1.
- If *everyone else* has broadcast "destroy" in round 1, stop and output invalid secret.
- Round 2: Broadcast input.
- Round 3: Reconstruct secret and output.

Overview of history

2004 2006 2008 ... 2022

Halpern-Teague: “broadcast-share is weakly dominated, here’s how to circumvent it.”

Gordon-Katz, Abraham-Dolev-Gonen-Halpern, Lysyanskaya-Triandopoulos: “Here are better workarounds.”

Kol-Naor: “actually, *nothing* is ever weakly dominated, there’s always a savior strategy.” [in restricted model]

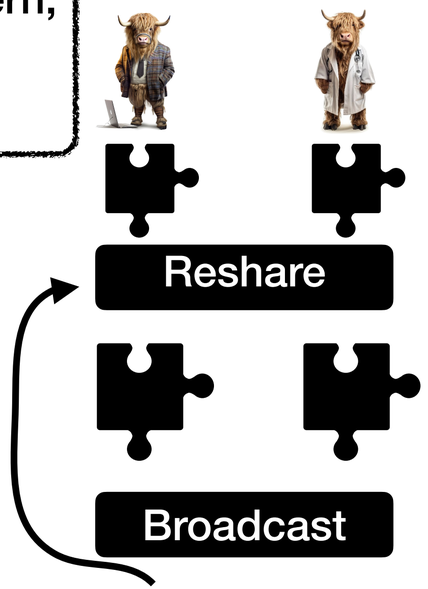

We: “actually, *all* strategies are weakly dominated if the shares are authenticated”

Our impossibility result

No mechanism survives weak domination

An arbitrary mechanism

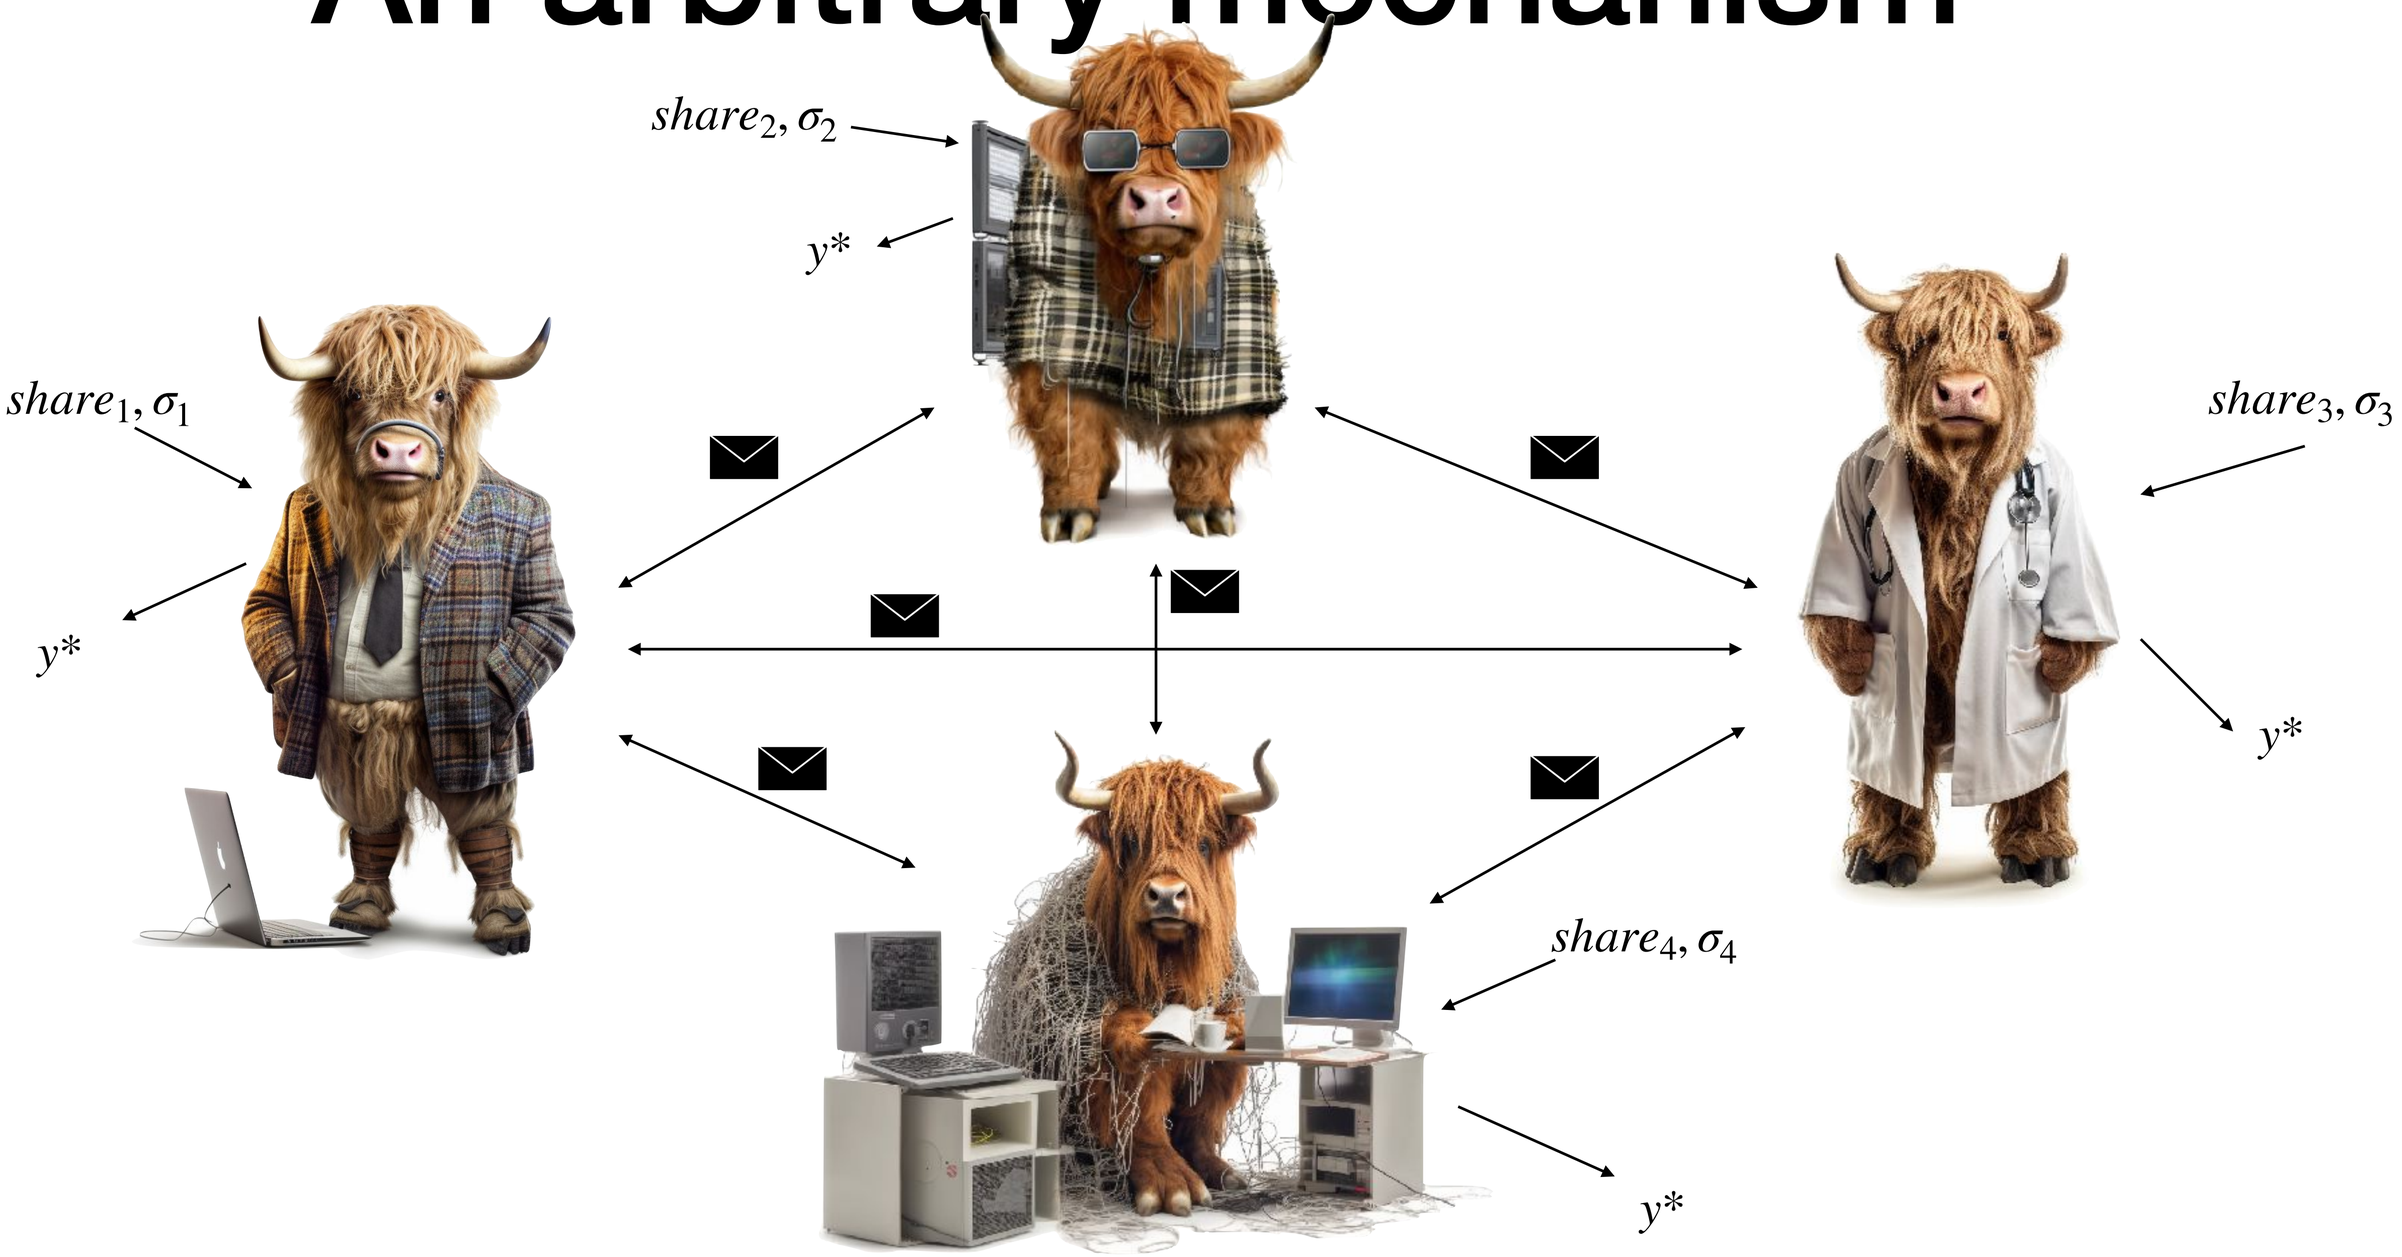

💡 A weakly dominating strategy

Behave exactly as original strategy, but at the end:

If in first round, every other party sent "(leak, $share_j, \sigma_j$)" with valid signatures, then reconstruct y^* and output it.

Otherwise, output the y^* that the mechanism would output.

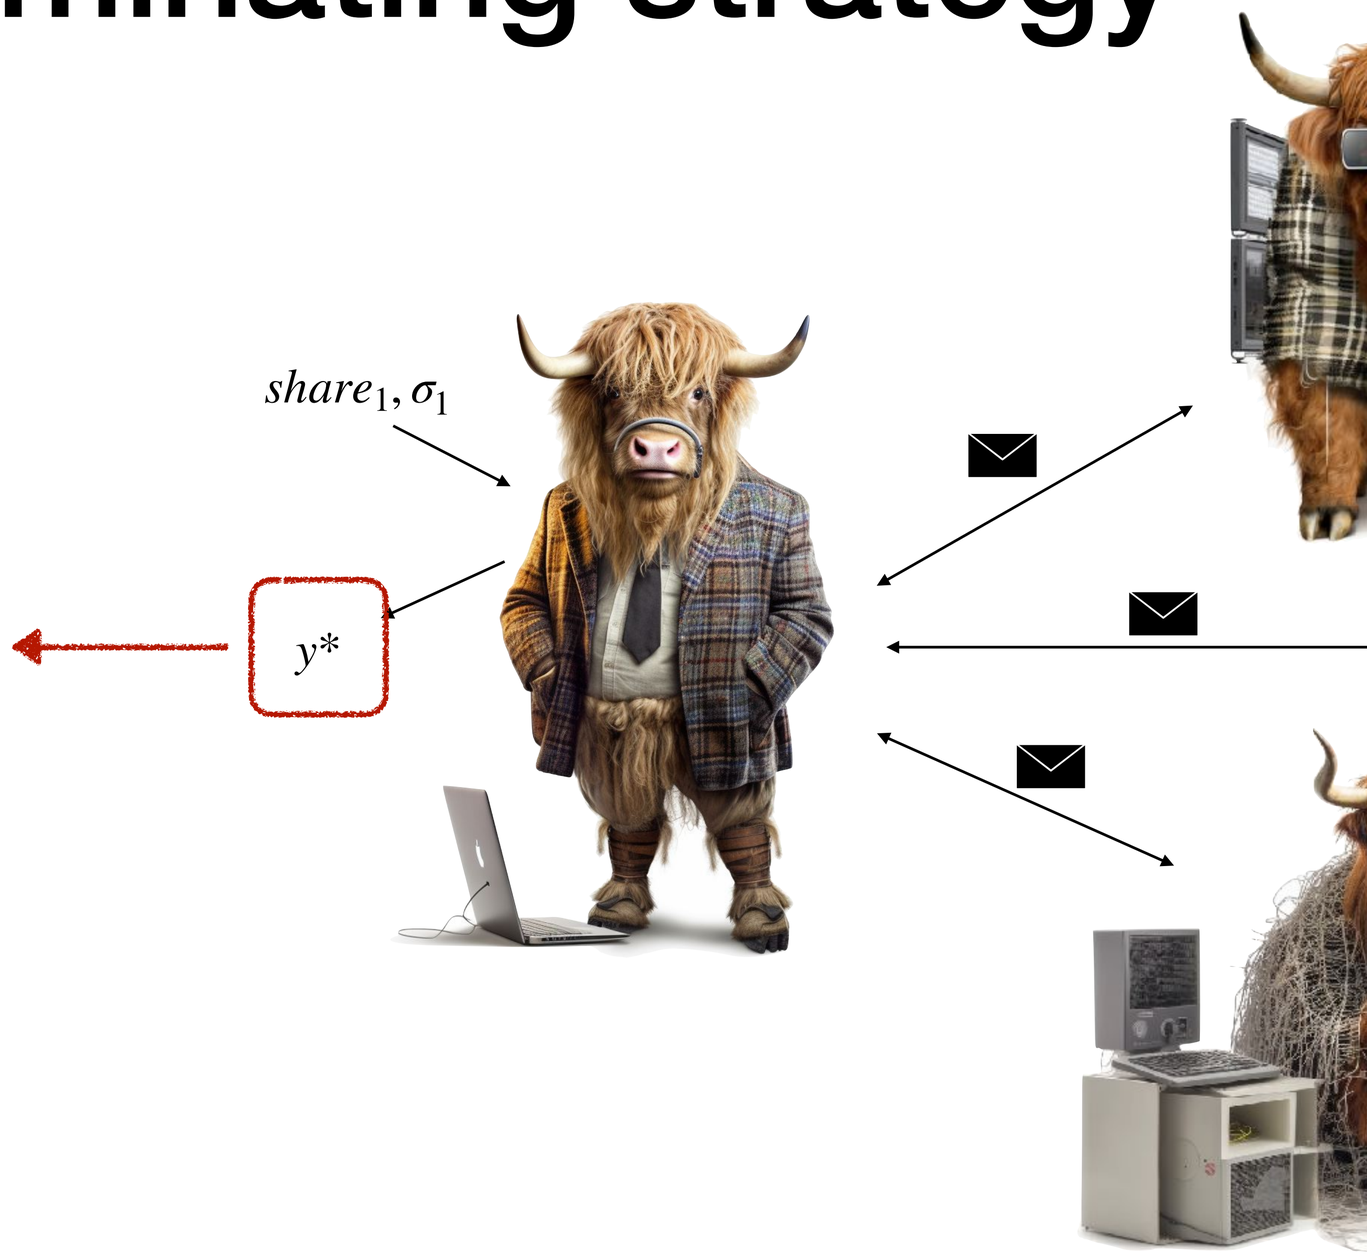

👍 Sometimes better

This strategy outputs the correct secret.

Original strategy outputs error because first message malformed.

If in first round, every other party sent "(leak, $share_j, \sigma_j$)" with valid signatures, then reconstruct y^* and output it.

Otherwise, output the y^* that the mechanism would output.

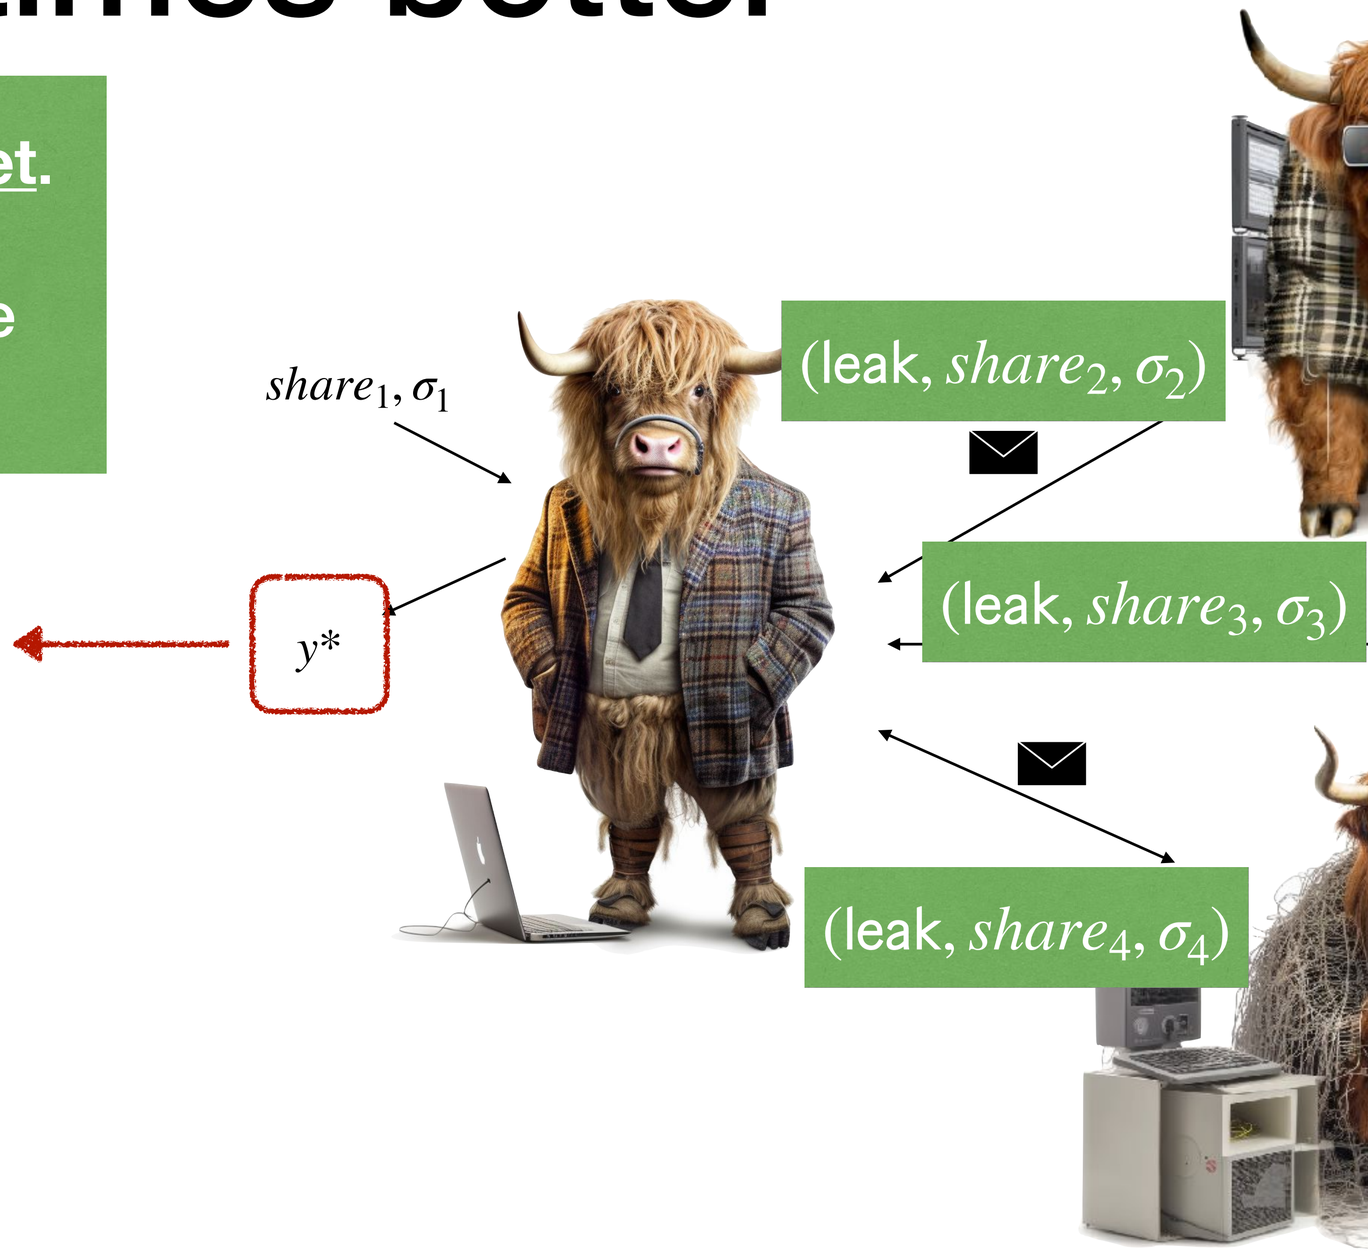

✋ Never worse

No matter what the other strategies do, we **only deviate by outputting the correct secret** (~ unforgeability of σ_j)

If in first round, every other party sent "(leak, $share_j, \sigma_j$)" with valid signatures, then reconstruct y^* and output it.

Otherwise, output the y^* that the mechanism would output.

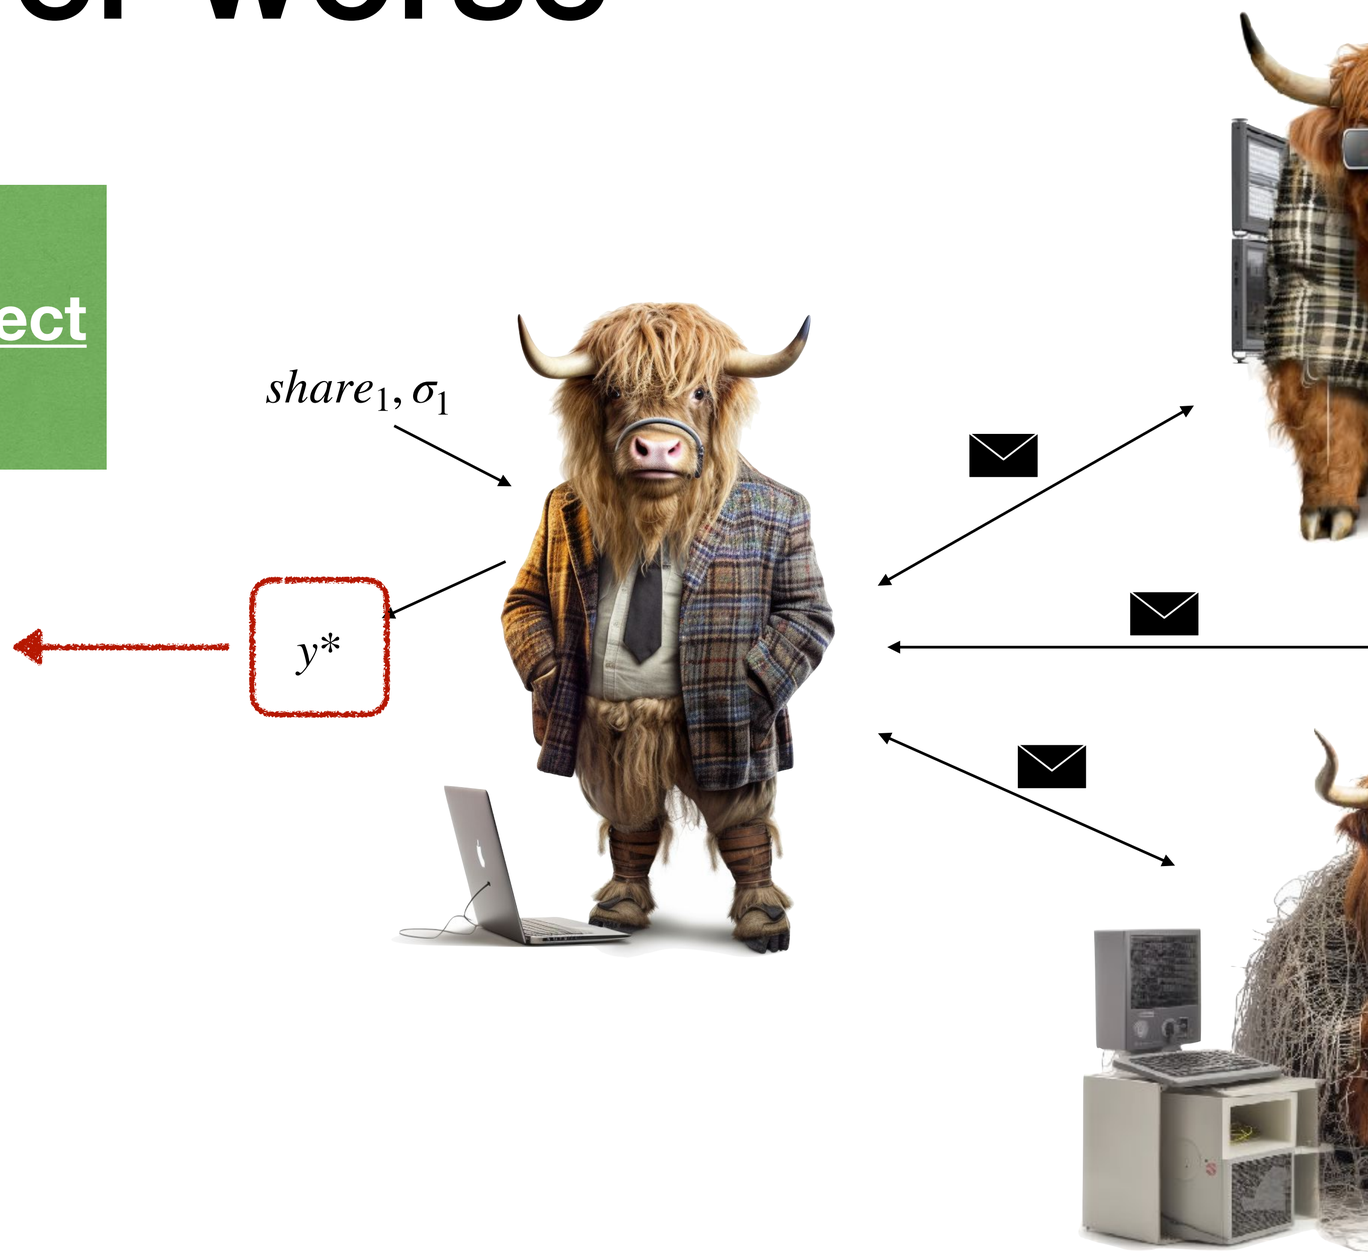

🏁 Conclusion

This strategy weakly dominates any "natural" mechanism strategy
(that doesn't *happen to* parse leak messages).

If in first round, every other party sent "(leak, $share_j, \sigma_j$)" with valid signatures, then reconstruct y^* and output it.

Otherwise, output the y^* that the mechanism would output.

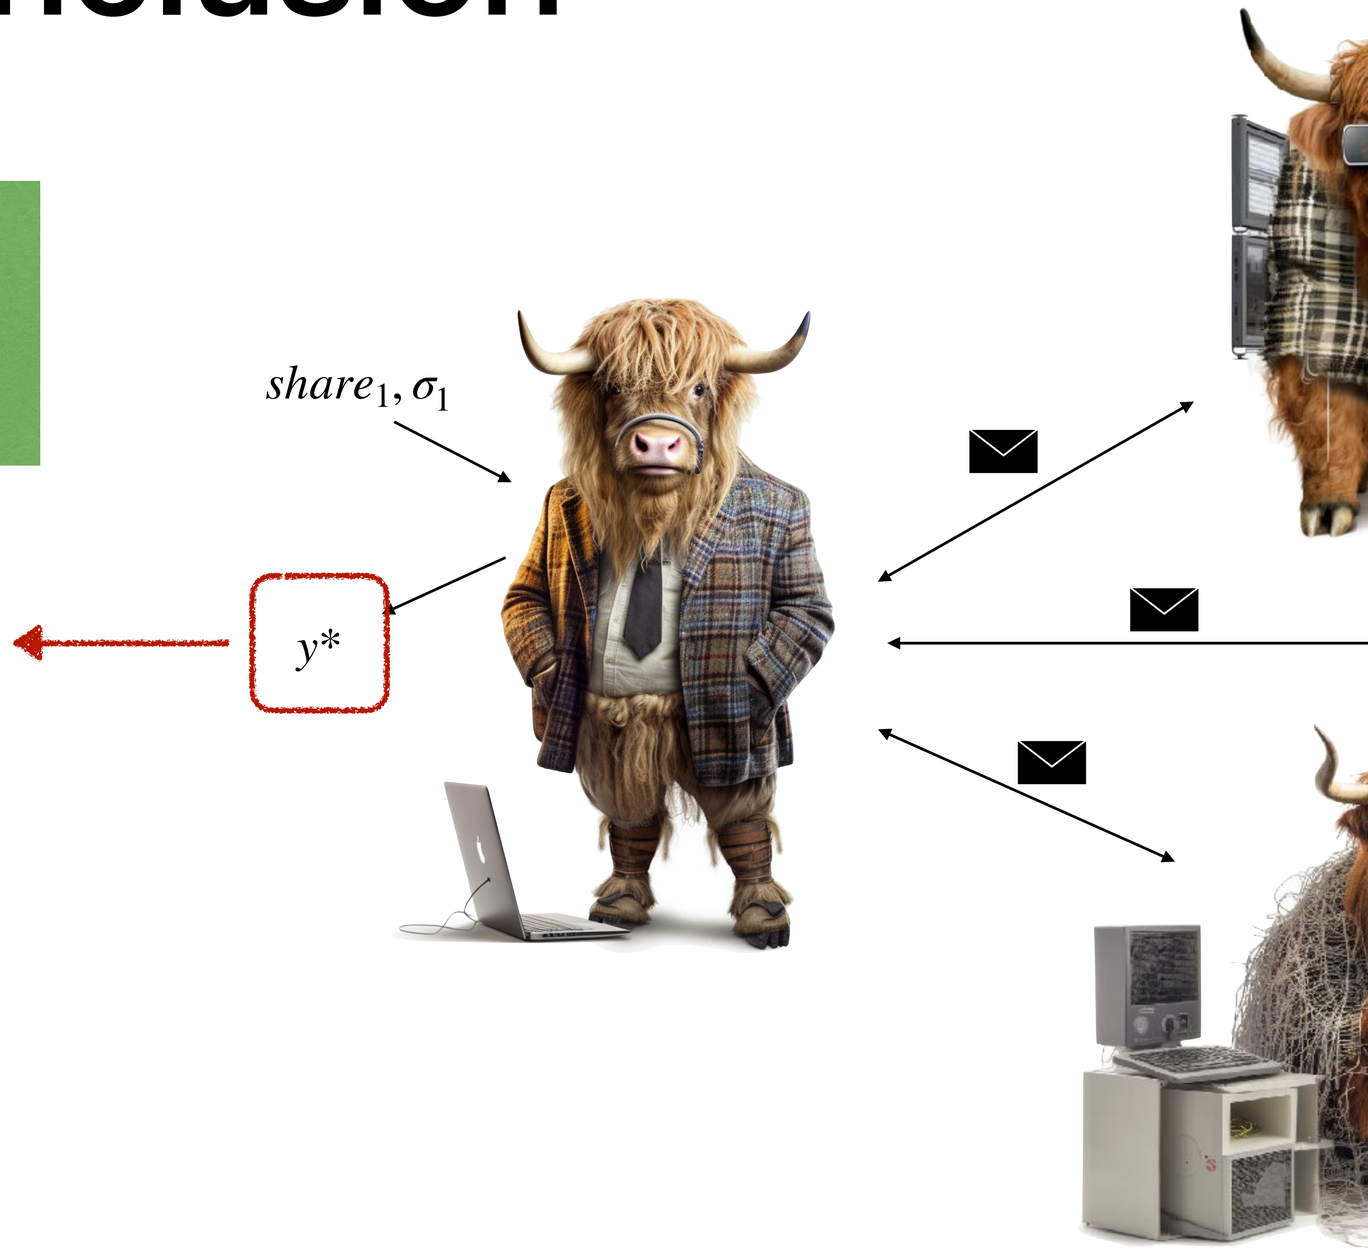

💡💡 An alternative weakly dominating strategy

This strategy weakly dominates any "natural" mechanism strategy
(that doesn't *happen to* parse leak messages).

If in first round, every other party sent "(leak, $\bar{share}_j, \bar{\sigma}_j$)" with valid signatures, then reconstruct y^* and output it.

Otherwise, output the y^* that the mechanism would output.

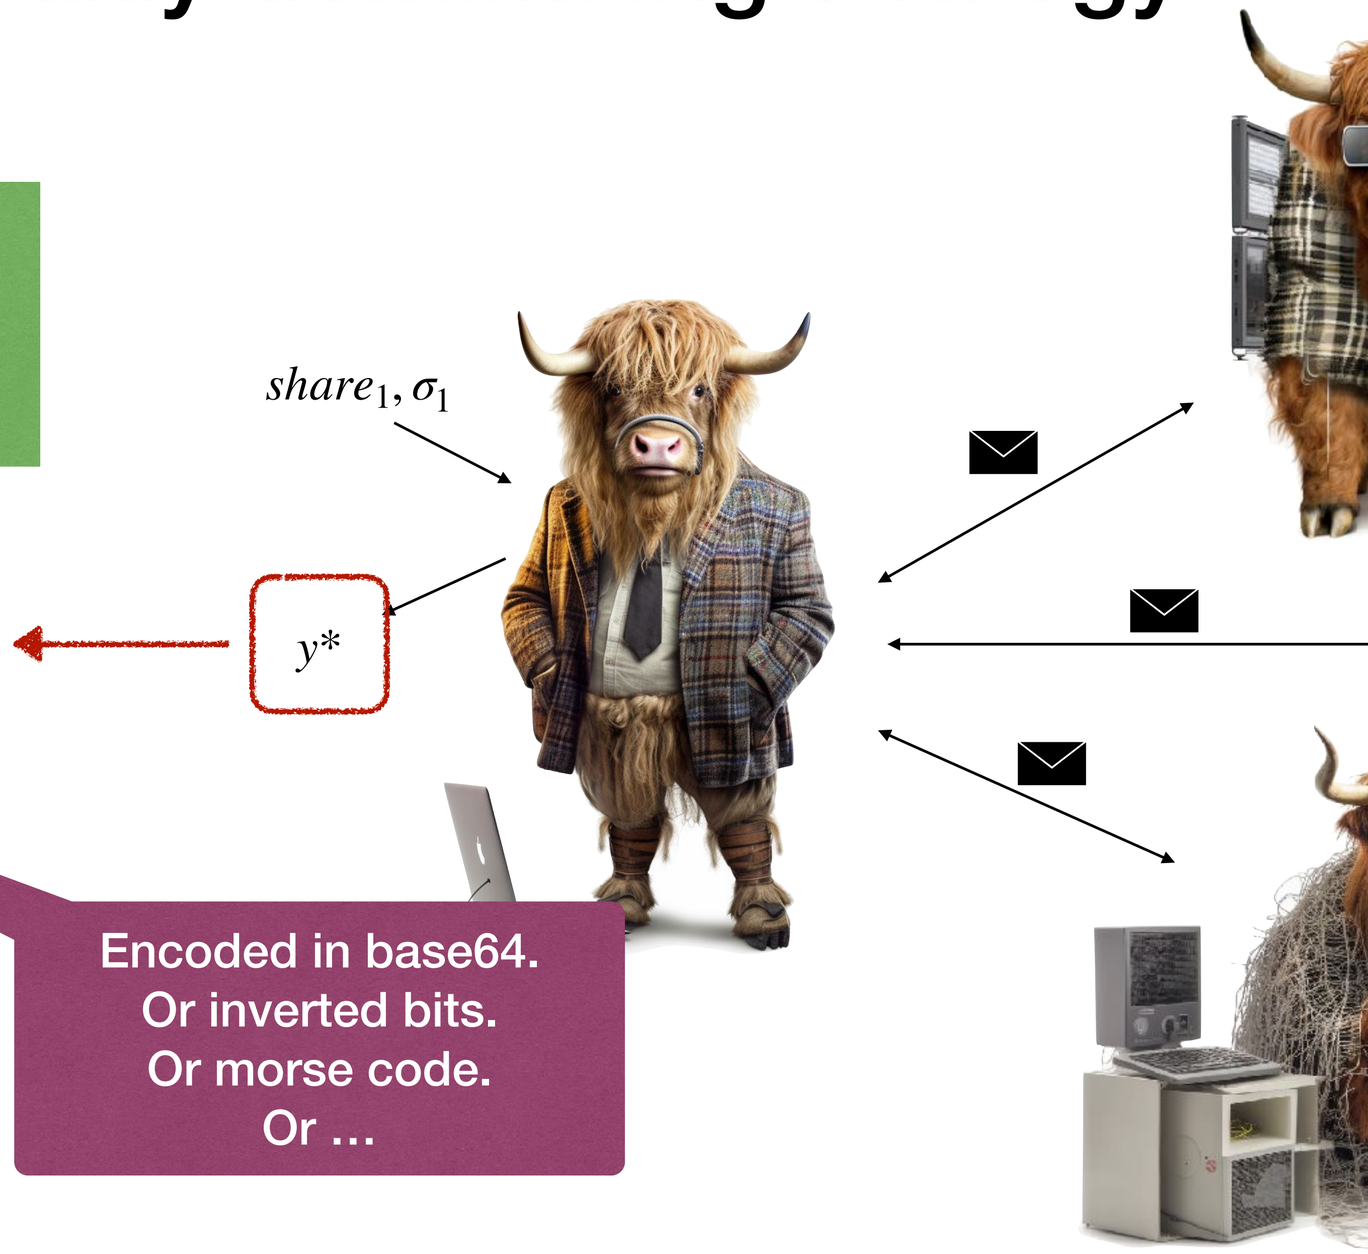

Encoded in base64.
Or inverted bits.
Or morse code.
Or …

Consequences

Every imaginable protocol with authenticated shares is weakly dominated.

Counterexamples: *what if* everyone leaks their input in some weird encoding? 🤯

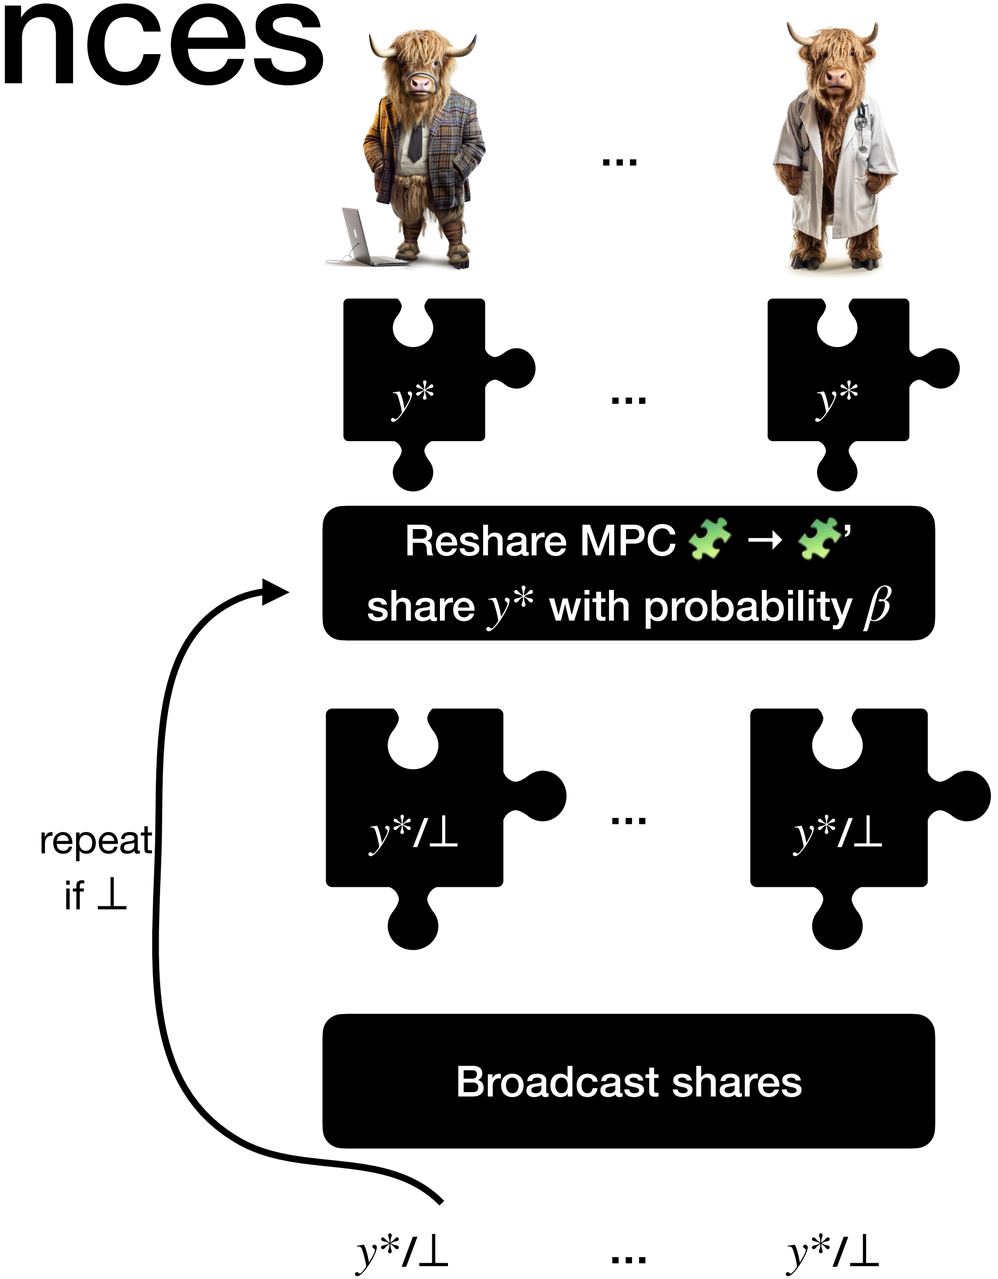

Takeaways

- 👍 Rational MPC enables good properties (fairness). Sometimes better model.
- 👎 Weak domination is too strict.
  - 🧠 We long believed it to be too weak.
  - 🔥 No known protocol survives weak domination
- 🧯 Lots of work to be done.
  - 📚 Replacement for weak domination?
  - 🛠️ Better protocols / settings?

Backup slides

Protocols without locally verifiable shares

🔍 What the counterexample needs

👍 Sometimes better

Against input-leaking strategies, we do better.

🖐 Never worse

Only deviation: output correct result.

What if I cannot check whether result is correct?

🤔 But without authentication, parties may lie about their shares

[Abraham, Dolev, Gonen, Halpern 2006]

ADGH06's third protocol

Idea: instead of signing the shares, use Reed-Solomon Codes

$f(0) = y^*, share_j = f(j)$

Degree of f: about n/3

No explicit authentication, but inconsistent shares of <n/3 parties can be detected.

Good enough for Nash equilibrium

Our counterexample cannot claim "never worse" anymore.

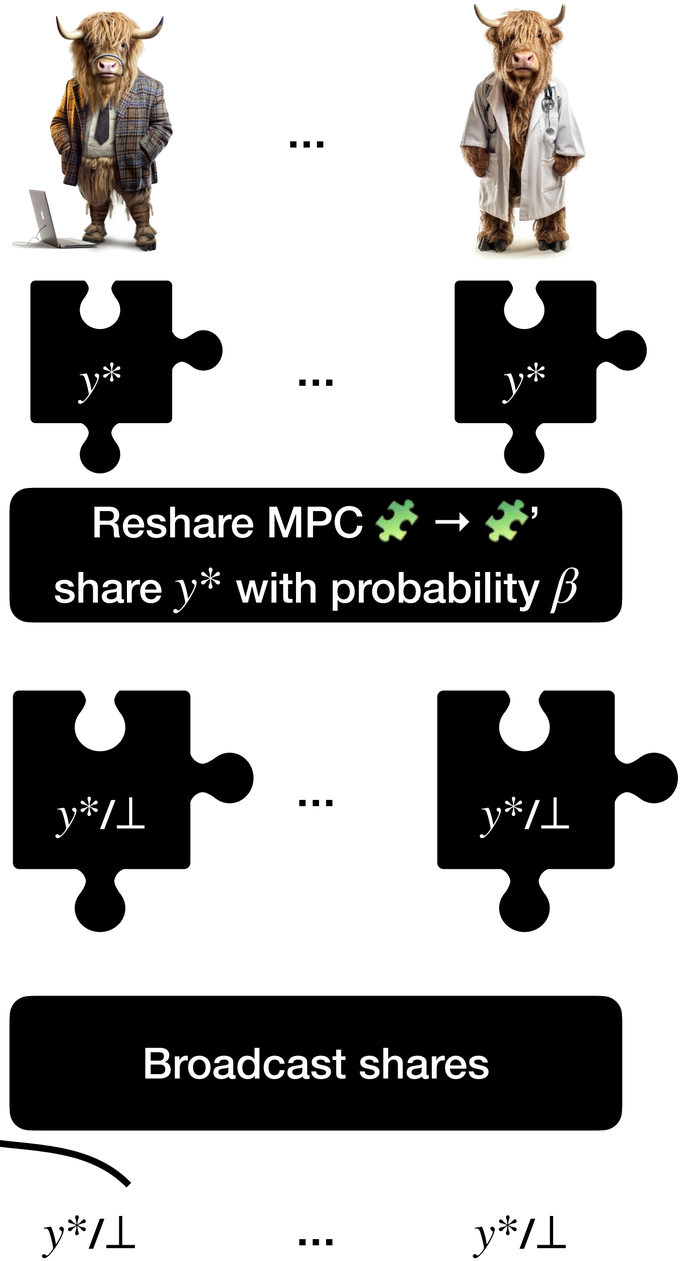

💡 ADGH06 is also weakly dominated

Assume secrets $y^* > 1000$ are never chosen.

Change:

If $y^* > 1000$, then return $y^* - 1000$

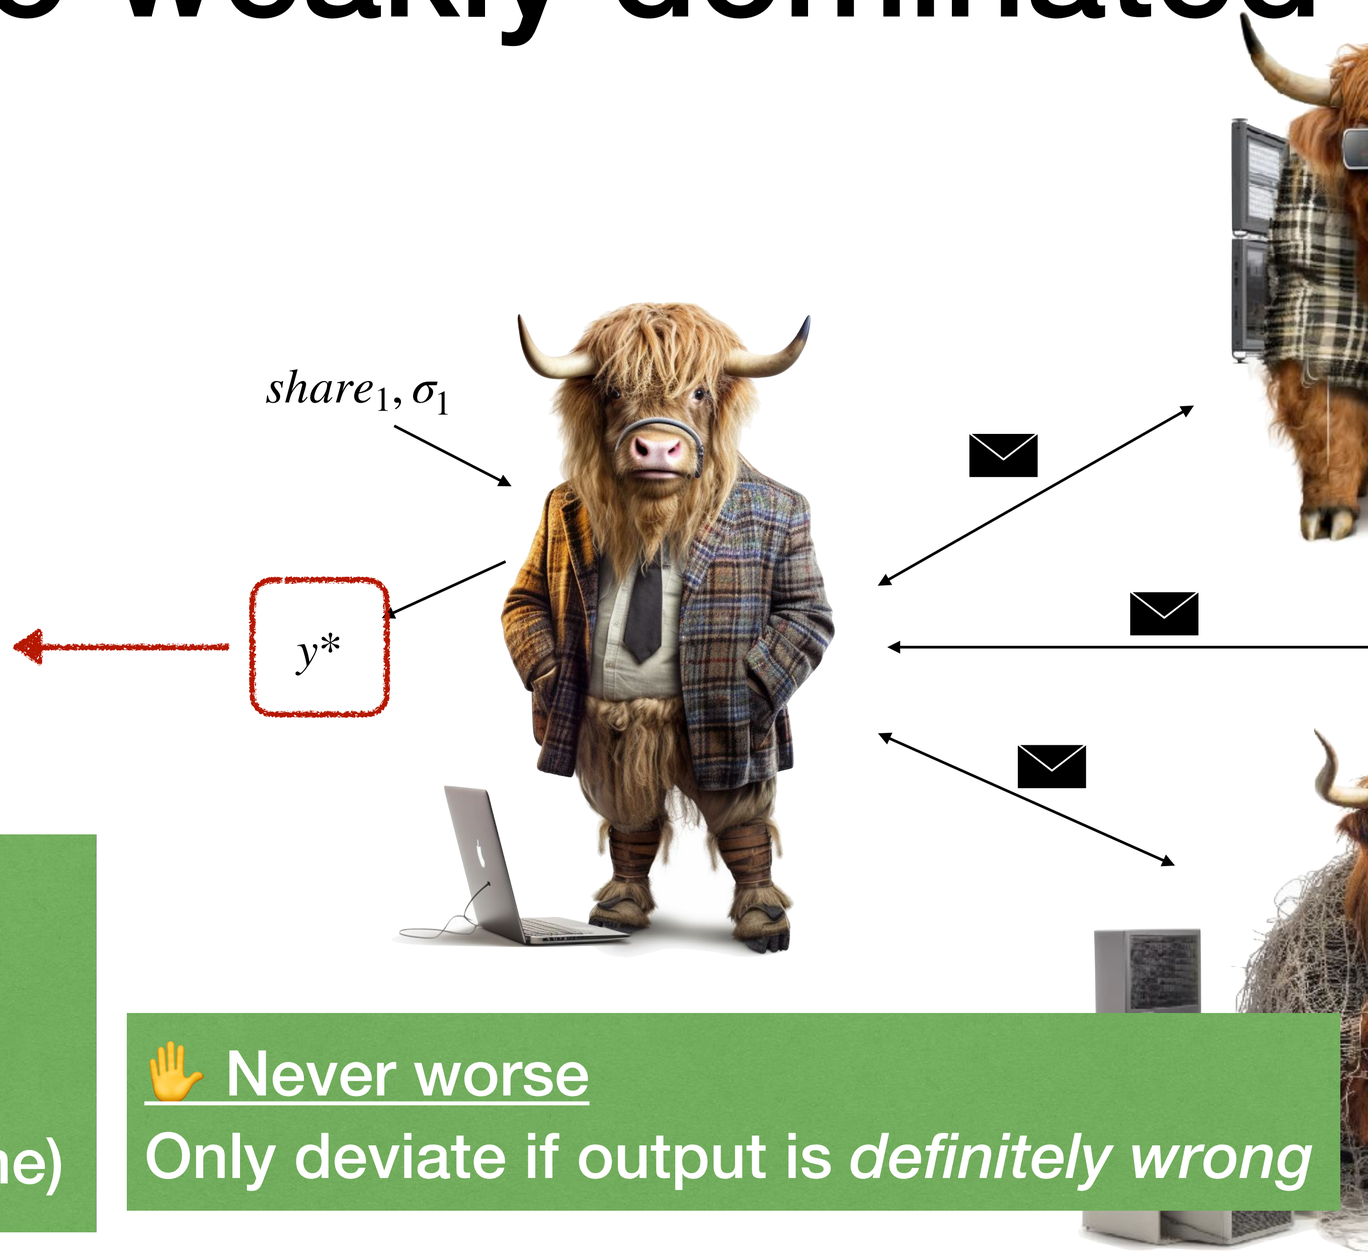

👍 Sometimes better

Against strategies that locally change their shares $[y^*]$ to $[y^* + 1000]$ (while keeping first user's share the same)

🖐 Never worse

Only deviate if output is *definitely wrong*

Secure coin flipping

$b_1 \leftarrow$ 🎲

$c \leftarrow \mathsf{Commit}(b_2; r)$

b_1

b_2, r

$b_2 \leftarrow$ 🎲

$b_1 \oplus b_2$

$b_1 \oplus b_2$

$c \stackrel{!}{=} \mathsf{Commit}(b_2; r)$

Can cheat: learns result first and can abort.
No *fairness*.

Rational coin flipping

$$u_1'(c, b_1, b_2, r) = \begin{cases} 10 & \text{if } b_1 \oplus b_2 = 0 \\ & \text{or } c \neq \text{Commit}(b_2; r) \\ -10 & \text{otherwise} \end{cases}$$

$u_2' = -u_1'$

$b_1 \leftarrow$ 🎲

$c \leftarrow \text{Commit}(b_2; r)$

$b_2 \leftarrow$ 🎲

b_1

b_2, r

$b_1 \oplus b_2$

$b_1 \oplus b_2$ if $c = \text{Commit}(b_2; r)$

0 if $c \neq \text{Commit}(b_2; r)$

If you abort, you “lose”
⇒ Abort allowed but irrational

Rational coin flipping

$$u_1'(c, b_1, b_2, r) = \begin{cases} 10 & \text{if } b_1 \oplus b_2 = 0 \\ & \text{or } c \neq \text{Commit}(b_2; r) \\ -10 & \text{otherwise} \end{cases}$$

$u_2' = -u_1'$

$b_1 \leftarrow$ 🎲

$c \leftarrow \text{Commit}(b_2; r)$

b_1

b_2, r

$b_2 \leftarrow$ 🎲

Incentive to send a non-random bit b_1?
Only negligibly
(e.g., try breaking hiding)

Incentive to abort? No.
Commit differently? Negligibly
(e.g., try breaking binding)

Impossibility result

Theorem 2. Let $\Pi = (\mathsf{Setup}_\Pi, \mathsf{Share}, \mathsf{Recon})$ be a secret sharing scheme (Definition 3) with perfect privacy (Definition 4). Consider a secret reconstruction game (Definition 12) for Π, with non-uniform strategies, non-trivial distribution of secrets (Definition 14), and reconstruction utilities preferring correctness (Definition 13). Let $(M_i, \omega'_1, \omega'_2, \dots)$ be a strategy for the secret reconstruction game, i.e. a non-uniform ppt ITM. If Π has (non-uniform) local $(n-1)$-verifiability (Definition 5), then there exists another strategy $(M_i^*, (\omega_1, \omega_2, \dots))$ which weakly dominates $(M_i, \omega'_1, \omega'_2, \dots)$ (Definition 10).

Theorem 4. Let Π be a secret sharing scheme (Definition 3) with perfect privacy (Definition 3) that is verifiable-or-fully-broken (Definition 16) for secret distributions $\mathcal{S}$. Consider the secret reconstruction game for secret sharing scheme Π with non-uniform strategies, non-trivial distribution of secrets $\mathcal{S}$ (Definition 14), and reconstruction utilities preferring correctness and exclusivity (Definition 13). Let $t \geq n/2$. Then there exists no mechanism with the following properties:

- If everyone follows the mechanism, the correct secret is reconstructed with probability 1.
- The mechanism is a t-resilient Nash equilibrium (Definition 9).
- There is no coalition $C \subseteq [n], |C| = t$ such that M_C is weakly dominated (Definition 15).