

Florian Tramer
Stanford

Phil Daian, Ari Juels
Cornell Tech, Jacobs, IC3

Lorenz Breidenbach
Cornell Tech, ETH, IC3

Enter the Hydra: Toward Principled Bug Bounties and Exploit-Resistant Smart Contracts

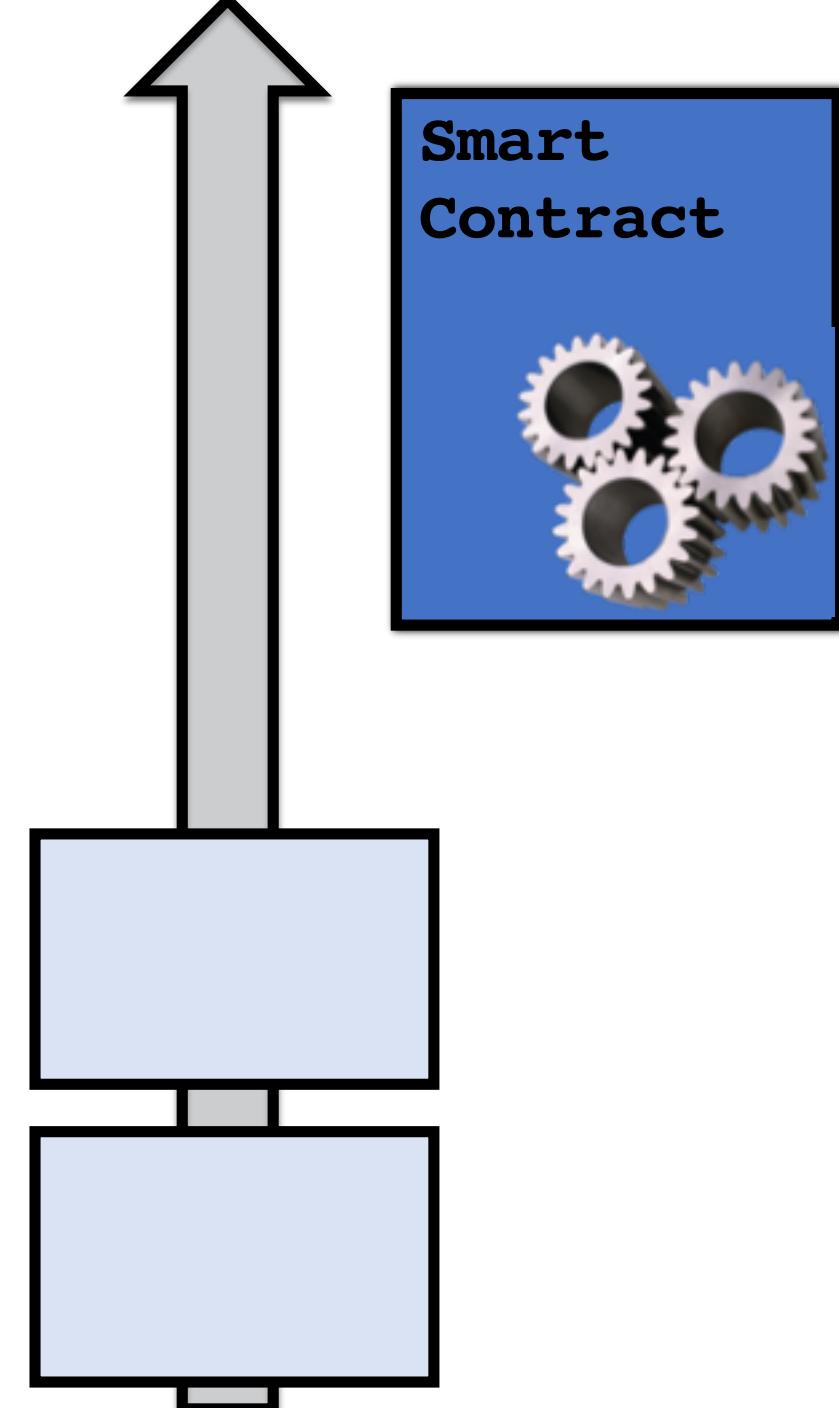
Summer School on Real-World Crypto and Privacy
Sibenik, Croatia, 14 June 2018

What's a Smart Contract?

- Type interpreted by operations
- Only stack & alt-stack
- No return stack (no calls)
- No heap
- Deterministic - No side effects or I/O

Smart contracts

- Small programs that run on *blockchains*
- Given trust in underlying blockchain, smart contracts are
 - Transparent
 - Irreversible
 - Tamper-resistant
- ...plus they can act upon **crypto tokens = \$money**



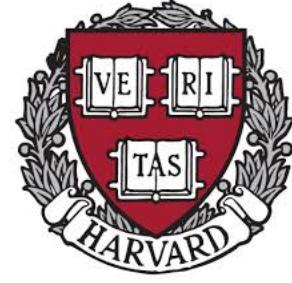
Lots of recent interest in ETH...



\$22 billion



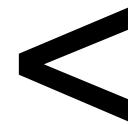
\$27 billion



\$35 billion

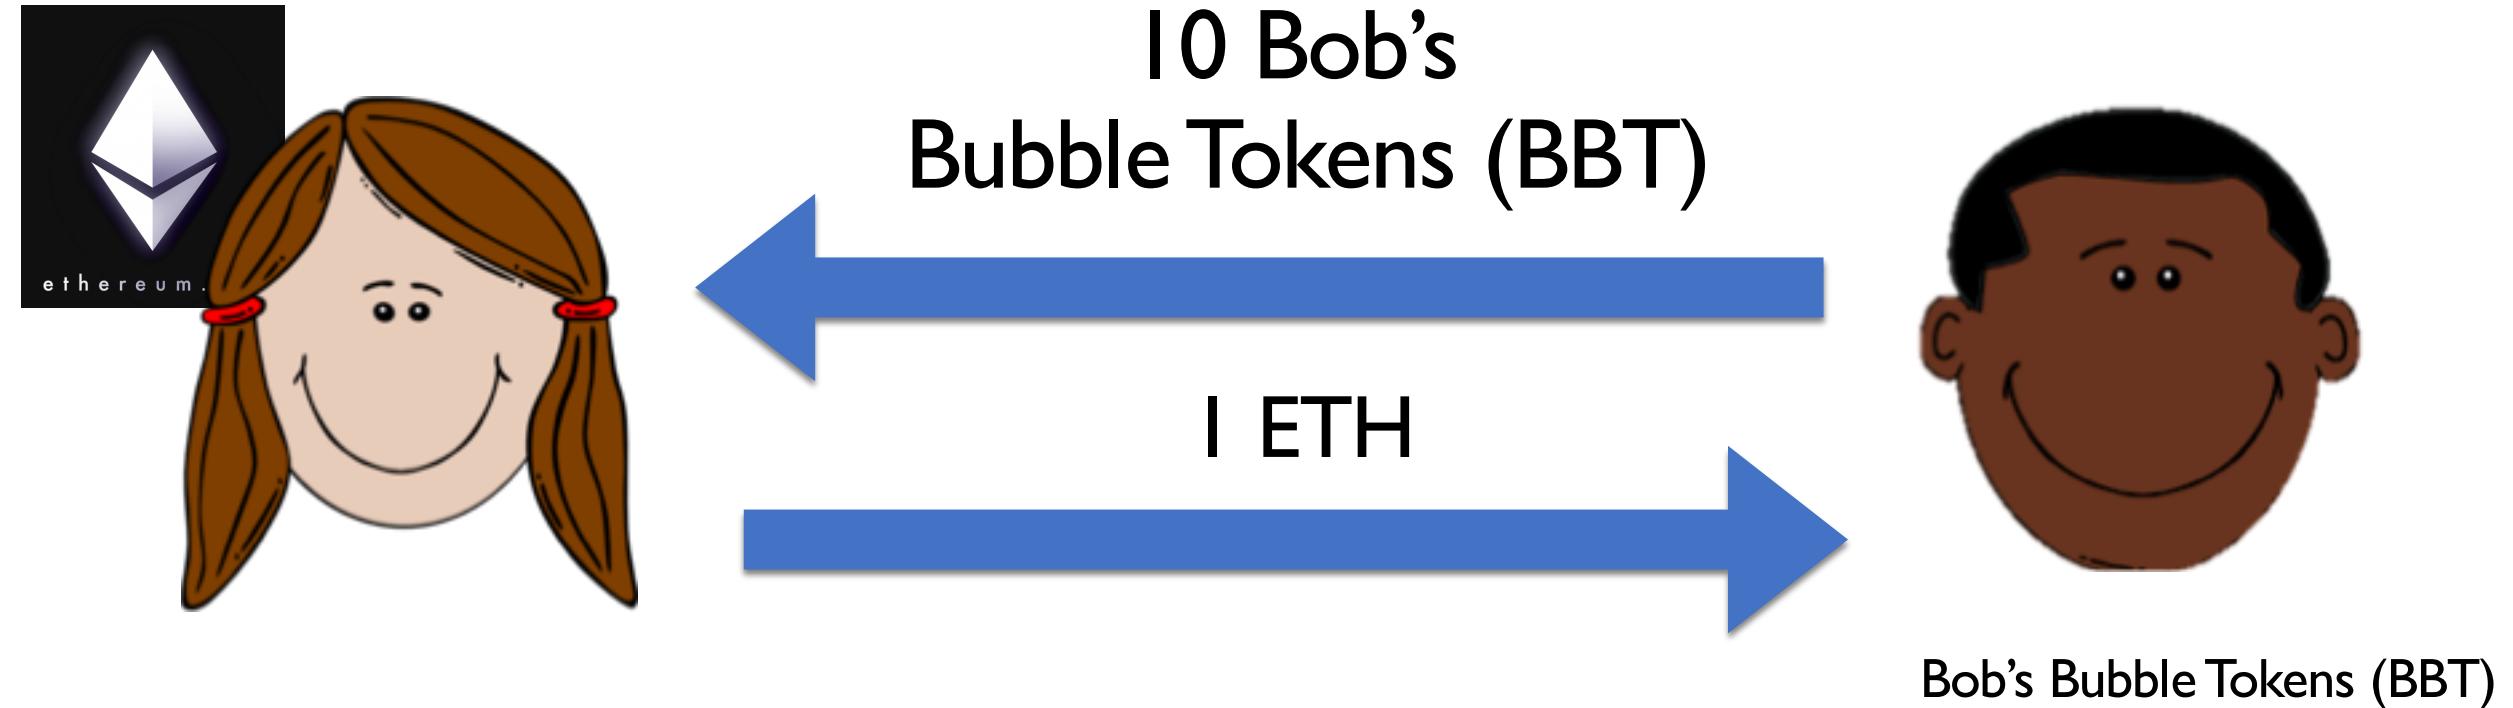


\$7 billion



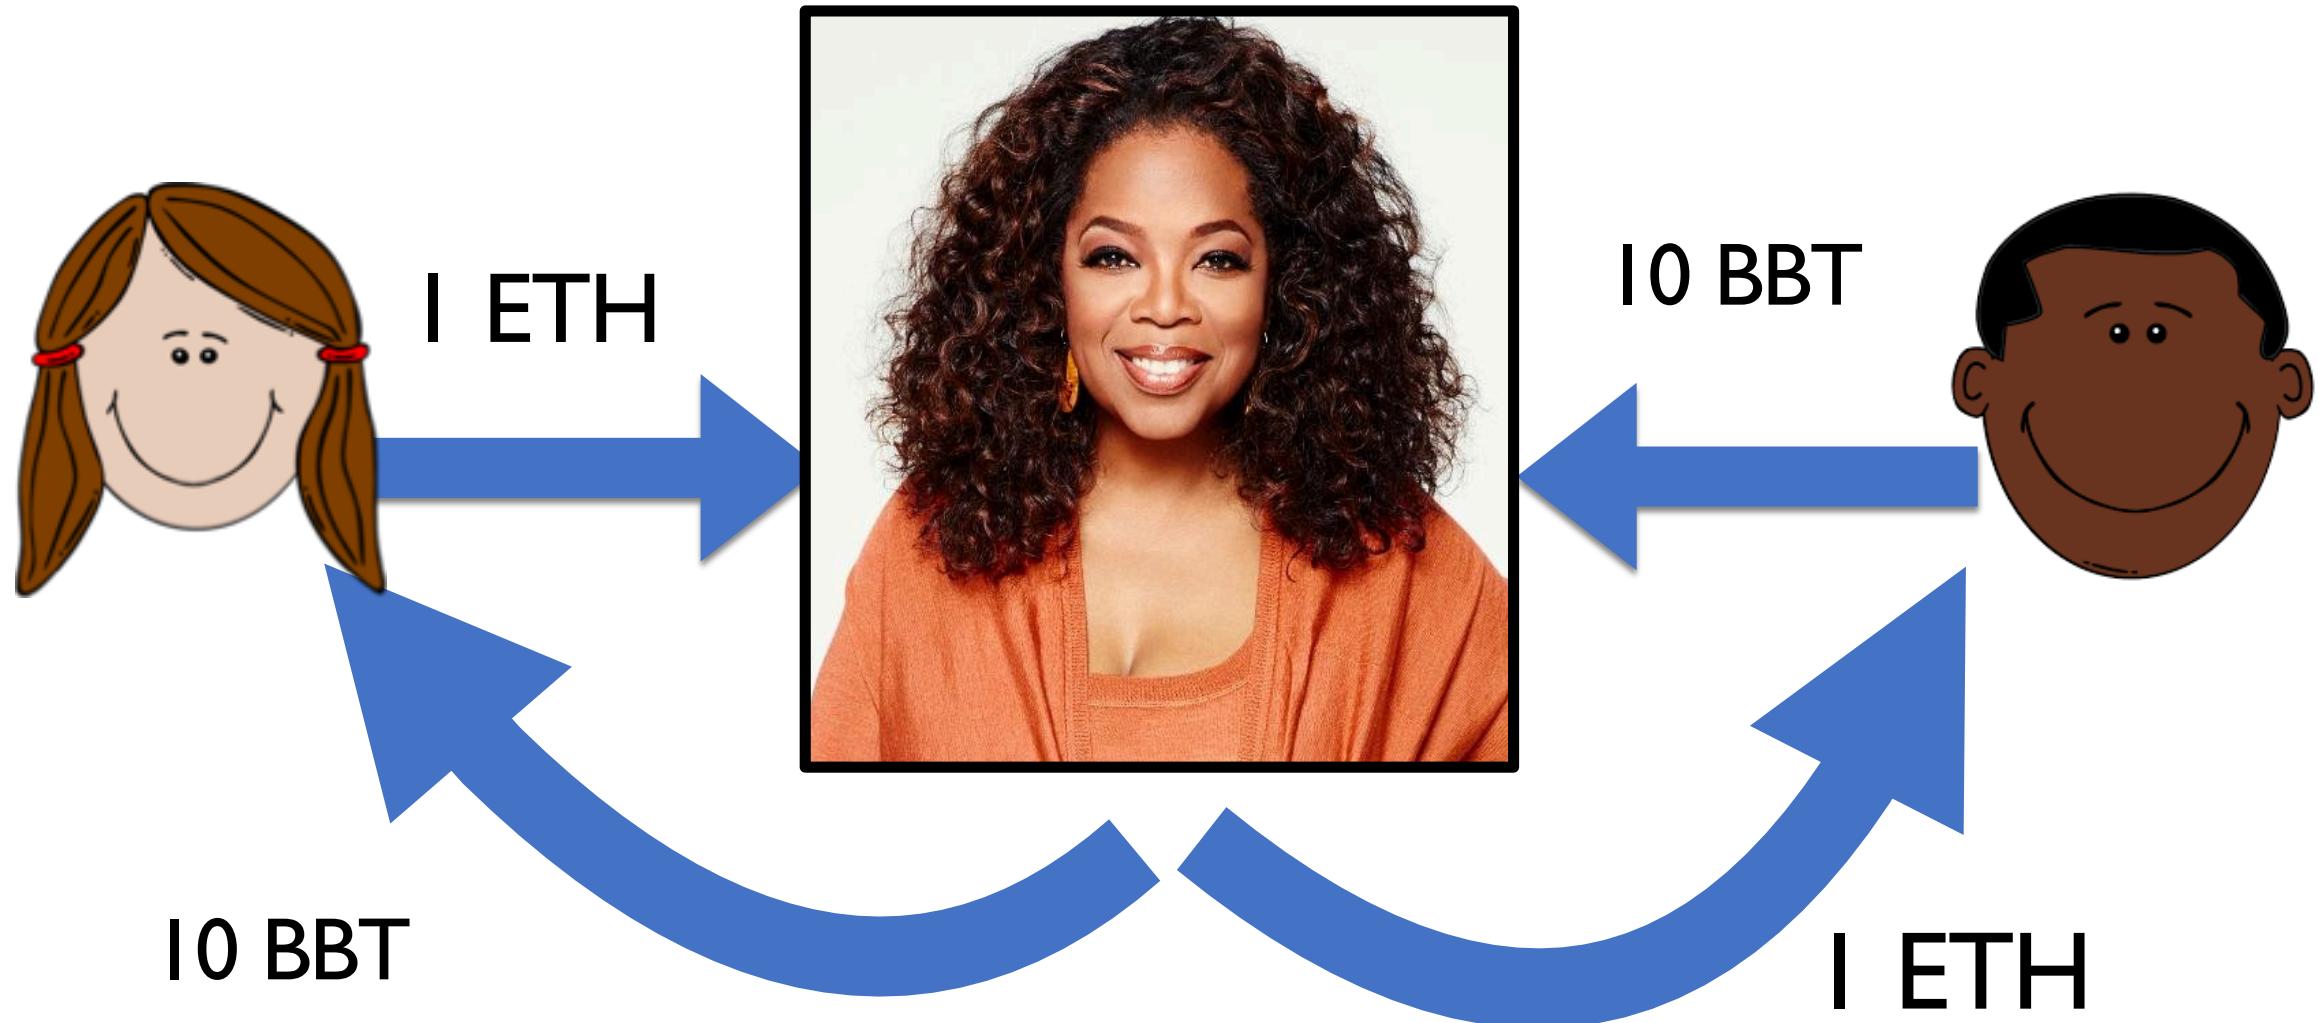
> \$48 billion

Why? Suppose Alice and Bob want to trade..

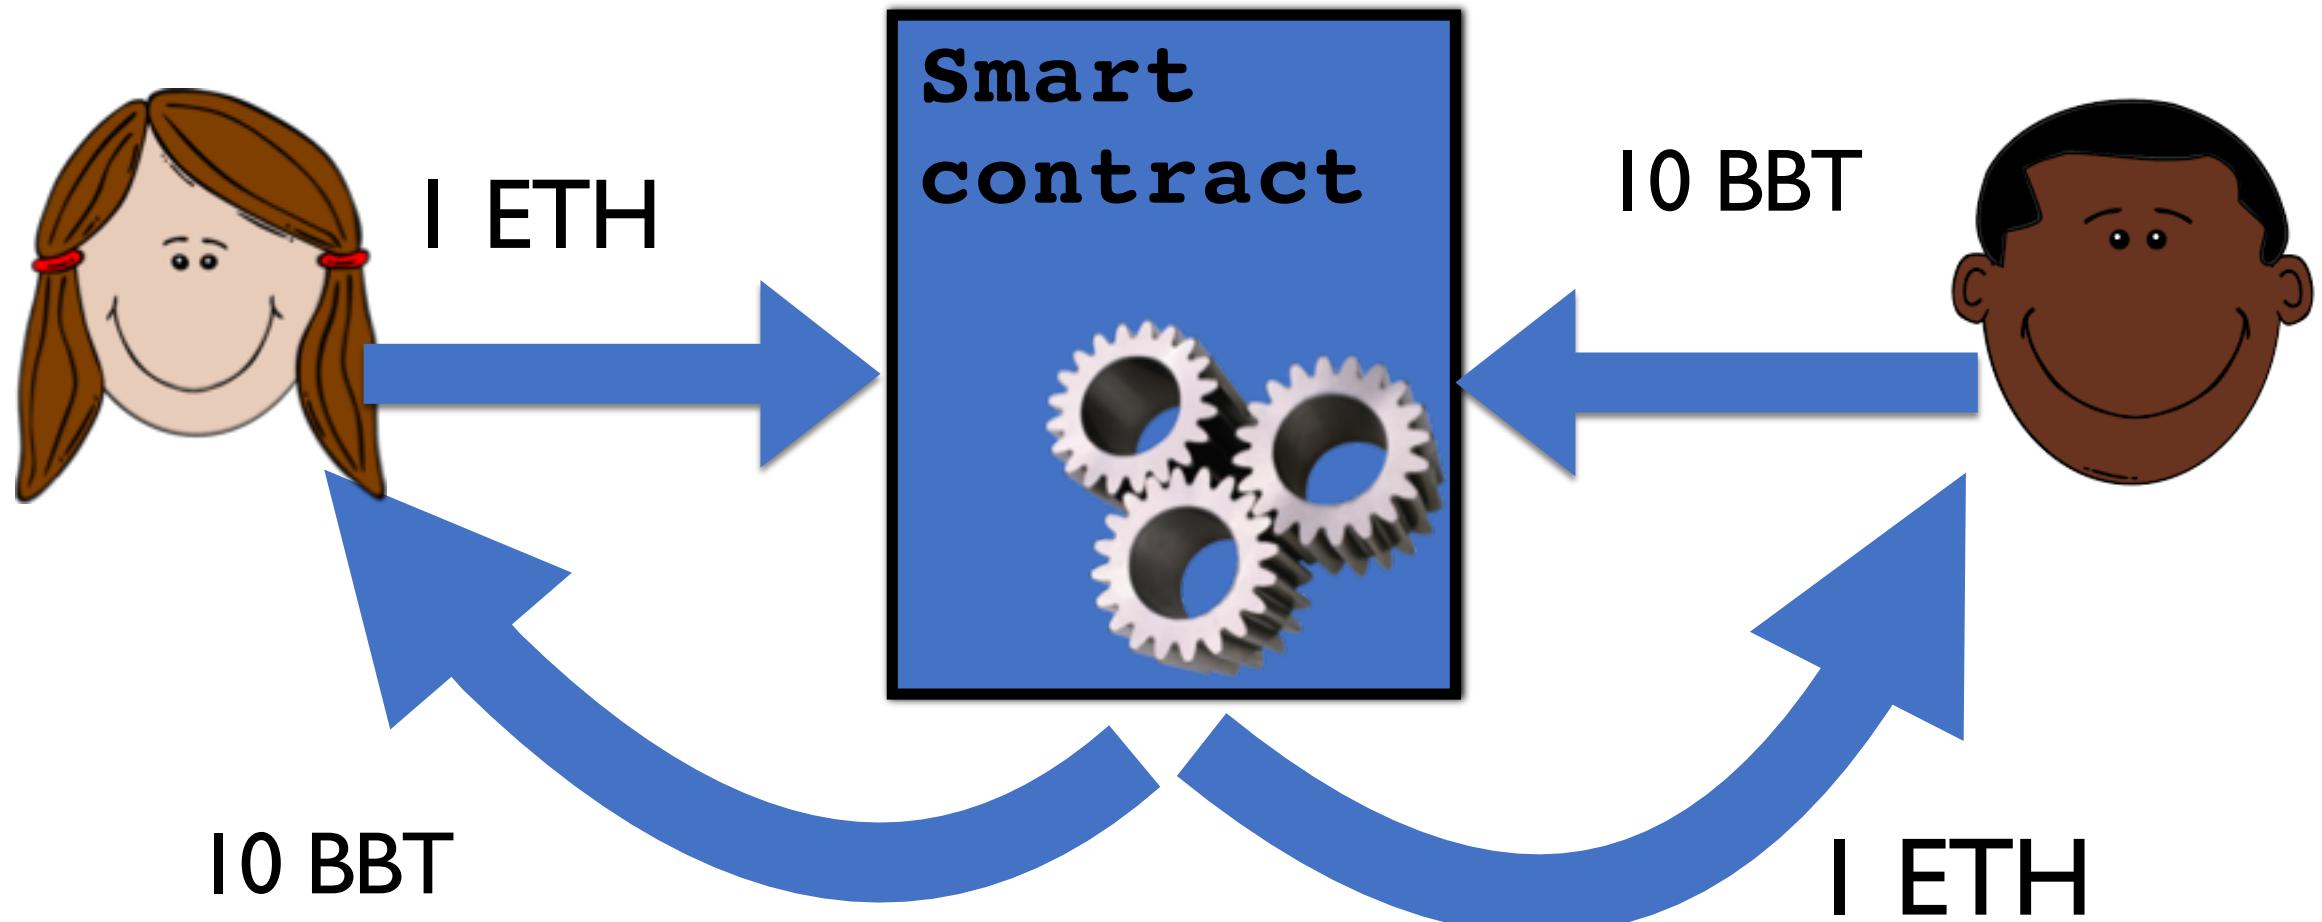


Problem of *Fair Exchange*!

Trusted third-party (with public state)



**Smart contract ≈
Trusted third-party (with public state)**





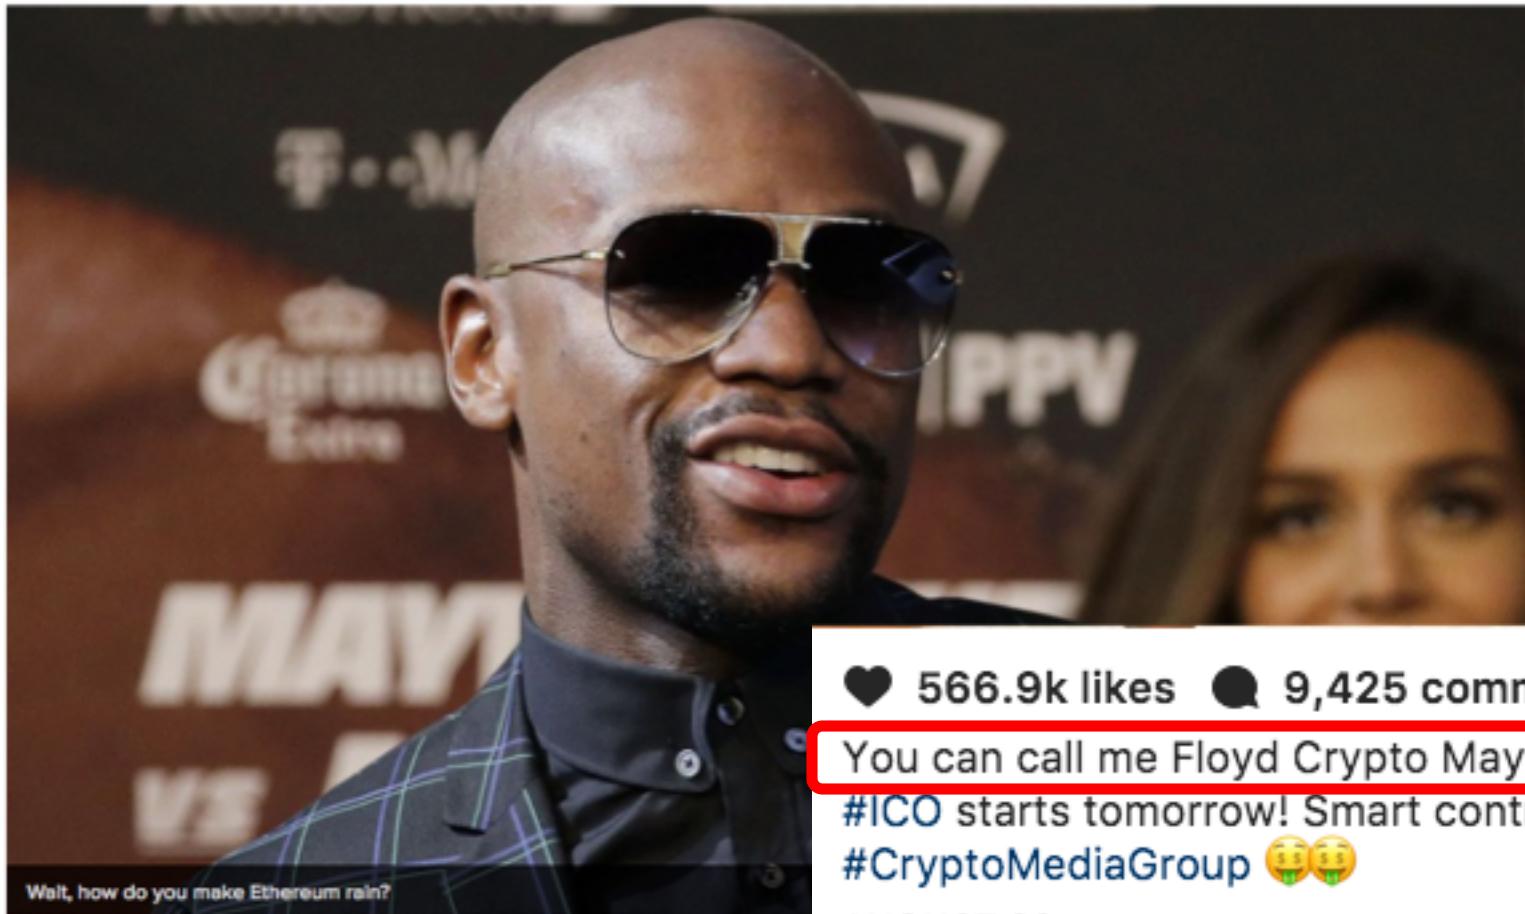
No, not
Floyd Mayweather...

Floyd 'Crypto' Mayweather promotes an ICO, again



Mashable

AUG 24, 2017



566.9k likes 9,425 comments

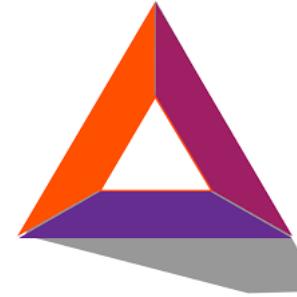
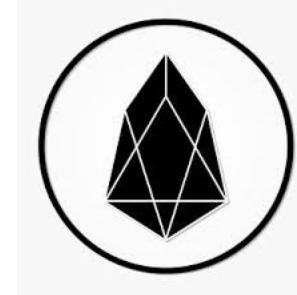
You can call me Floyd Crypto Mayweather from now on..Hubii.Network
#ICO starts tomorrow! Smart contracts for sports?! #HubiiNetwork
#CryptoMediaGroup 💰💰

AUGUST 23



Crypto Tokens

- Application-specific cryptocurrency
- Mainly ERC20 tokens
 - Managed in Ethereum smart contracts
- \$38+ billion token market cap

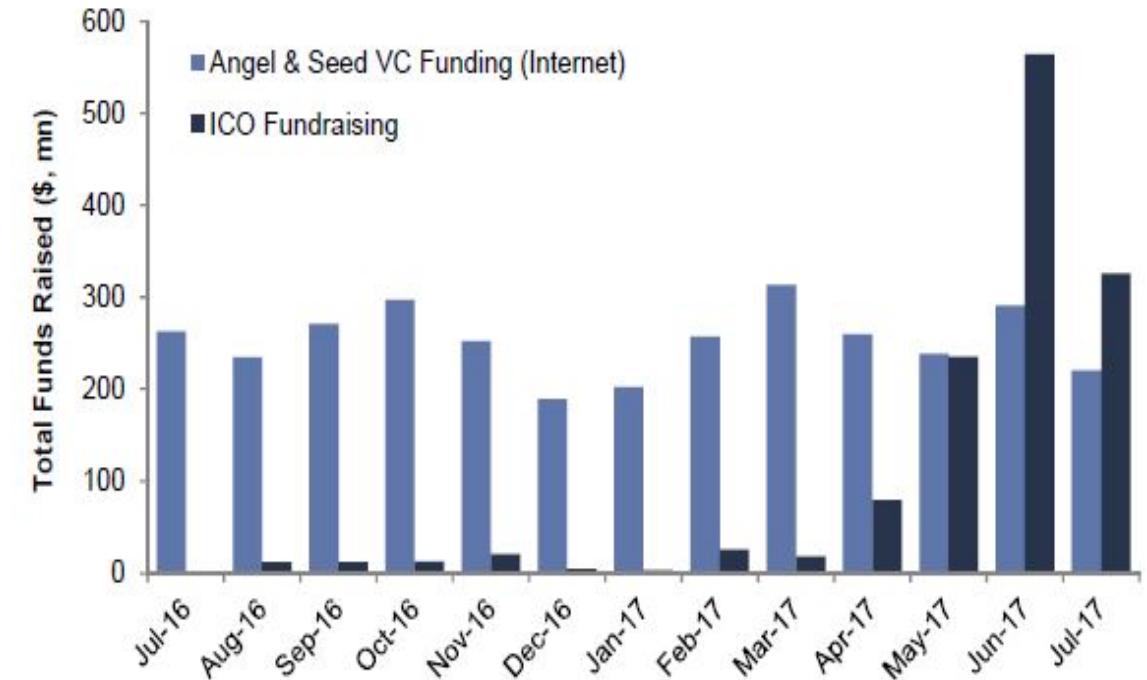


Crypto Tokens

- Sold in Initial Coin Offerings (ICOs)
 - a.k.a. Token Launch, Token Generation Events (TGEs), etc.
 - Like unregulated VC
 - Token like a share (kind of...)
- Since mid-2017, ICO funding outstripping early-stage Internet VC (!)

Exhibit 8: The pace of ICO fundraising has now surpassed Angel & Seed stage Internet VC funding globally

Total Funds Raised by month (\$, millions)

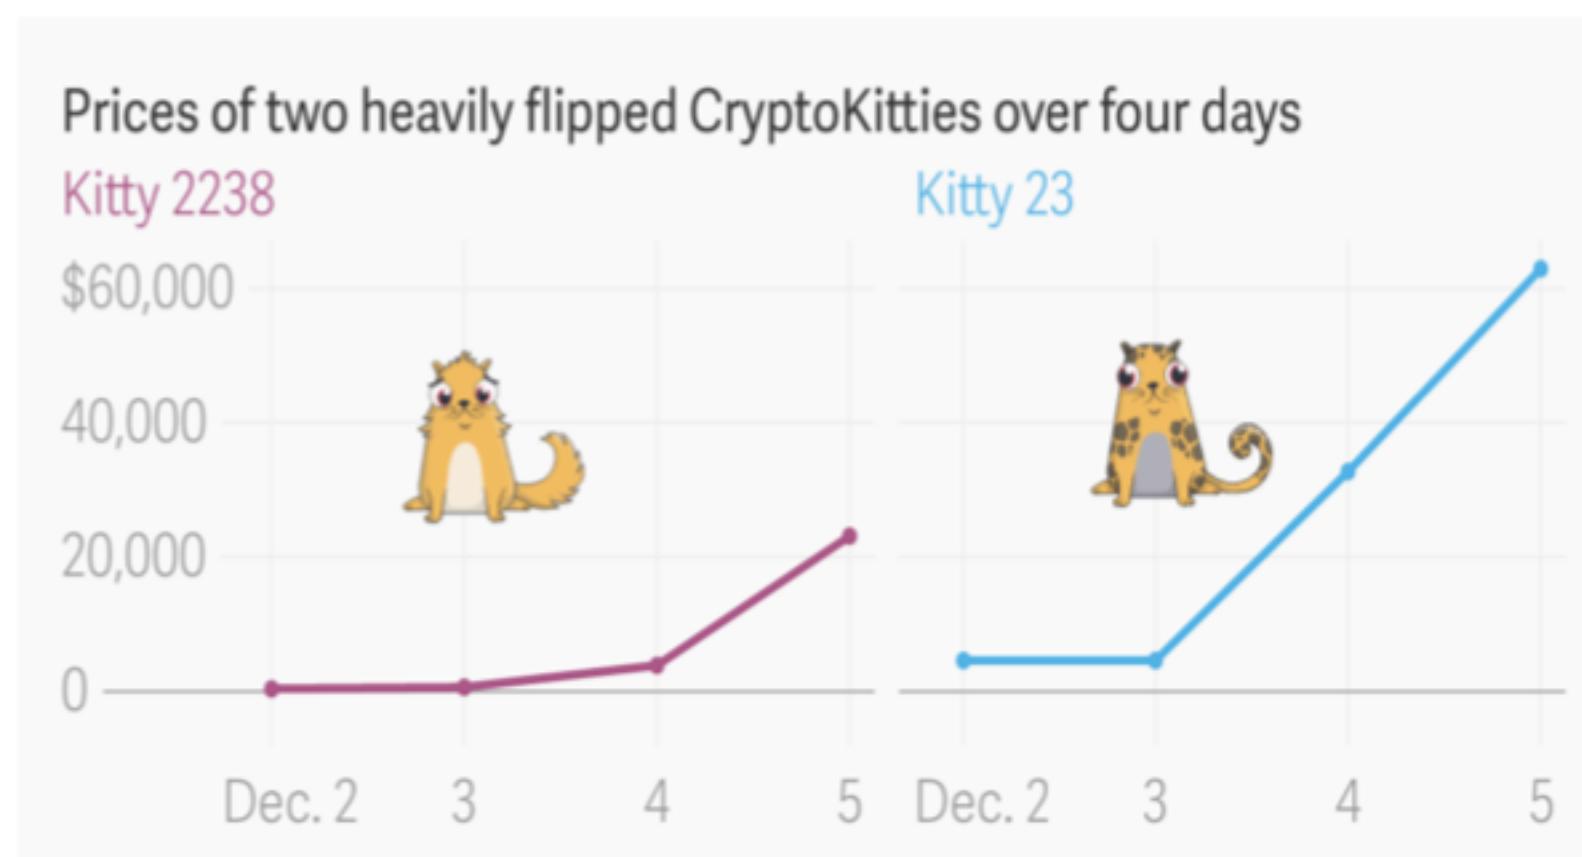
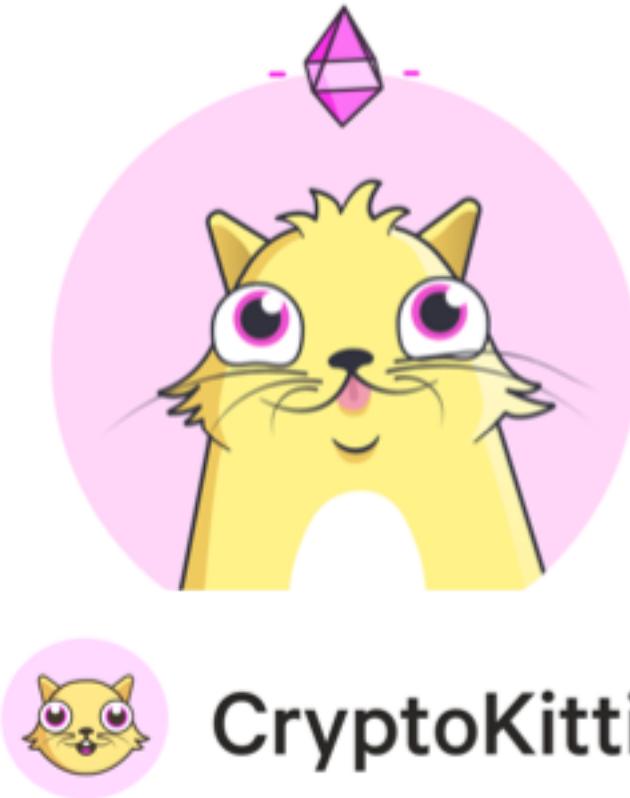


Note: ICO fundraising as of July 18th, 2017, per Coin Schedule. Angel & Seed VC funding data as of July 31st, 2017 and does not include "crowdfunding" rounds.

Source: CoinSchedule, CB Insights, Goldman Sachs Global Investment Research.

Crypto Tokens: ERC721

- “Non-fungible tokens”: Represent unique objects



SMART CONTRACT CHALLENGES



- 1. Correctness:** Contracts often have fatal bugs!
- 2. Confidentiality:** No private data.
- 3. Authenticated data:** No good, trustworthy access to real-world data!

Side effects of the token mania

- Token smart contracts are compact
- Lots of money per contract
- Astonishing value per line of code
- Which makes for juicy targets...

Token	Lines of Code	Value per line
OmiseGo (OMG)	396	~\$2.4M
Tether (USDT)	423	~\$5.9M
EOS (EOS)	584	~\$15.8M

Sources: coinmarketcap.com, 14 June 2018., and published contract source code

Some (in)famous smart contracts

- The DAO (June 2016)
 - Reentrancy bug \Rightarrow \$50+ million stolen
- Parity multisig hack (July 2017)
 - Parity 1.5 client's multisig wallet contract
 - Bad use of **delegatecall** \Rightarrow \$30 million stolen
 - ...from 3 ICO wallets (Edgeless Casino, Swarm City, and æternity)
- Parity multisig hack—Redux! (Nov. 2017)
 - Bad use of **delegatecall** \Rightarrow >\$150 million frozen
 - ...much from ICO wallets (Polkadot, \$98 million)

Why not try to address correctness with...

- Formal verification
 - Absolutely!
 - But limited scaling
 - What if there's a bug in the formal spec? (Turtles!)
- Static and dynamic verification
 - Absolutely!
 - But limited scope



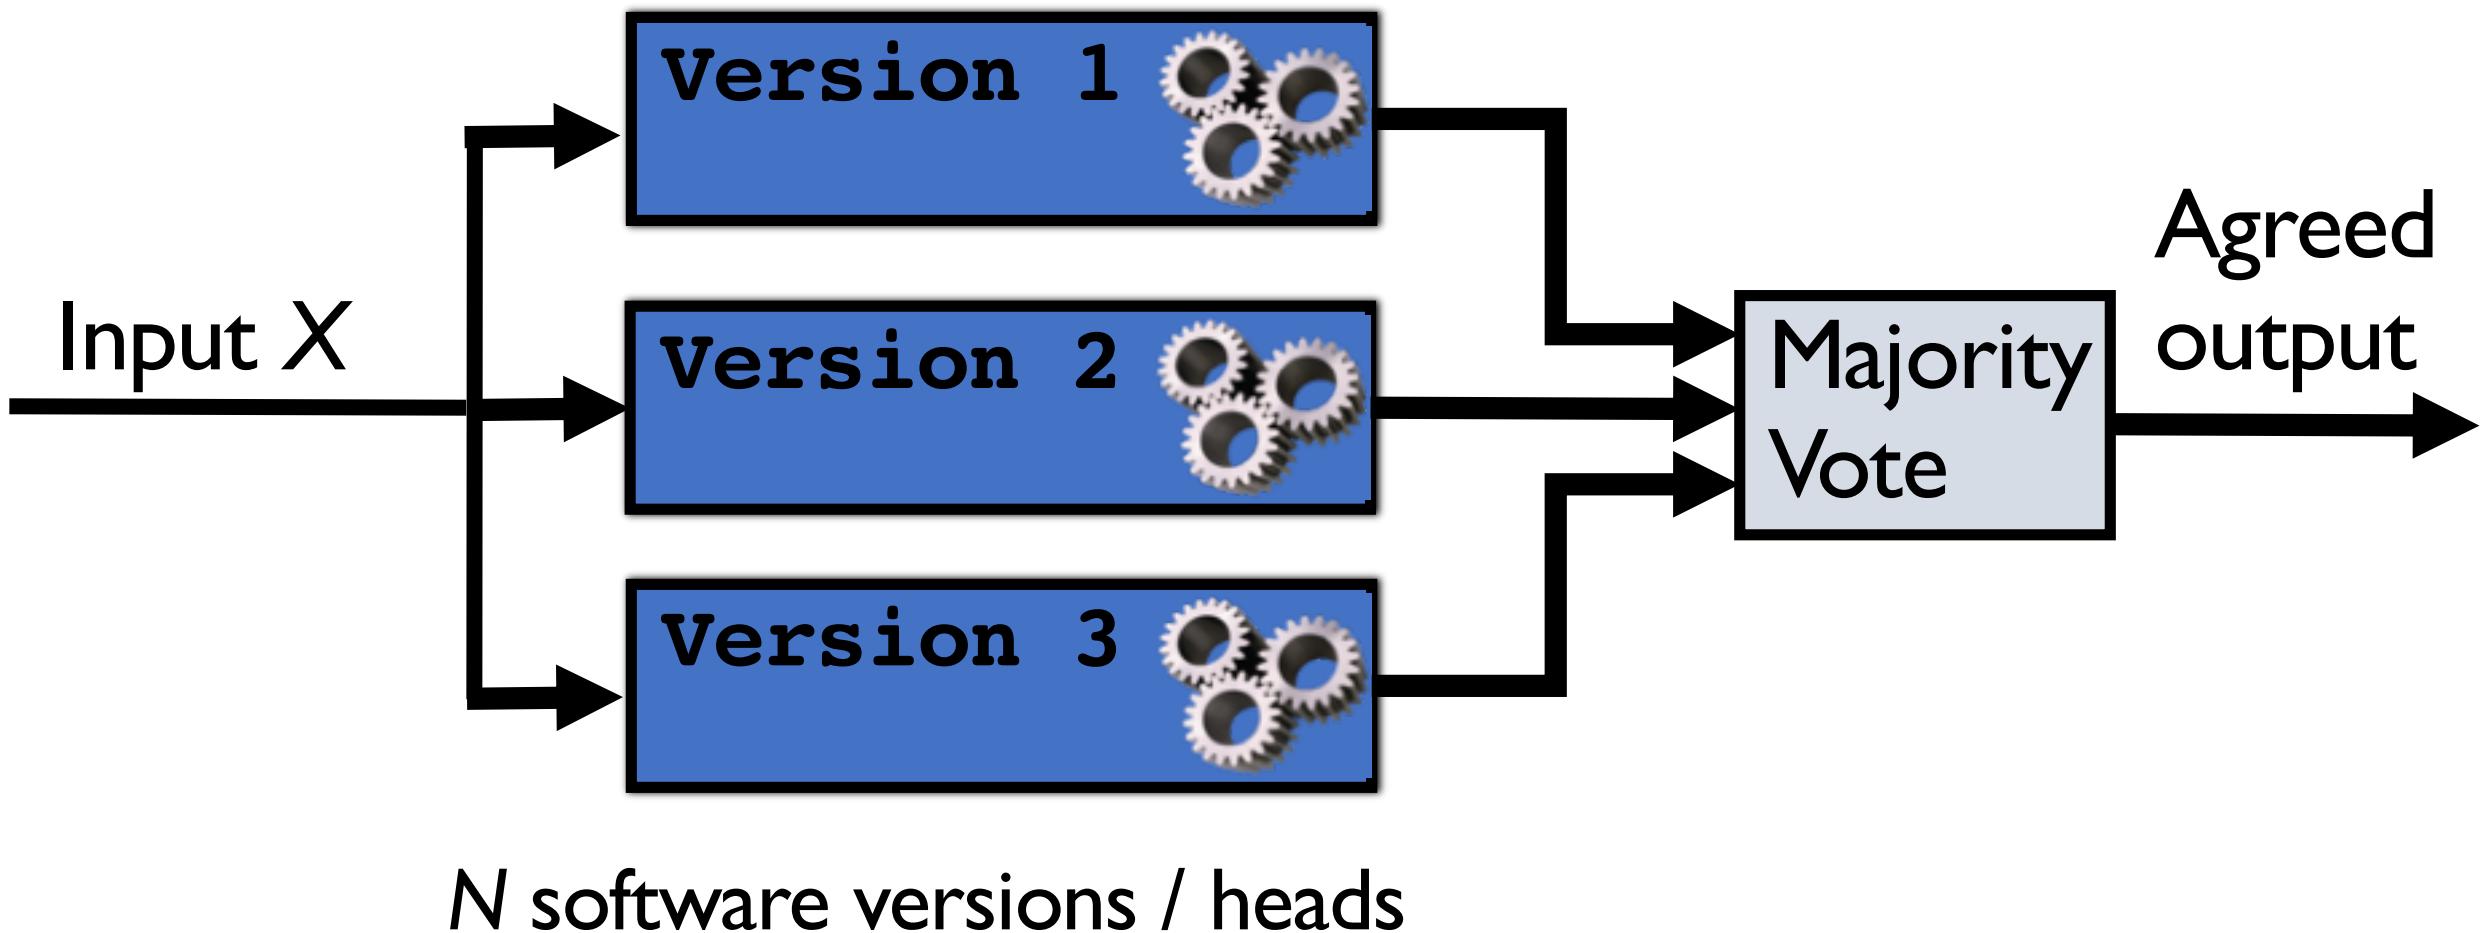
N-Version programming

(Chen & Avizienis '78, Knight-Leveson '86)

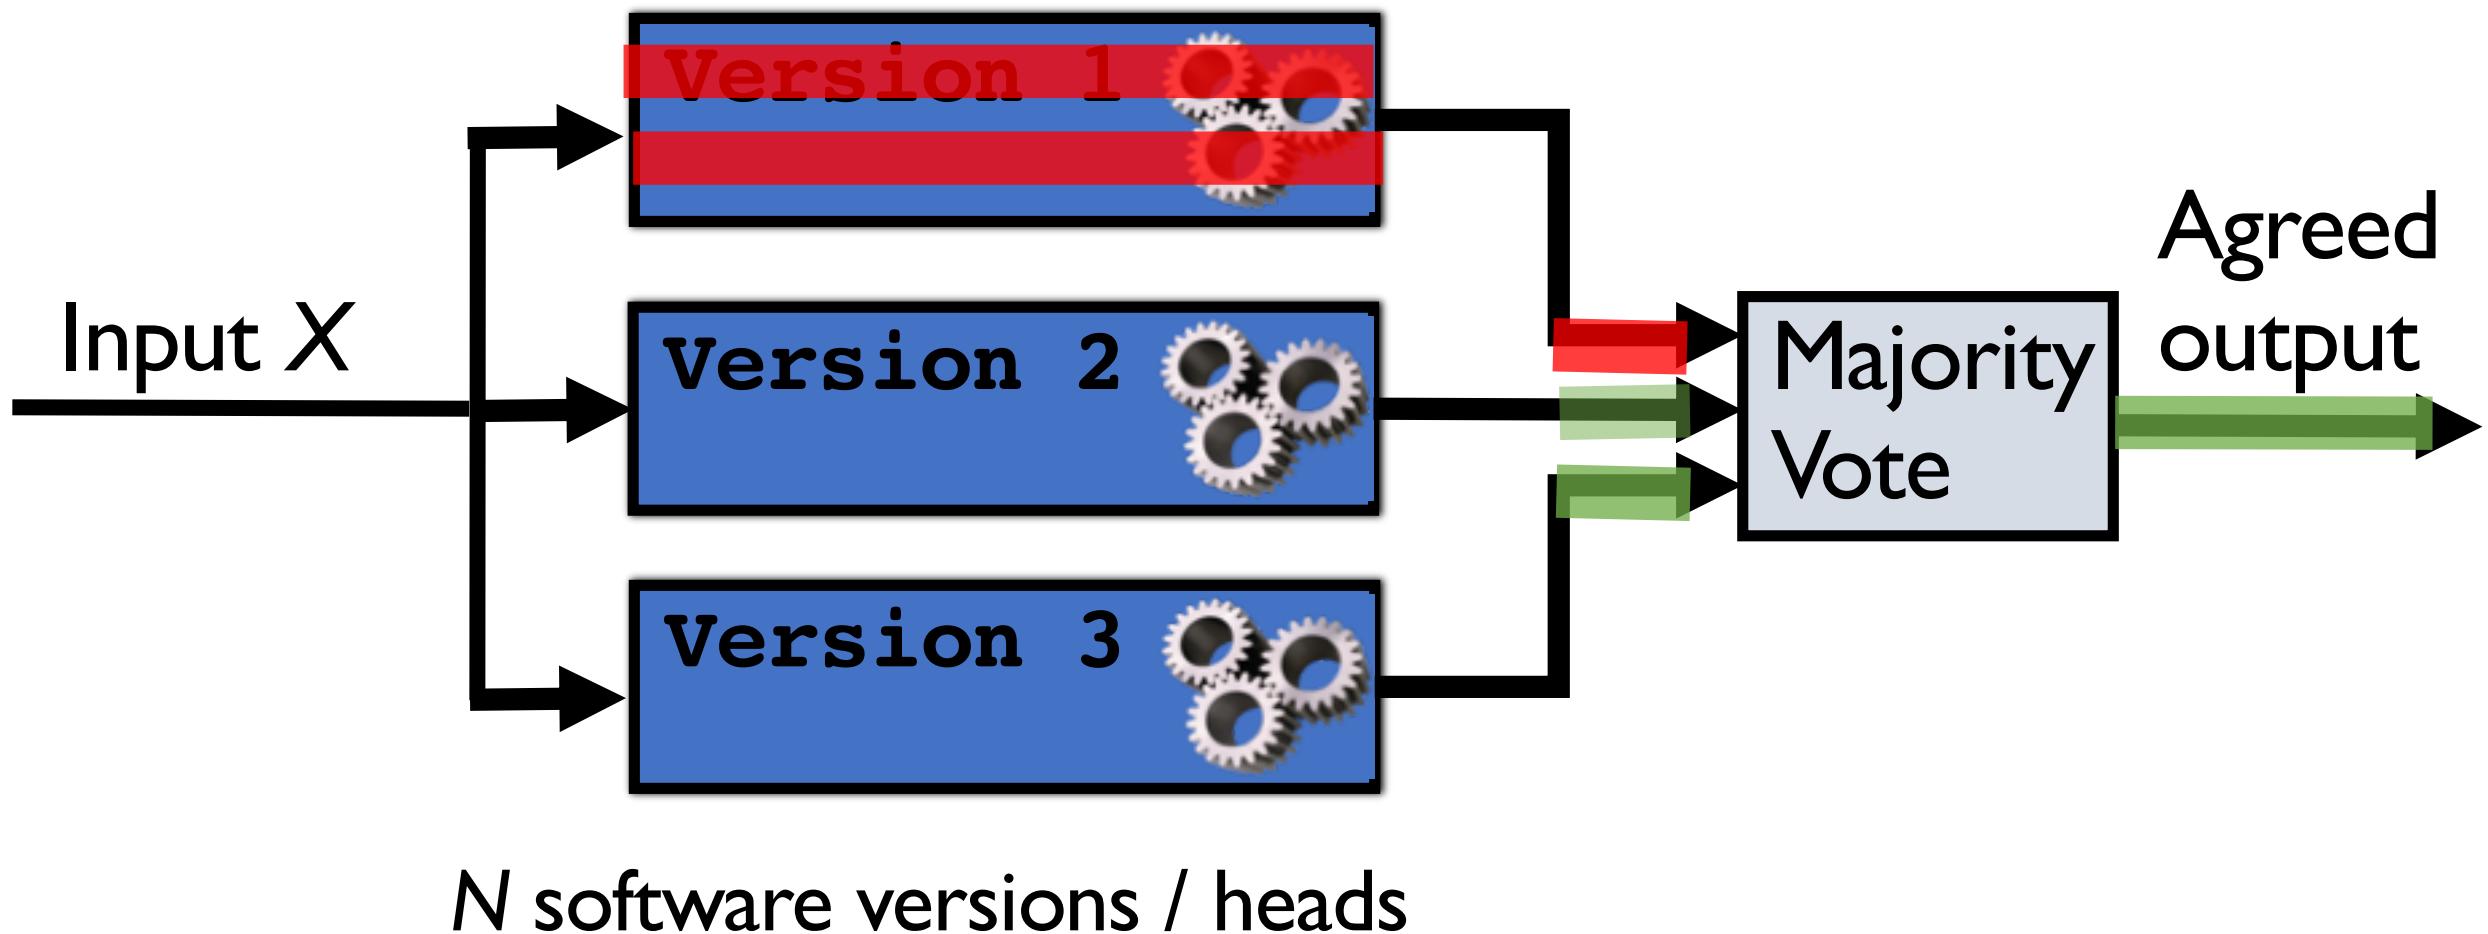


N-Version programming

(Chen & Avizienis '78, Knight-Leveson '86)

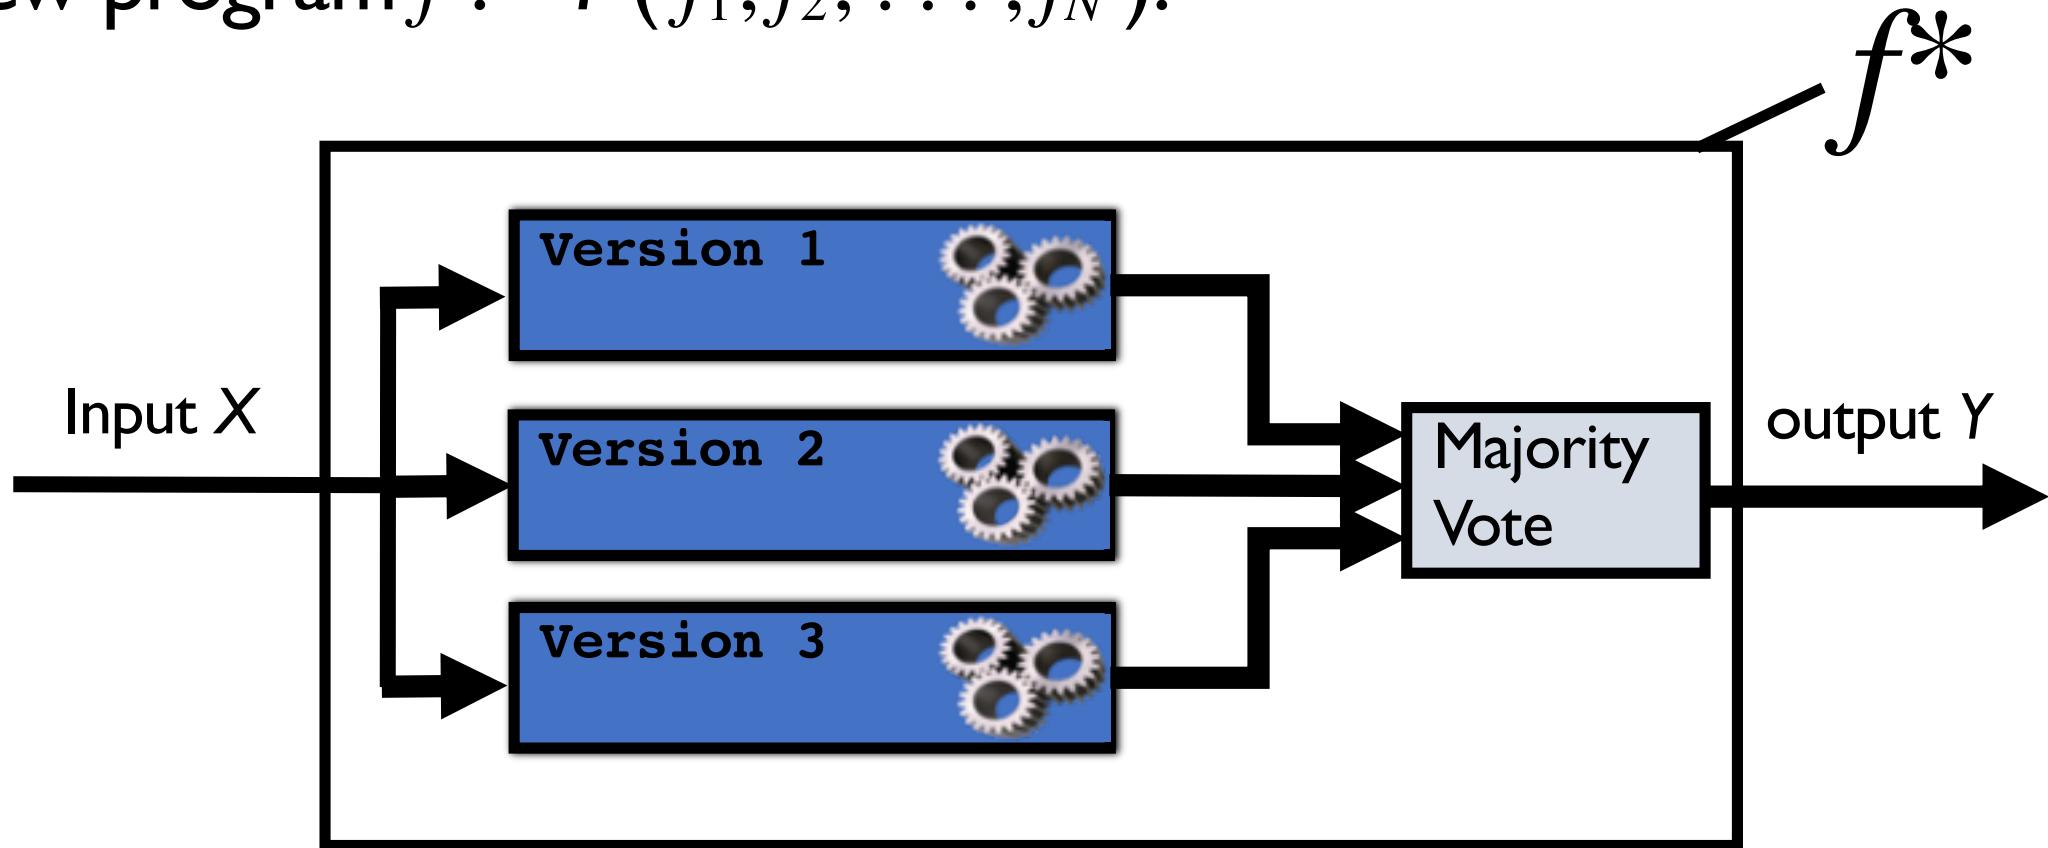


If something goes wrong...



What is N-version programming doing?

A *program transformation* T takes $N \geq 1$ programs and creates new program $f^* := T(f_1, f_2, \dots, f_N)$.



Some more definitions

- Let \mathcal{I} be an *ideal* program specification
 - Conceptual! Doesn't actually exist...
- Let f be an implemented program
- An *exploit* is an input X such that $\mathcal{I}(X) \neq f(X)$
- Intuition: Any deviation from *intended behavior* is a potentially serious bug
- *Exploit set* $E(f, \mathcal{I})$: set of exploits X for f and \mathcal{I}

Mind the gap

- Let \mathcal{D} be a distribution over inputs X
- Definition of **exploit gap**:

$$\text{gap} := \frac{\Pr_{X \in \mathcal{D}} \left[X \in \bigcup_{i=1}^N E(f_i, \mathcal{I}) \right]}{\Pr_{X \in \mathcal{D}} [X \in E(f^*, \mathcal{I})]}$$

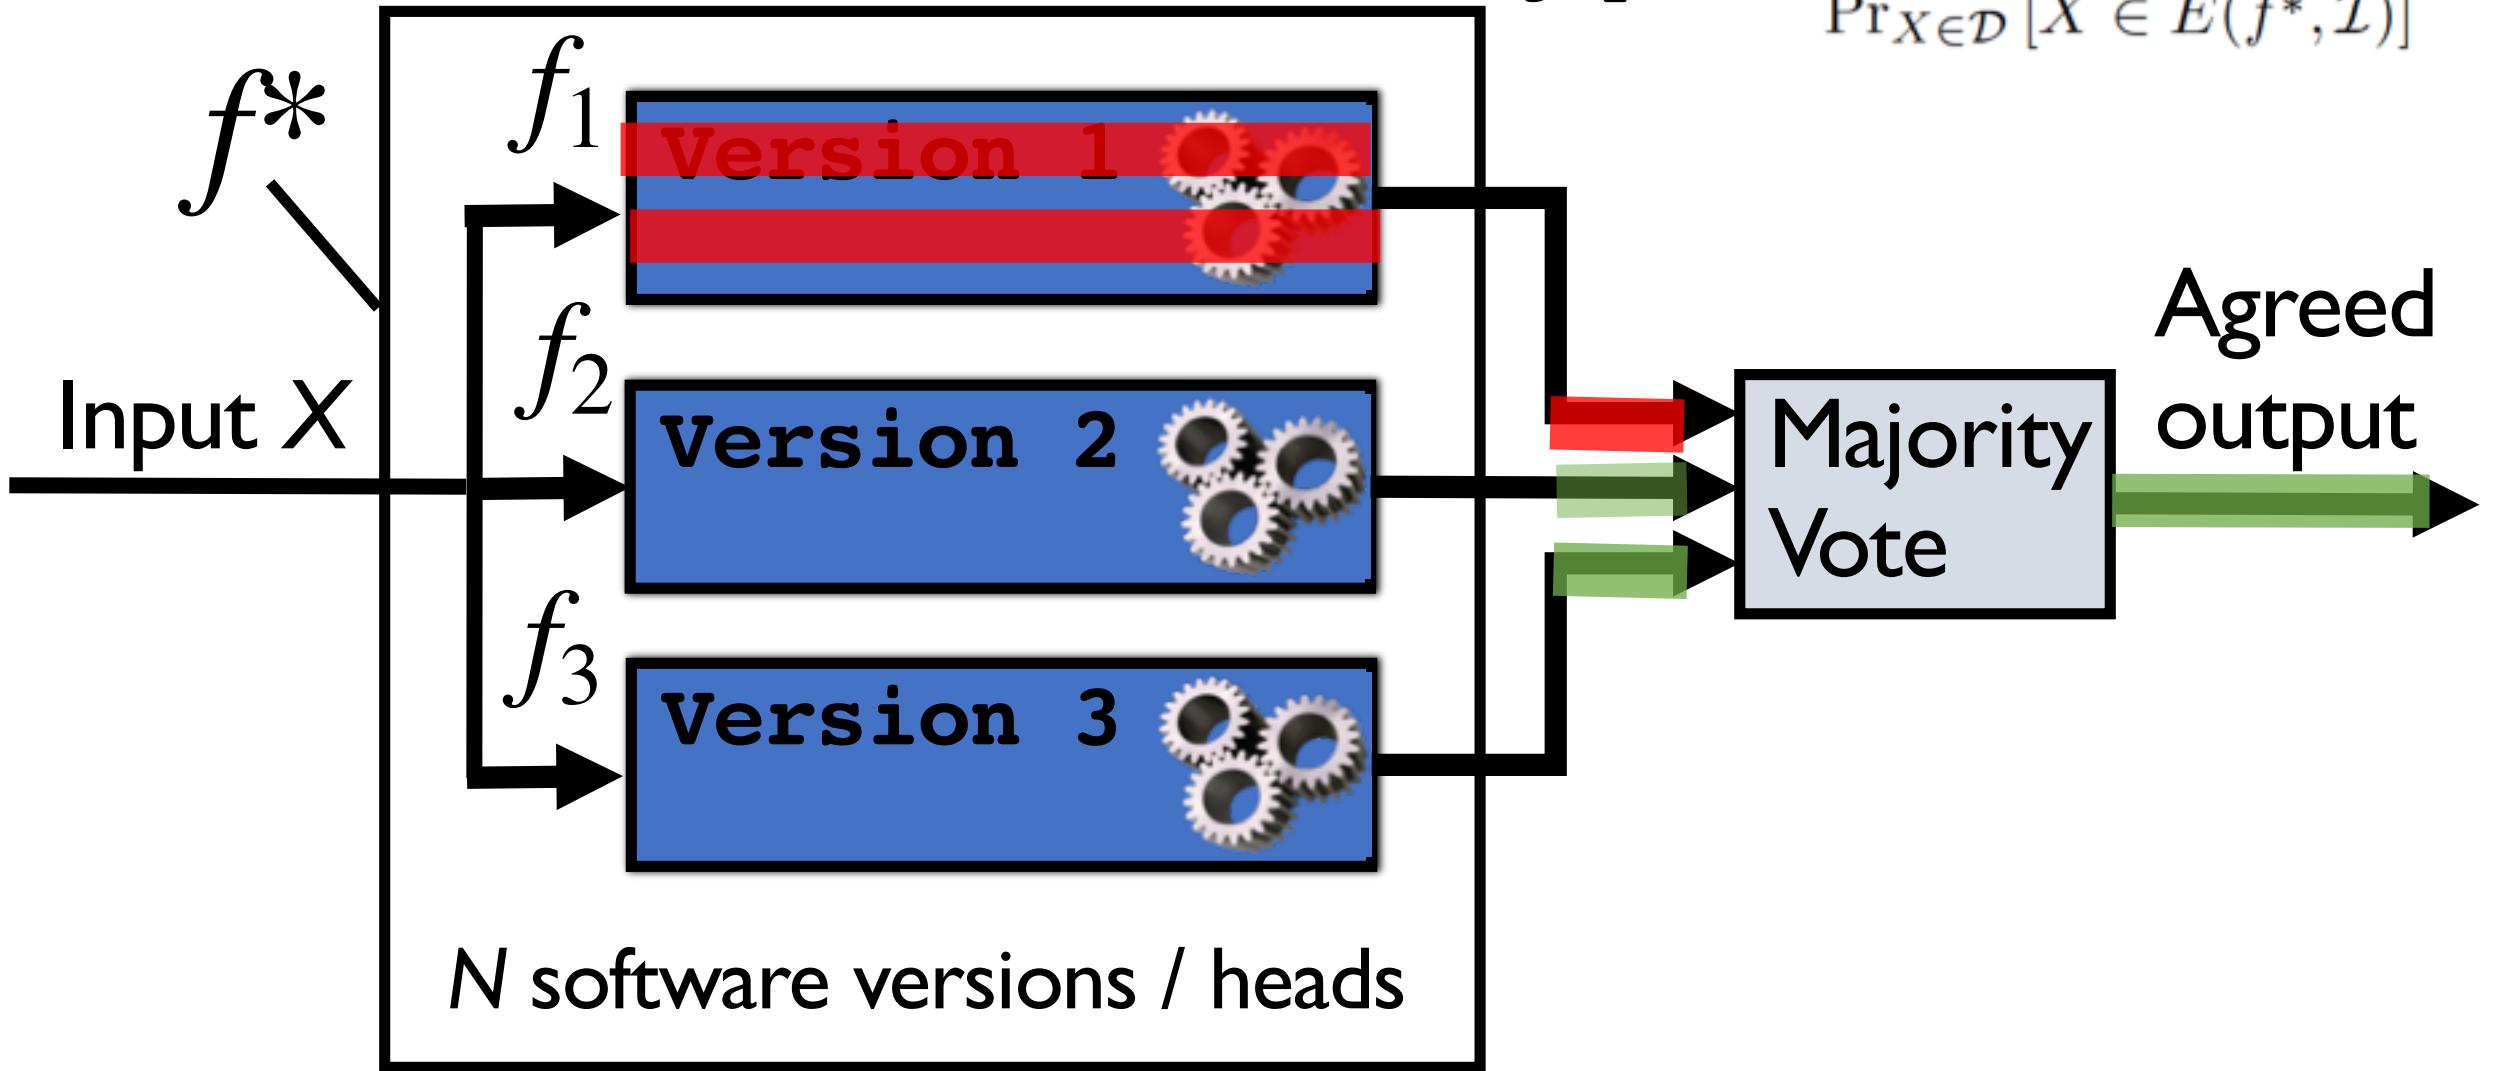
Exploits against
 $f_1, f_2, f_3 \dots$

Exploits against f^*

- *Affirmative* gap (> 1) means T reduces exploits
- Bigger gap \Rightarrow fewer relative bugs in f^*
- gap captures dependencies among heads

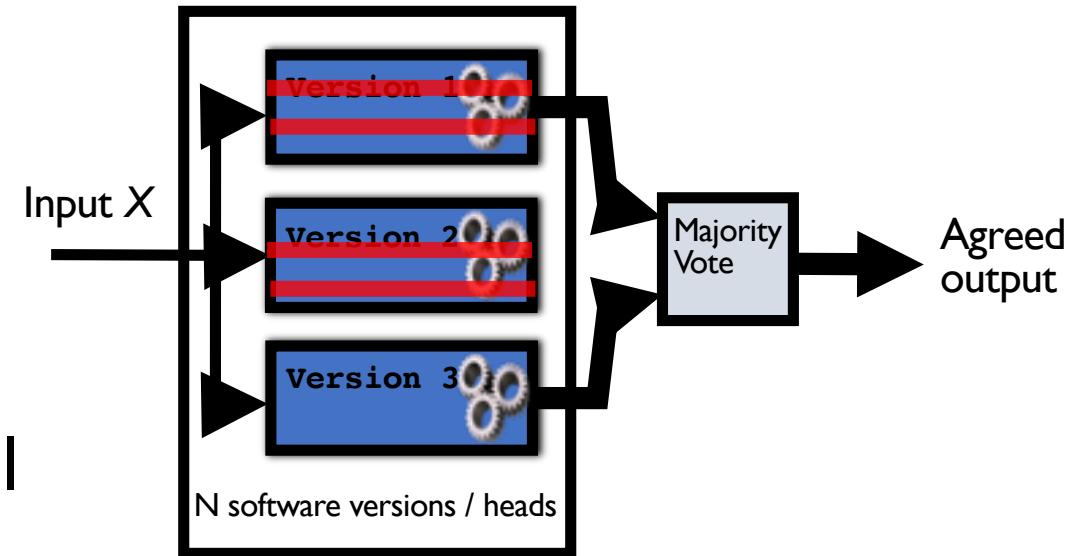
Houston... we have a gap

$$\text{gap} := \frac{\Pr_{X \in \mathcal{D}} [X \in \bigcup_{i=1}^N E(f_i, \mathcal{I})]}{\Pr_{X \in \mathcal{D}} [X \in E(f^*, \mathcal{I})]}$$



N-version-programming criticism

- Strong gap requires independence among heads
 - Correlations hurt!
- Knight-Leveson (1986):
 - “We reject the null hypothesis of full independence at a p-level of 5%”
- Eckhardt et al. (1991):
 - “We tried it at NASA and it wasn’t cost effective”
 - Worst case: 3 versions \Rightarrow 4x fewer errors

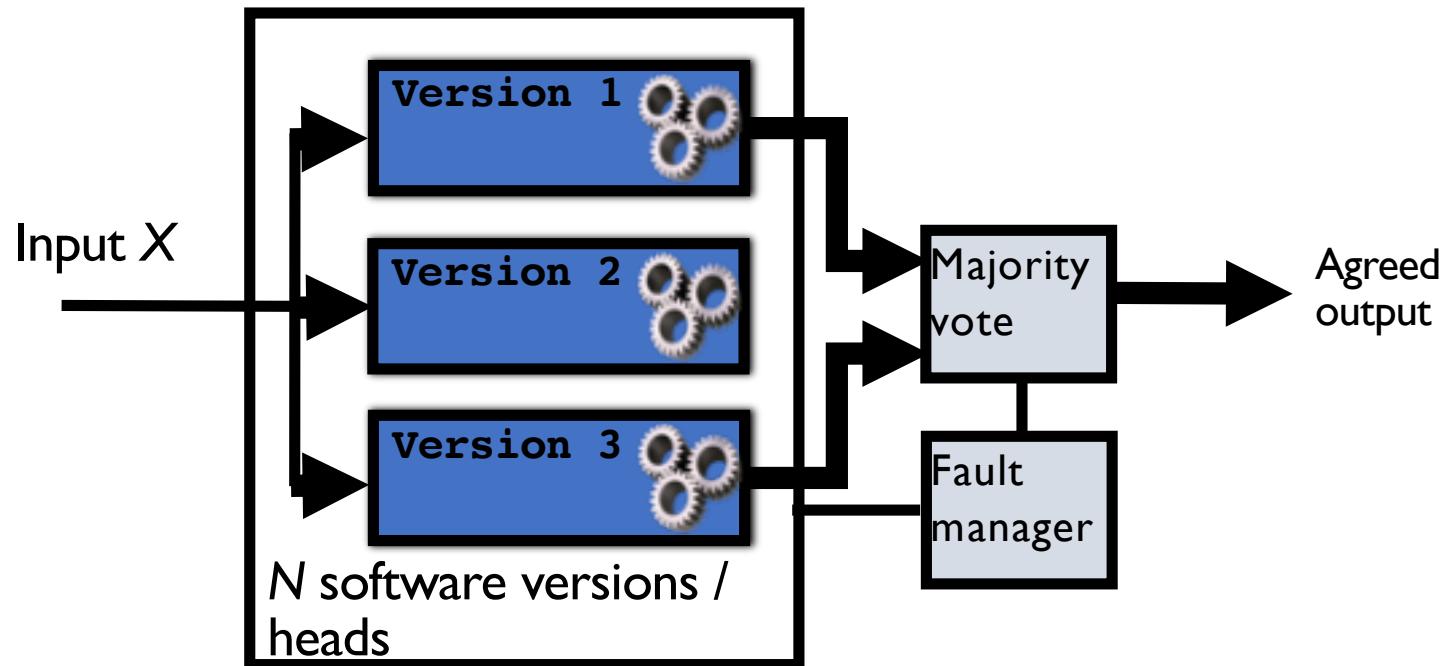


But not everything is a space shuttle...

- Not all software needs to be available at all times!
 - E.g., Smart contracts: How bad if it's down for a while?
- In fact, often ***better no answer than the wrong one***
 - Bugs are often harmful
- ***N-of-N-Version Programming (NNVP)***

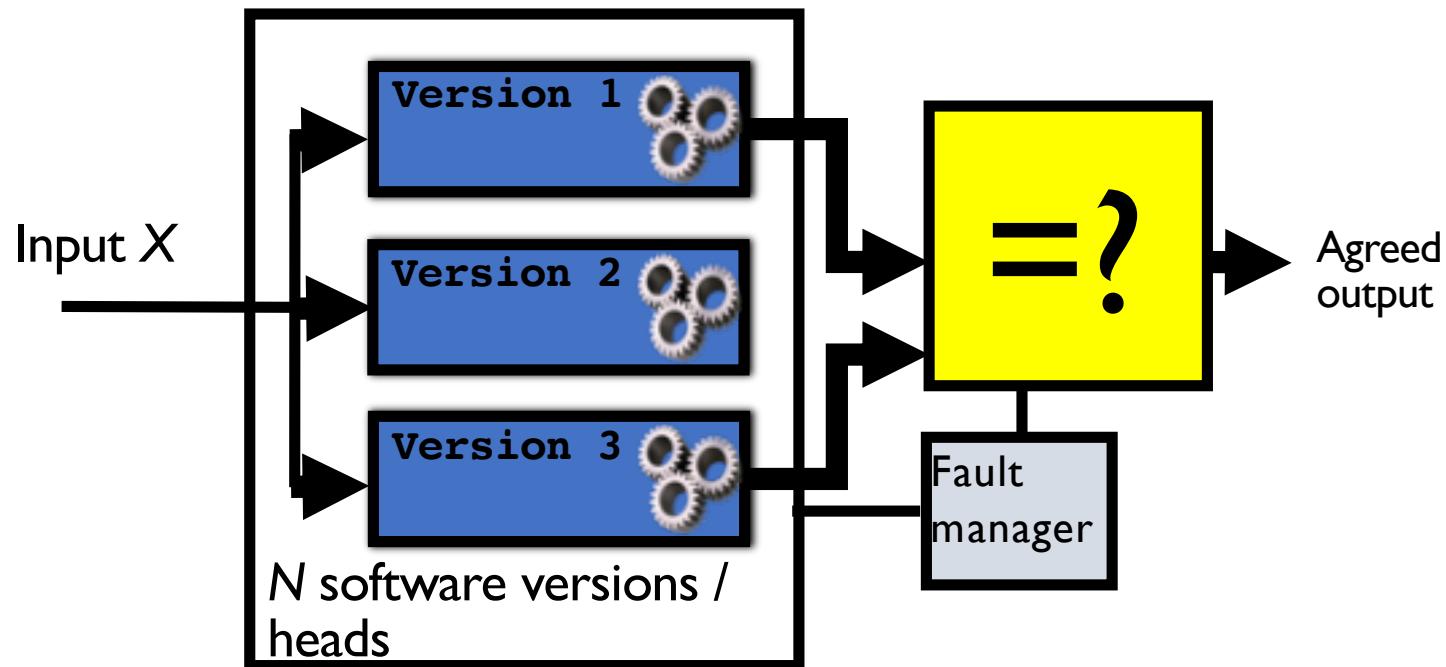


NNVP a.k.a. Hydra Framework



Idea: Strengthen majority vote of N-Version Programming

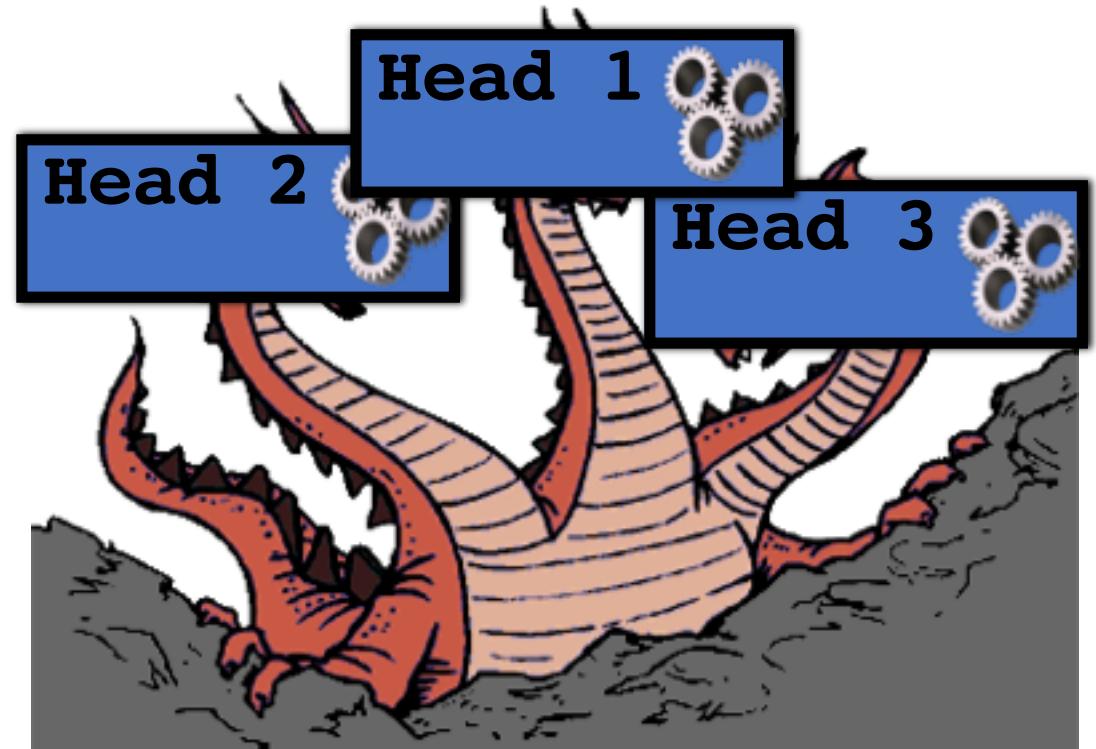
NNVP a.k.a. Hydra Framework



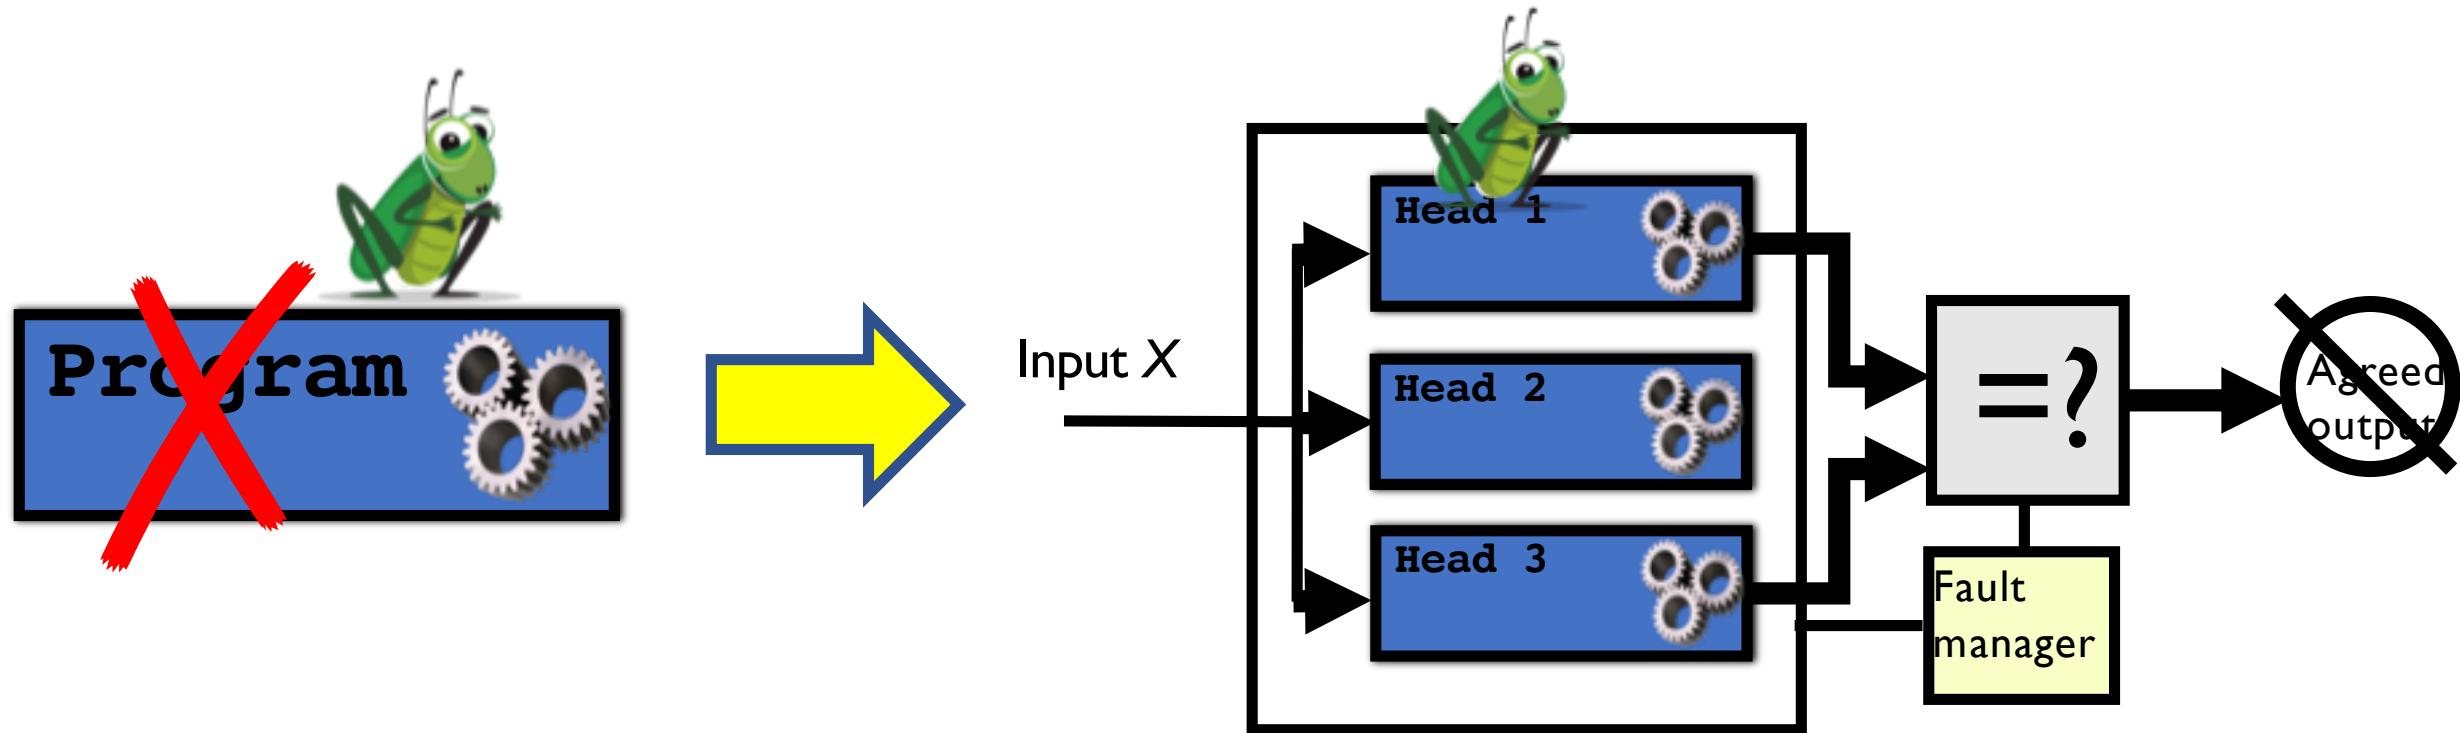
Unless *all versions agree, abort!*

NNVP a.k.a. Hydra

- Aborting in NNVP:
 Correctness \leftarrow Availability
- NASA numbers much better for NNVP
 - Some availability loss, but...
 - gap = 4,409 for $N = 3$ heads
 - gap = 34,546 for $N = 4$ heads
 - Probably even better!



Hydra creates a (strong) gap...

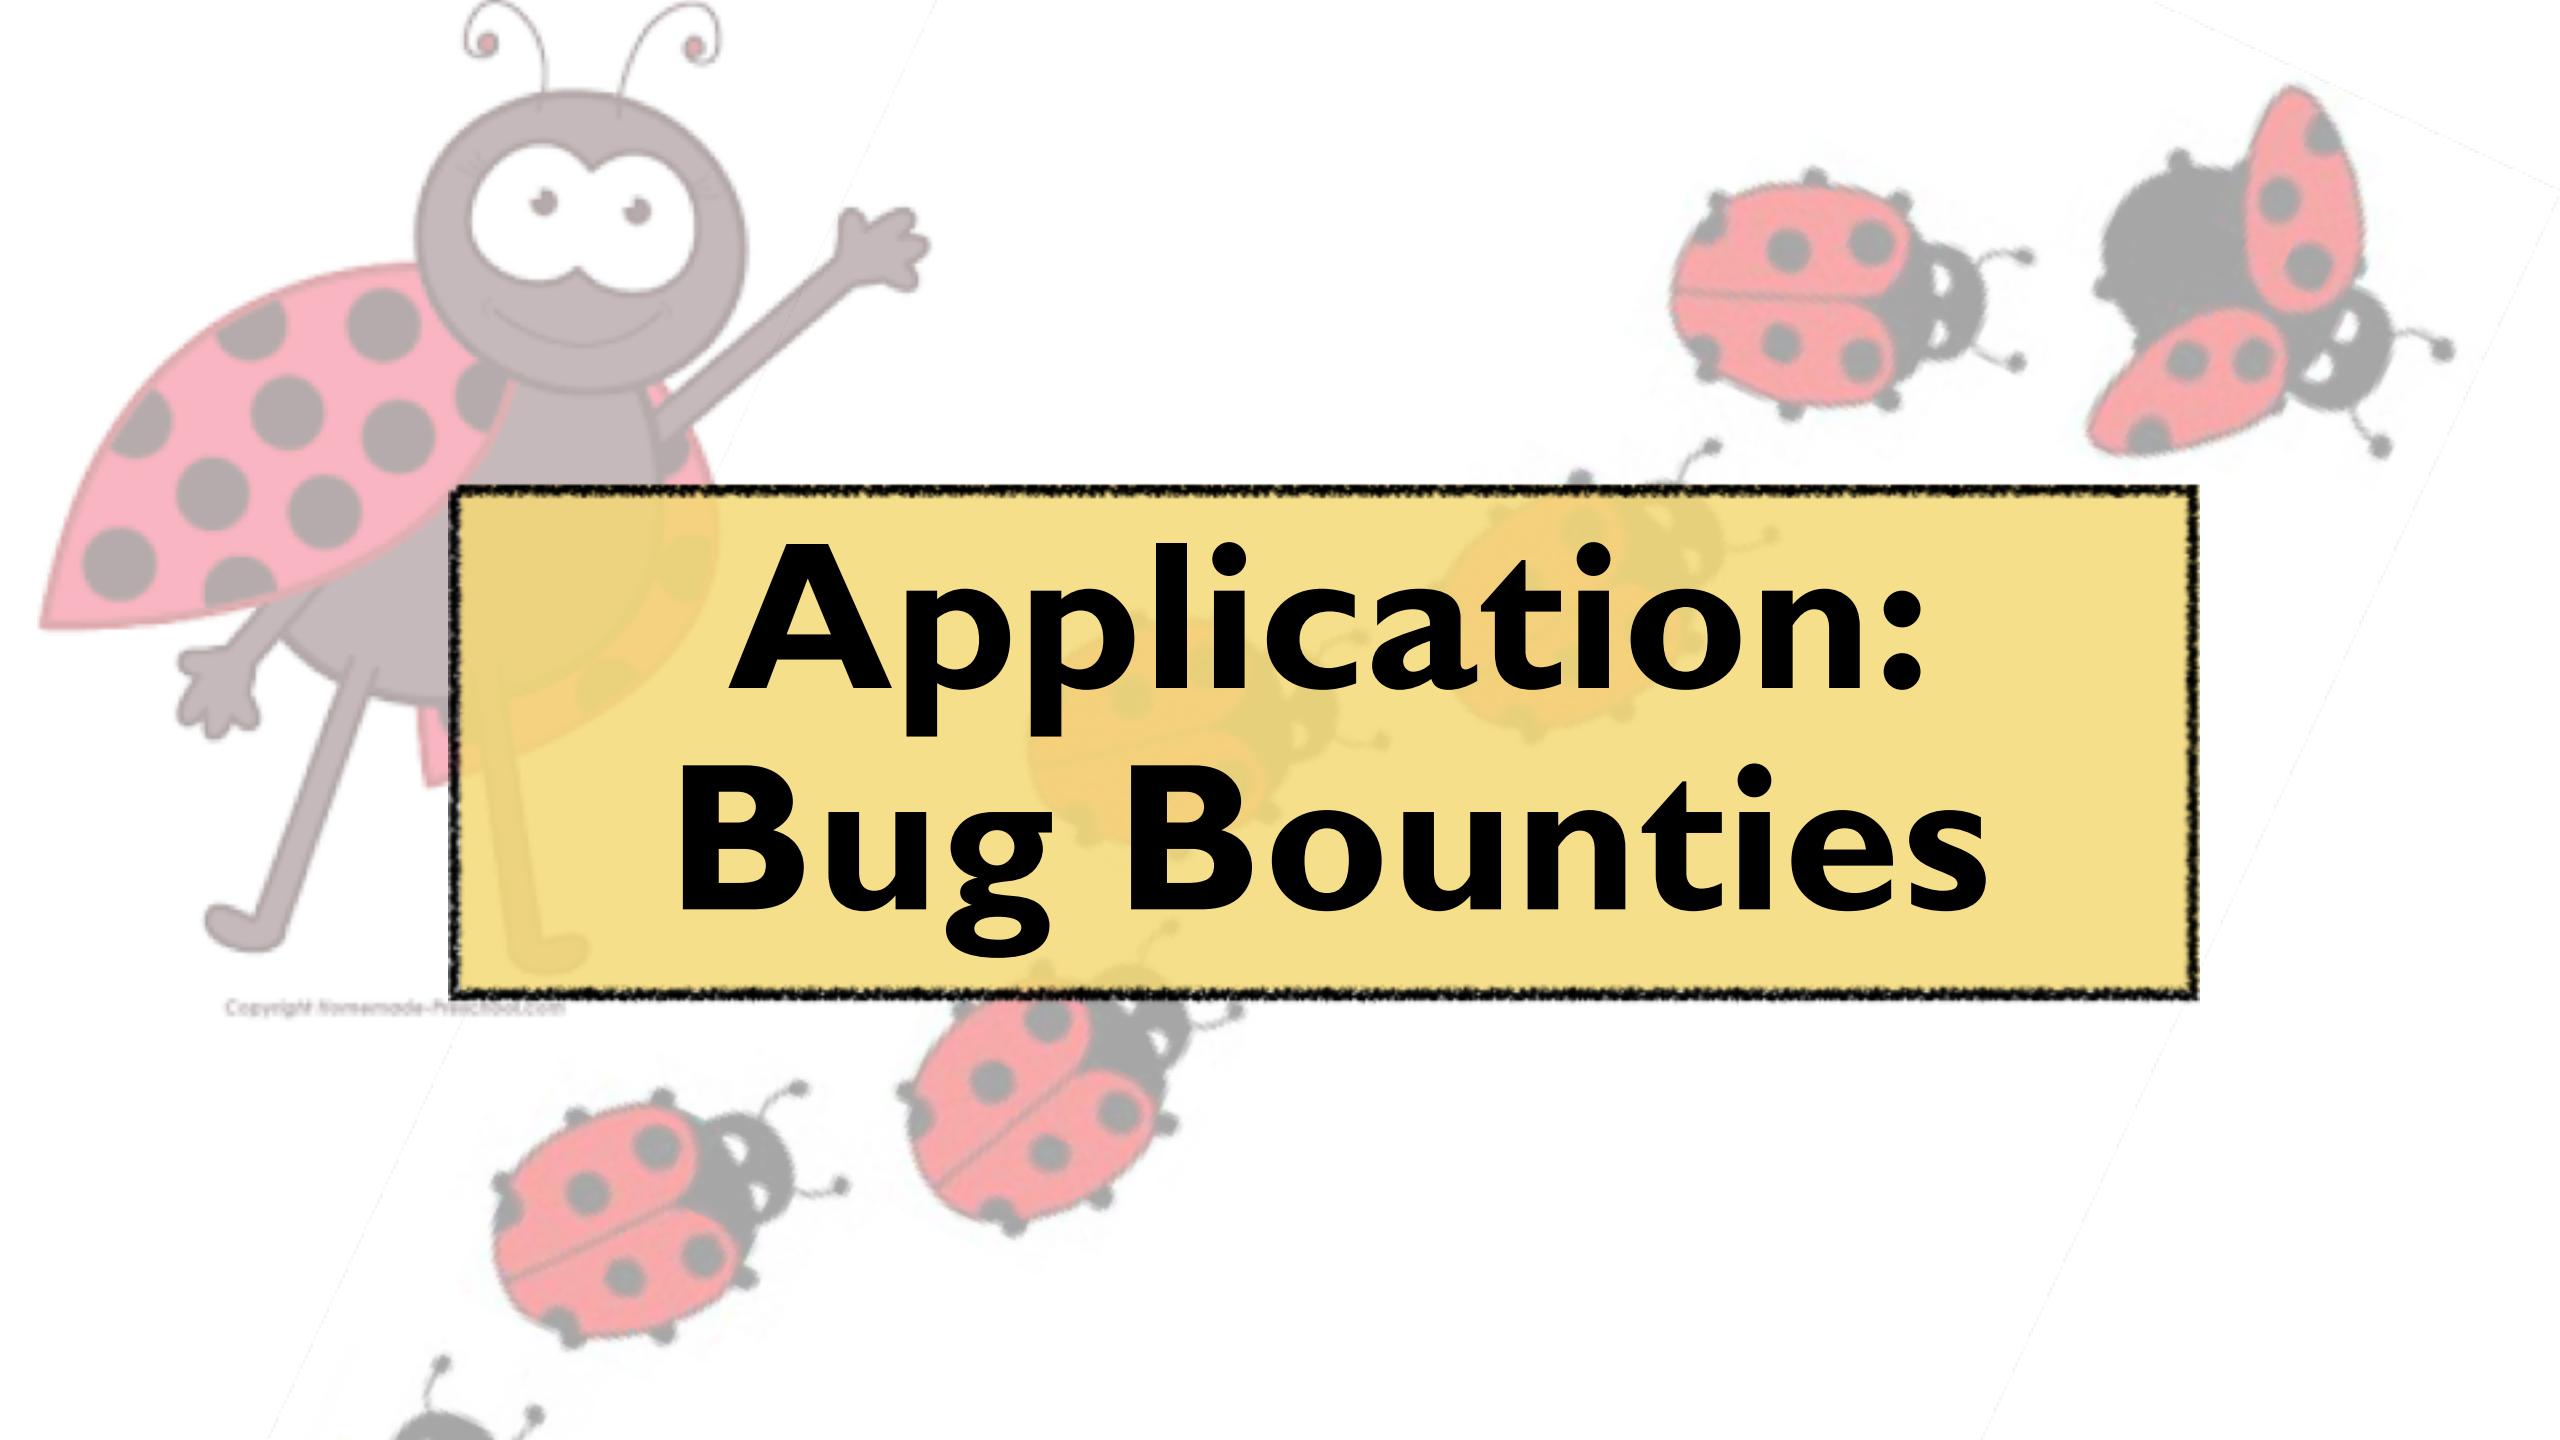


Serious bug in one head now rarely fatal...

Smart contracts are Hydra-friendly!

Contract name	Exploit value (USD)	Root cause	Independence source	Exploit gap
Parity Multisig [3]	180M	Delegate call+unspecified modifier	programmer/language?	✓/✗
The DAO* [19]	150M	Re-entrancy	language	✓
SmartBillions [20]	500K	Bug in caching mechanism	programmer	✓
HackerGold (HKG)* [21]	400K	Typo in code	programmer+language	✓
MakerDAO* [22]	85K	Re-entrancy	language	✓
Rubixi [23]	<20K	Wrong constructor name	programmer+language	✓
Governmental [23]	10K	Exceeds gas limit	None?	✗

Hydra could probably have addressed cases in green and yellow vulnerabilities



Application: Bug Bounties

Copyright Nomadahouse.com

Bug bounties

- Reward for responsible disclosure of software vulnerabilities
- Key element of nearly all security assurance programs
 - E.g., Apple (up to \$200k)

bugcrowd

COMPANY	NEW	RWARD	SWAG	HALL OF FAME
1Password	✓	✓		✓
123 Contact Form				✓
99designs		✓		✓
Abacus				✓
ABN Amro				
Acorns LLC		✓		✓
Acquia				✓
Active Campaign				✓

Some problems with bug bounties:

1. Bounties often fail to incentivize disclosure
 - Apple: $\leq \$200k$ bounty
 - Zerodium: \$1.5 million for certain iPhone jailbreaks
2. Time lag between reporting and action
 - Weaponization can happen *after* disclosure
3. Bounty administrator doesn't always pay!

Home > Vulnerabilities



Researchers Claim Wickr Patched Flaws but Didn't Pay Rewards

By [Ionut Arghire](#) on October 31, 2016

3. Bounty administrator doesn't always pay!

The perfect bug bounty

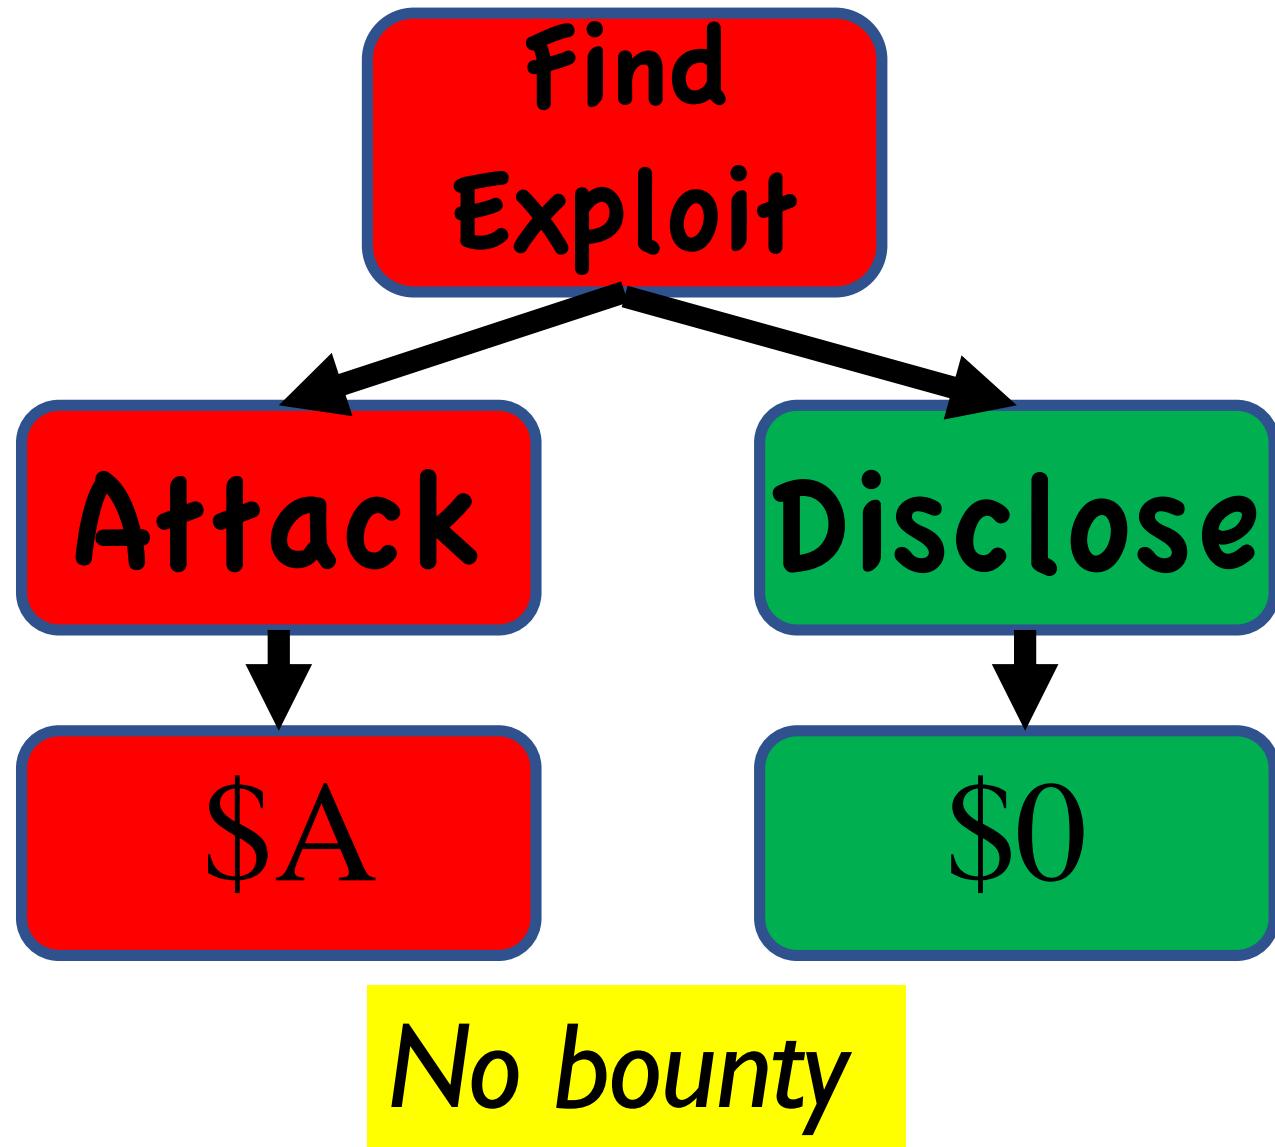


- 1. High leverage:** Small bounty incentivizes disclosure for valuable program
- 2. Automatic payout:** Bounty hunter need not trust bounty administrator to pay
 - Censorship-resistant, verifiable
- 3. Automatic remediation:** Immediate intervention in affected software

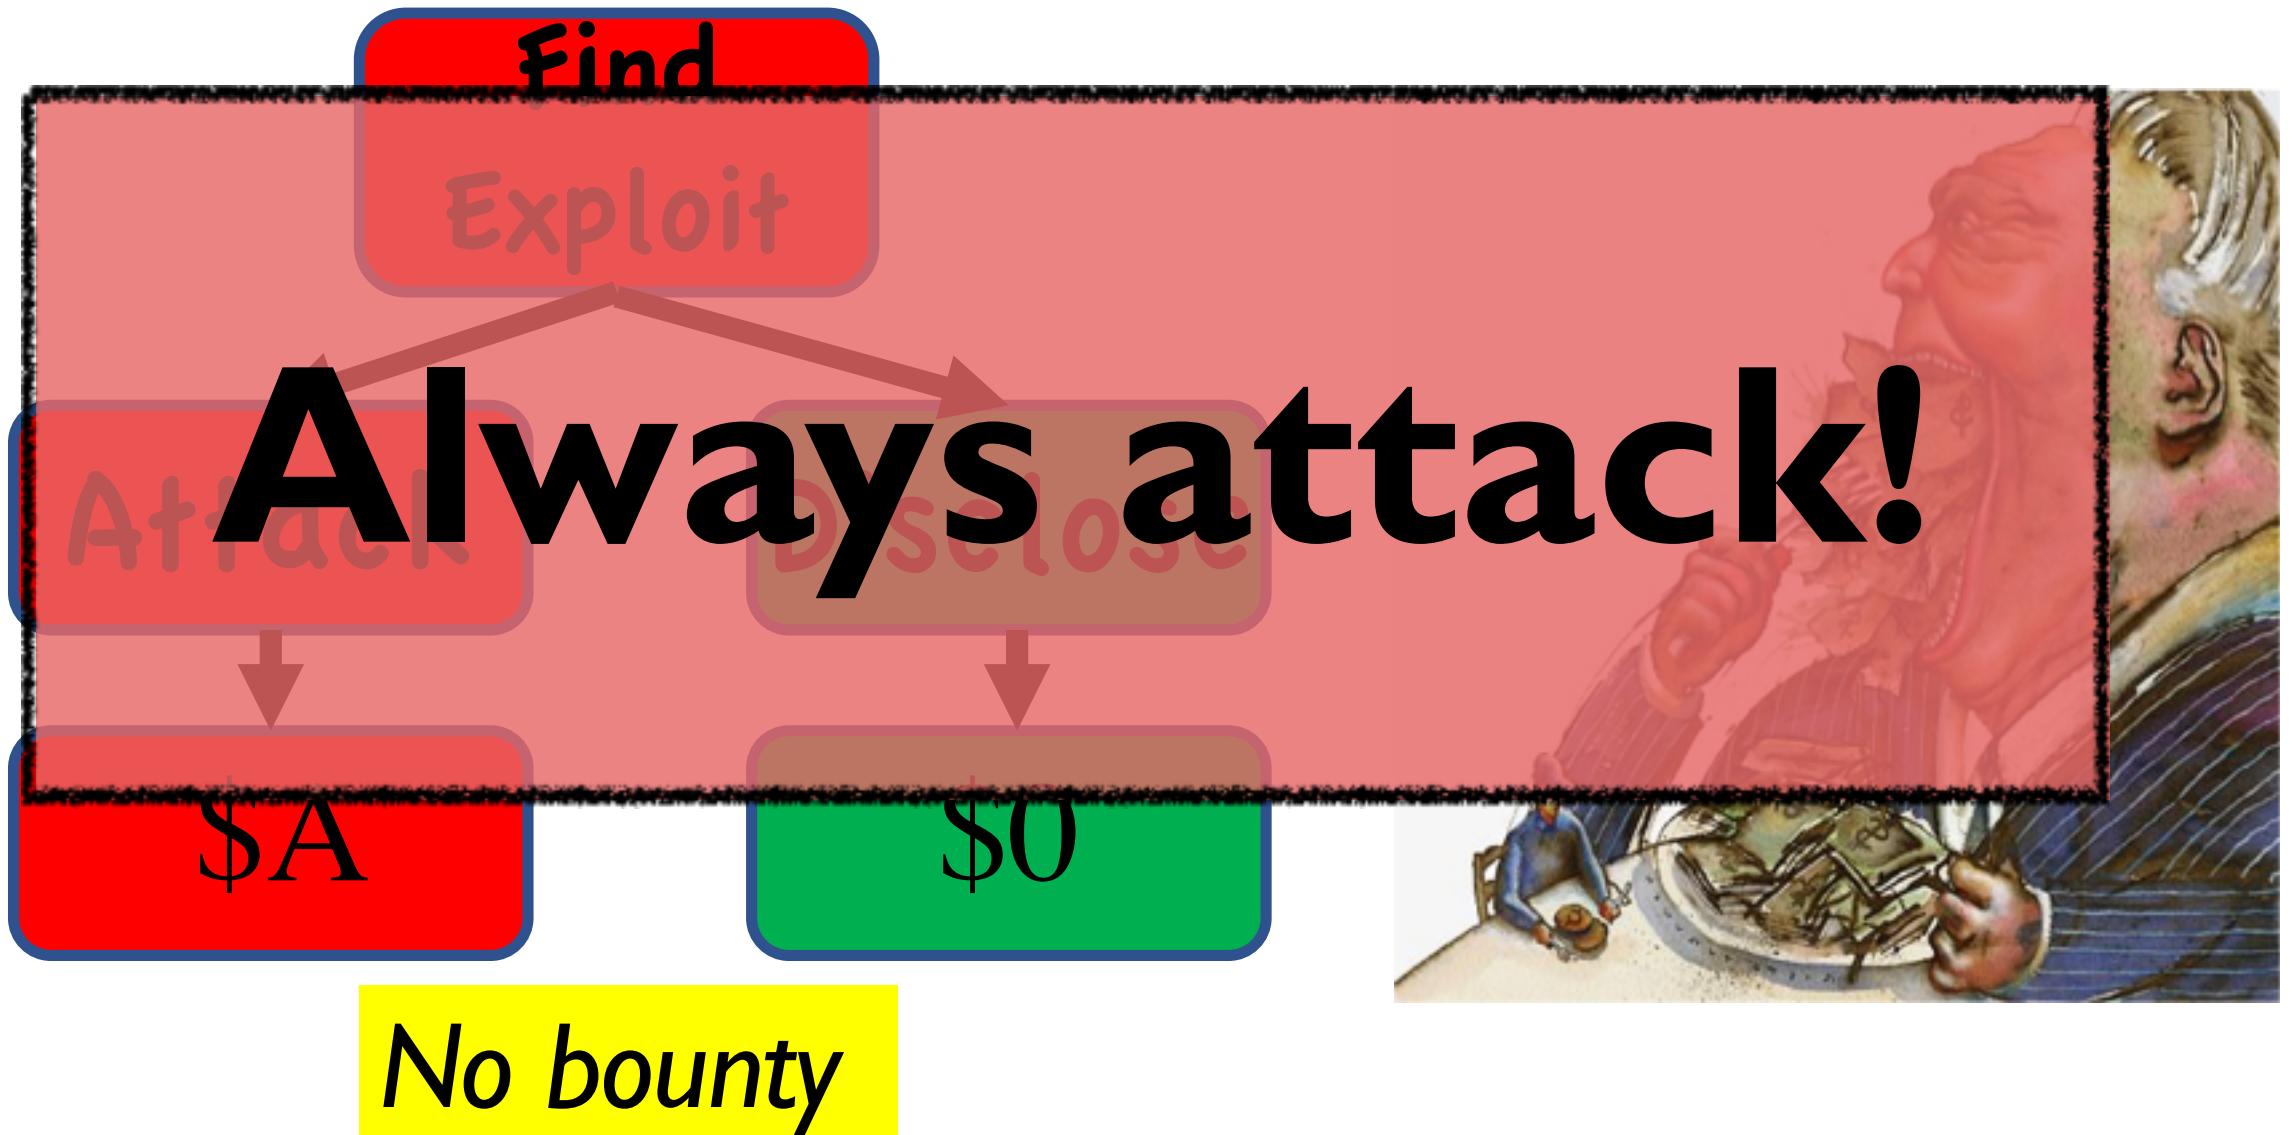
Bug bounties: The Rational Attacker's Game



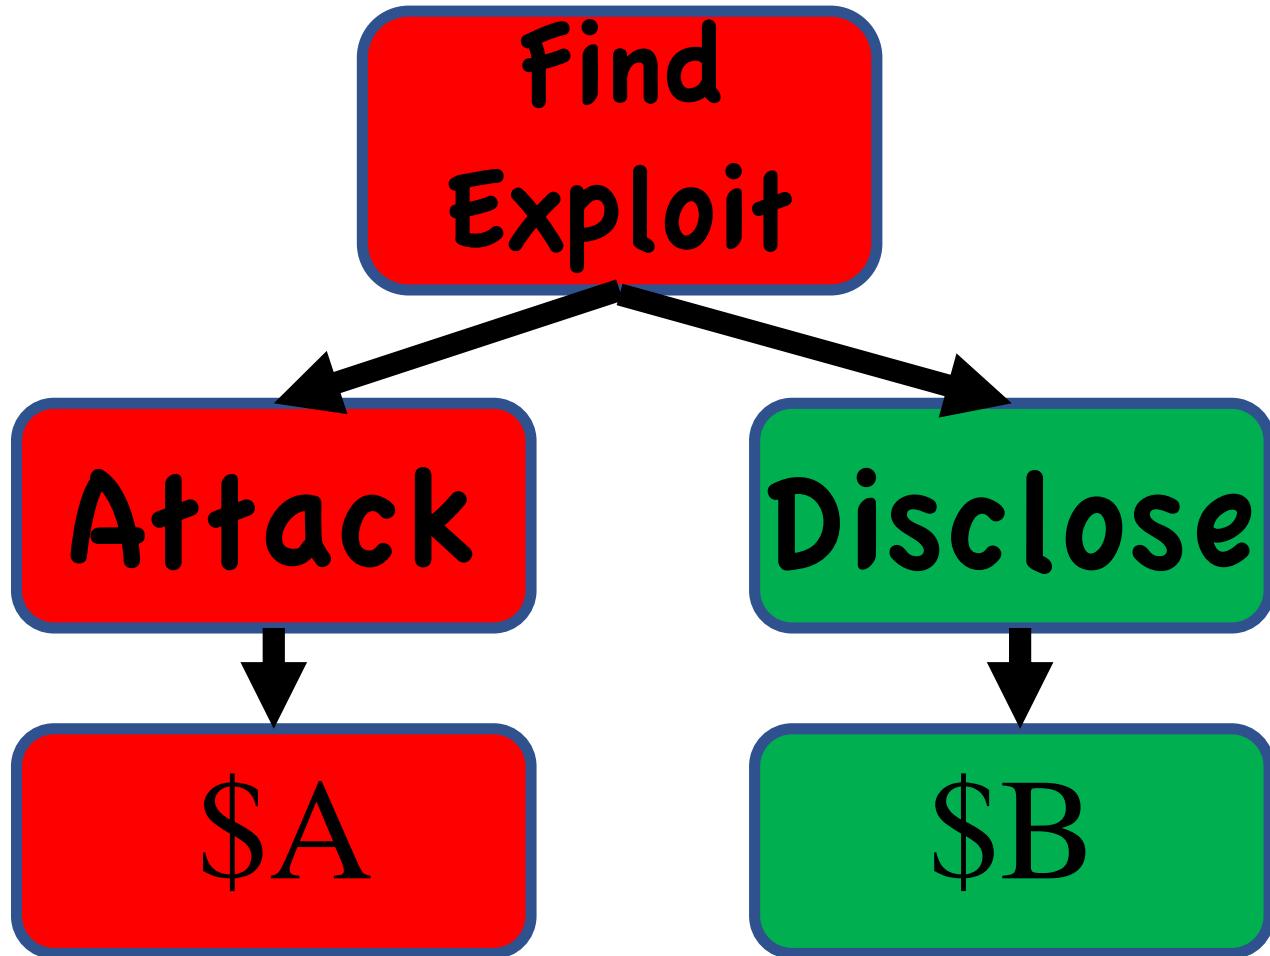
Bug bounties: The Rational Attacker's Game



Bug bounties: The Rational Attacker's Game



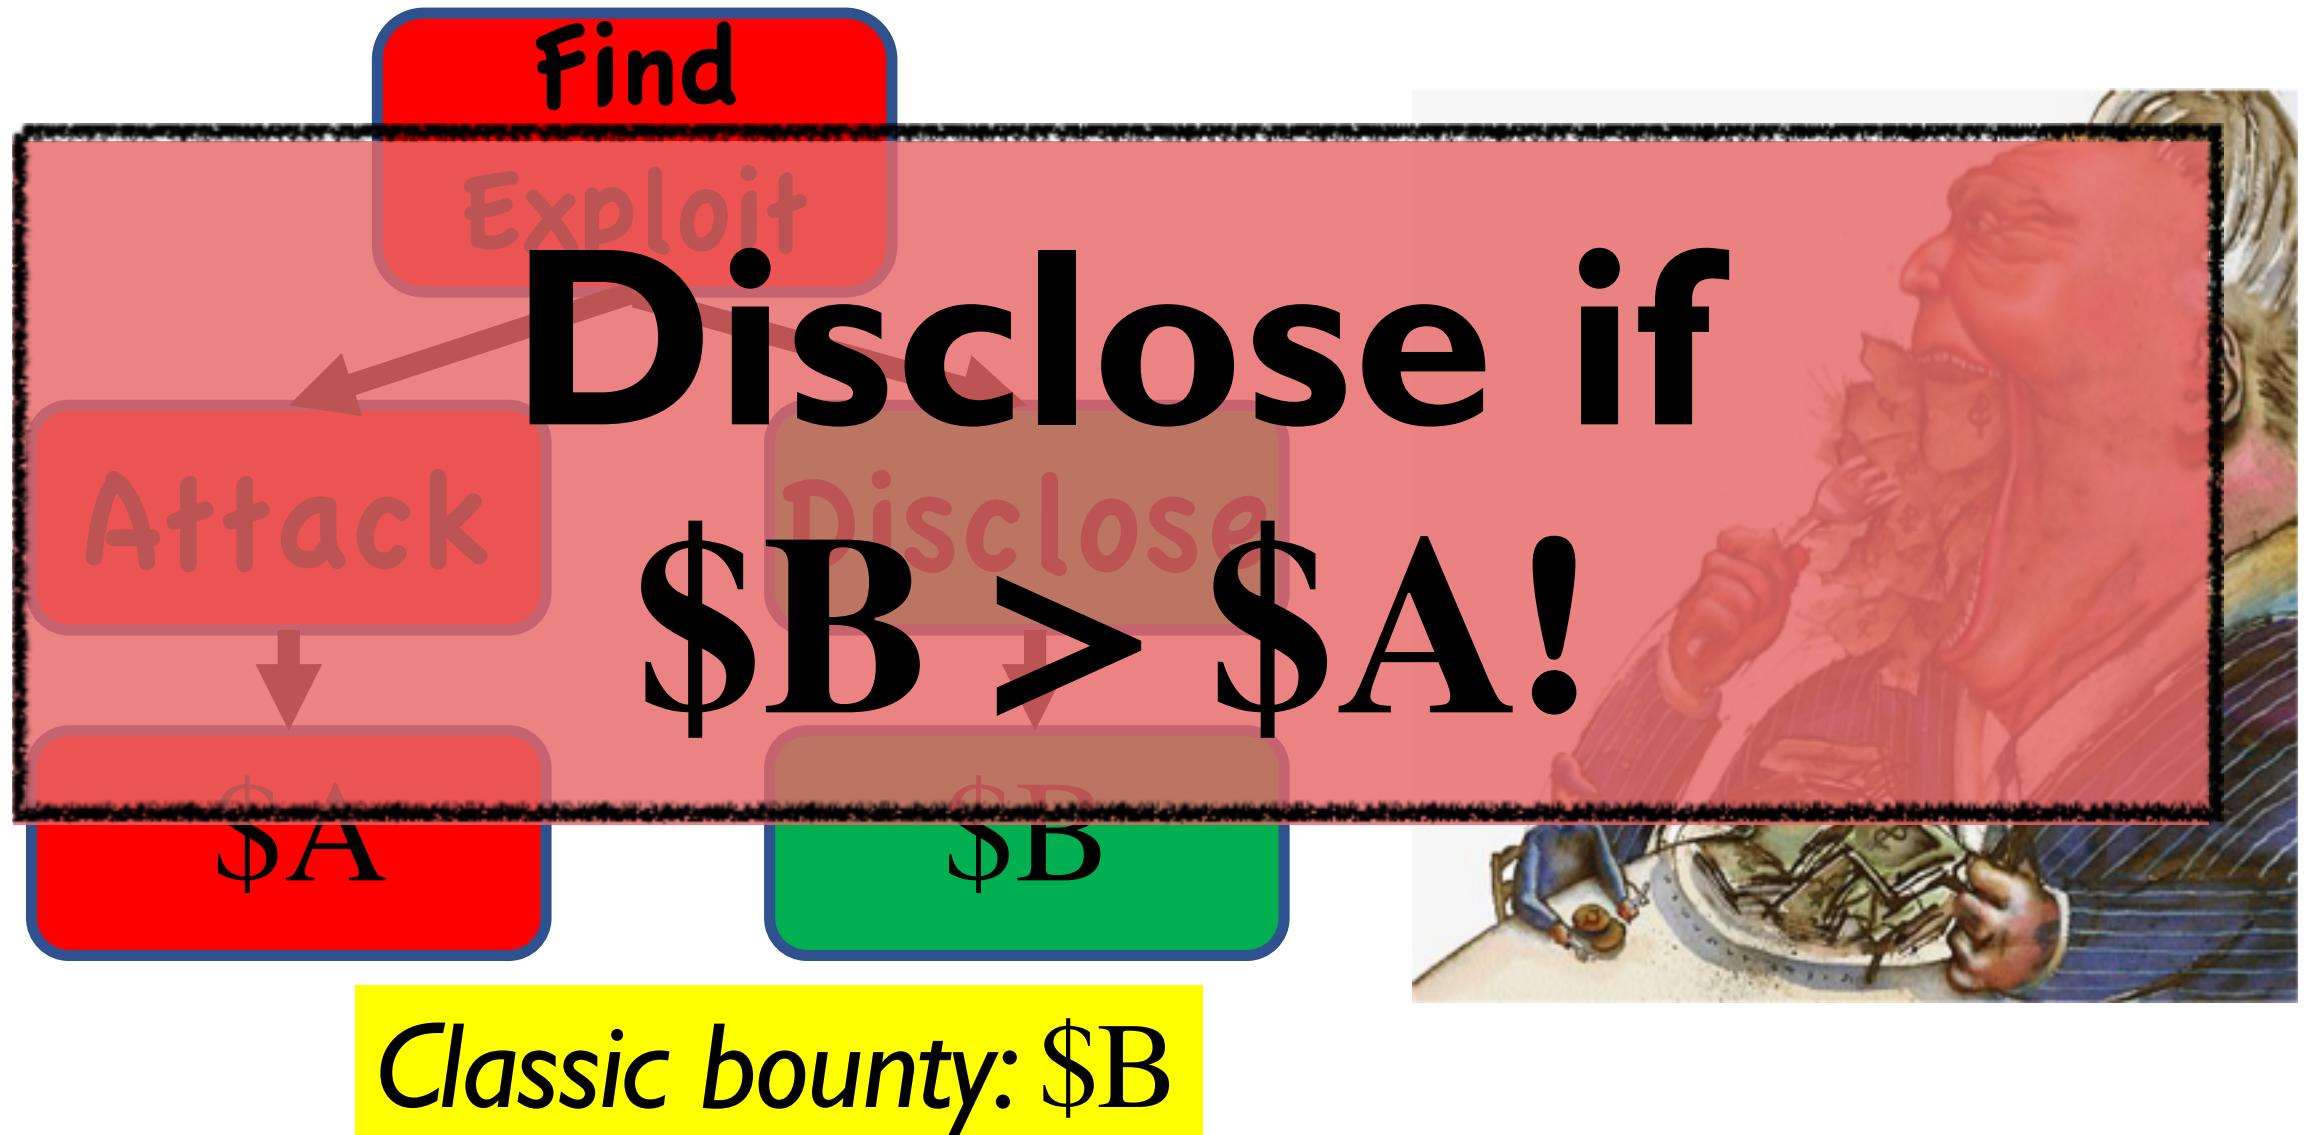
Bug bounties: The Rational Attacker's Game



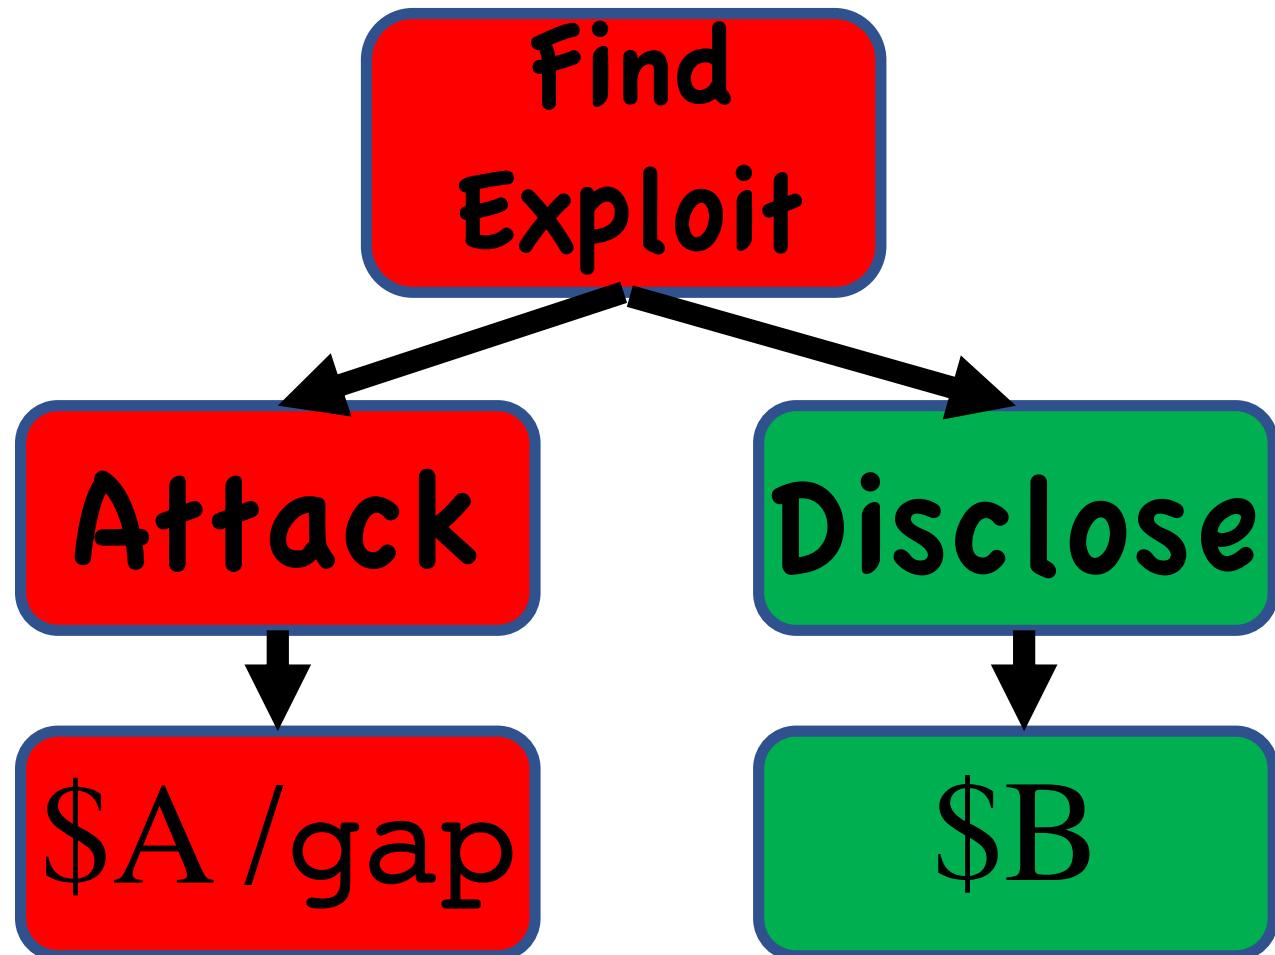
Classic bounty: \$B



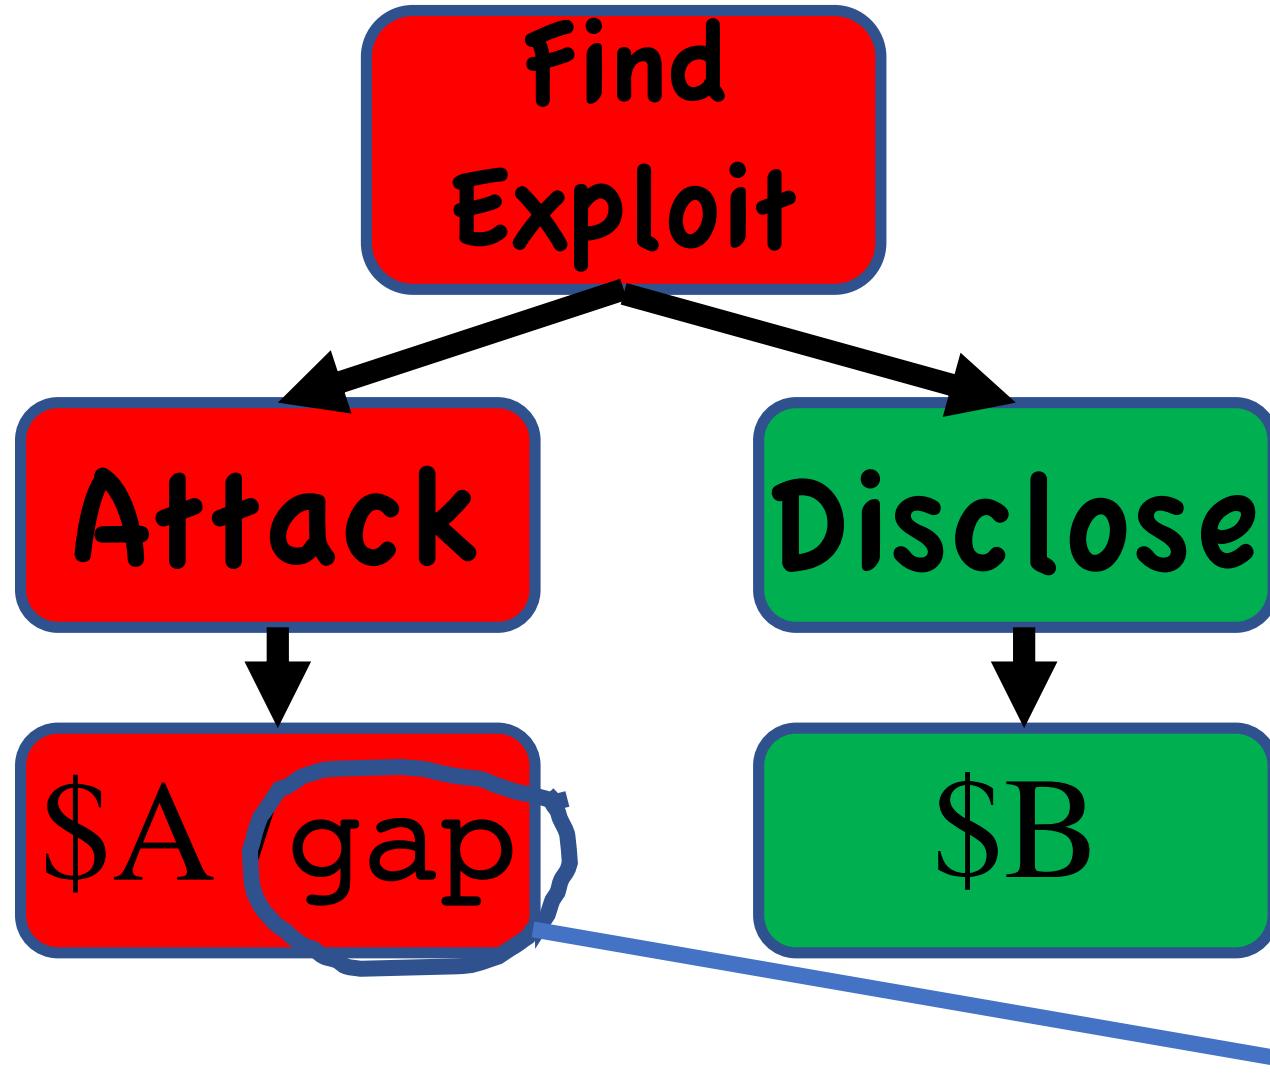
Bug bounties: The Rational Attacker's Game



Our goal: High leverage



Our goal: High leverage



Our goal: High leverage

find

Exploit

Disclose if
 $\$B > \$A / \text{gap!}$

Attack

Disclose

$\$A / \text{gap}^*$

$\$B$

Exploit
gap



Wait a minute...

Program

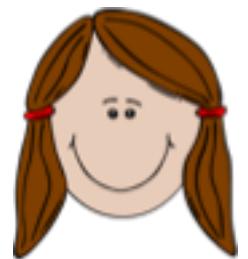
Value: \$A



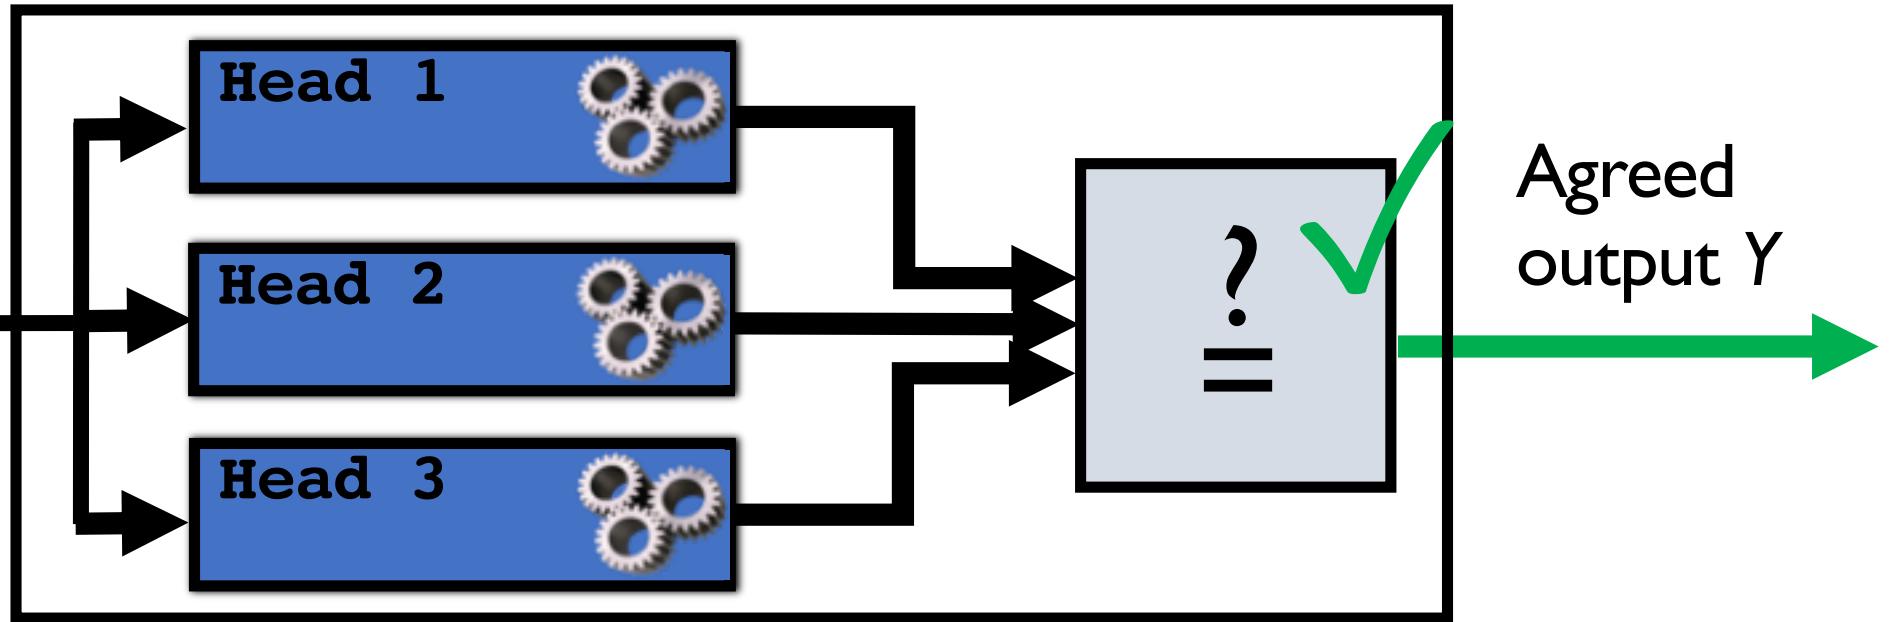
Disclose, i.e.,
don't attack
even though
 $\$B < \A ?!



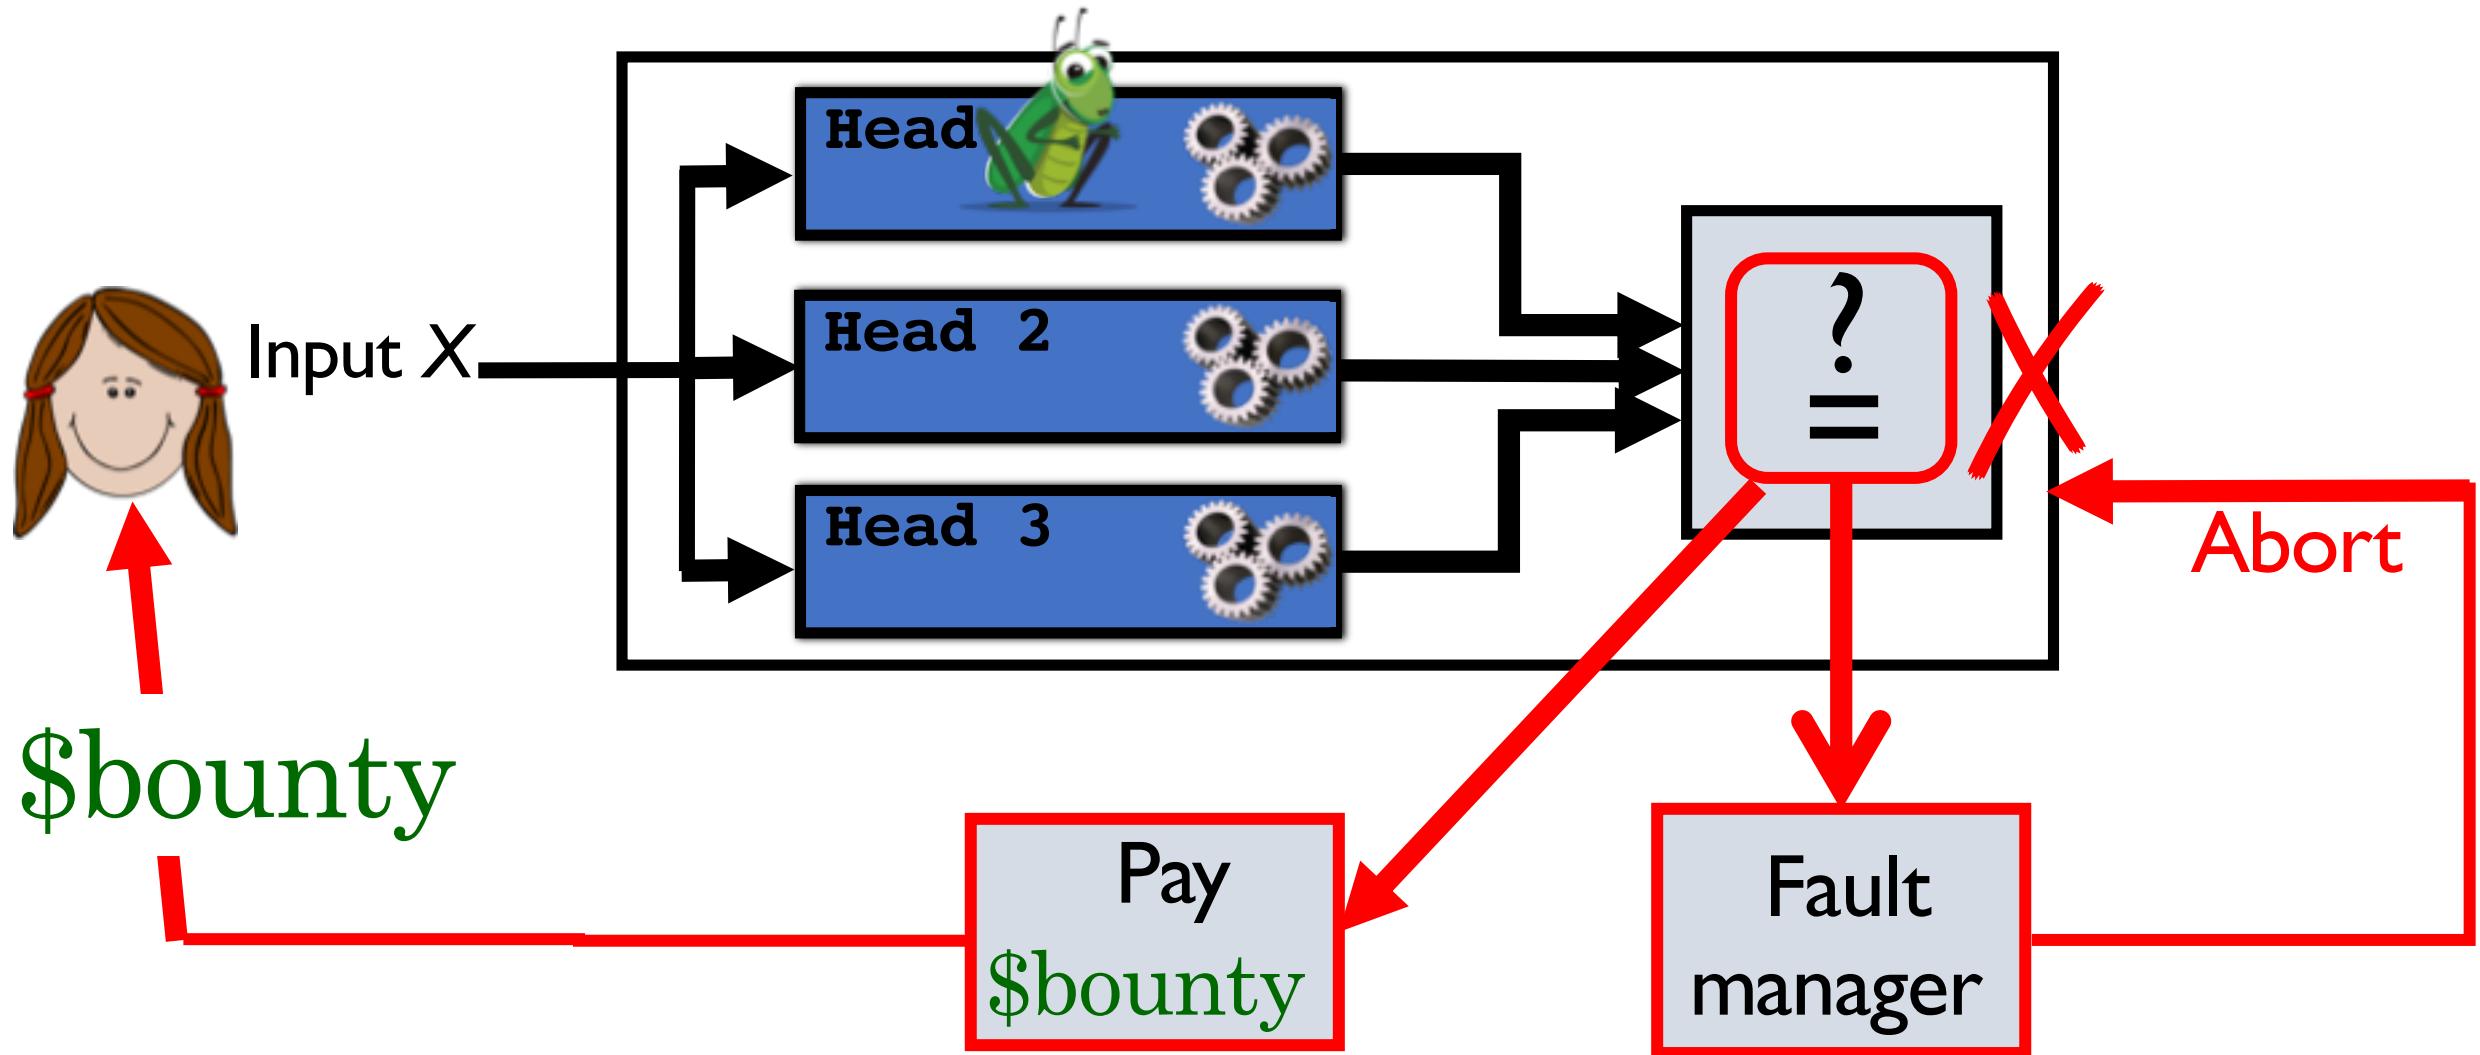
The Hydra Framework for Bug Bounties



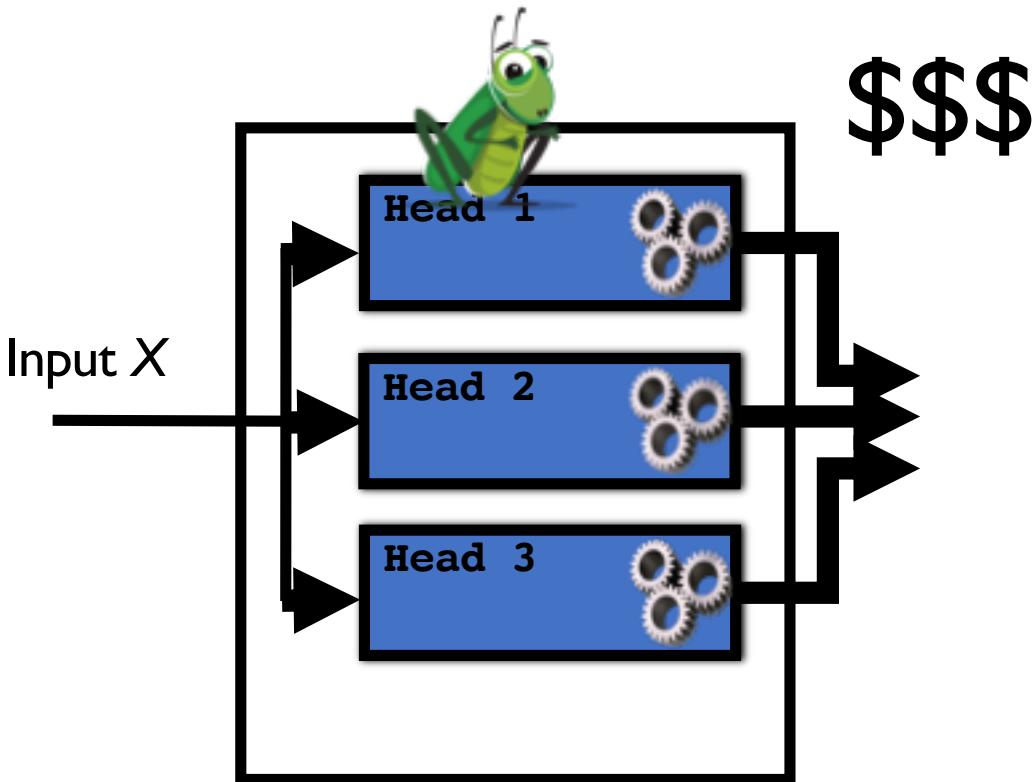
Input X



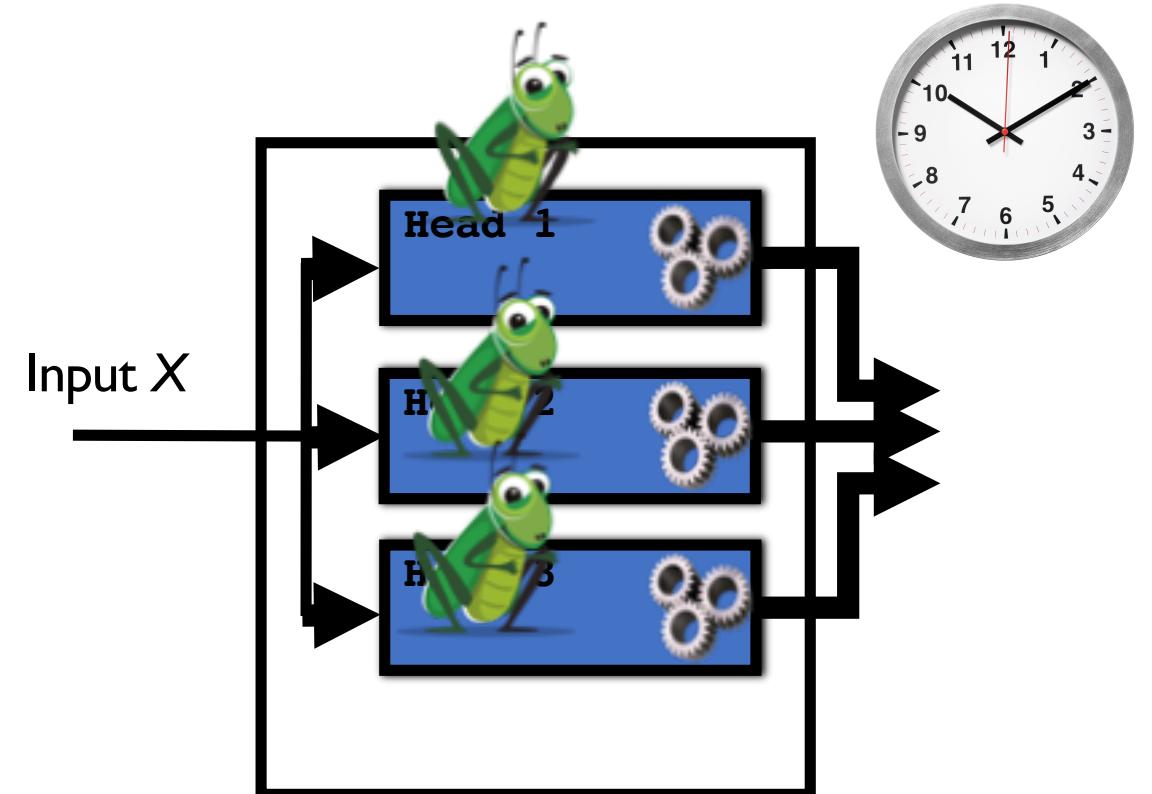
The Hydra Framework for Bug Bounties



The Hydra Hacker's Dilemma



Claim bounty (\$B) now?



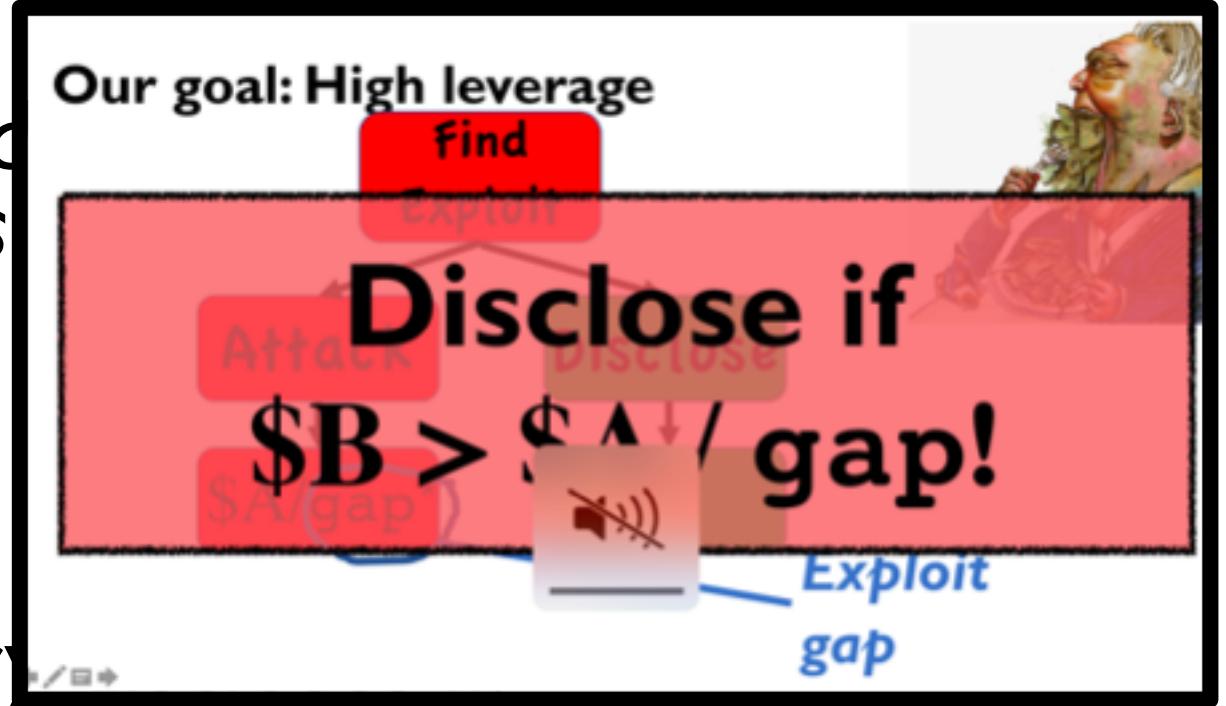
Try to break all heads (\$A)?

Recall:

$$\text{gap} := \frac{\Pr_{X \in \mathcal{D}} \left[X \in \bigcup_{i=1}^N E(f_i, \mathcal{I}) \right]}{\Pr_{X \in \mathcal{D}} [X \in E(f^*, \mathcal{I})]}$$

Hydra Framework → High leverage

- Suppose strong rational actors as *all honest bounty hunters*
- Suppose:
 - Contract worth $\$A$
 - Bounty $\$B$
- Then (we prove) adversary



$$\$B > \$A / (\text{gap} + 1).$$

Example

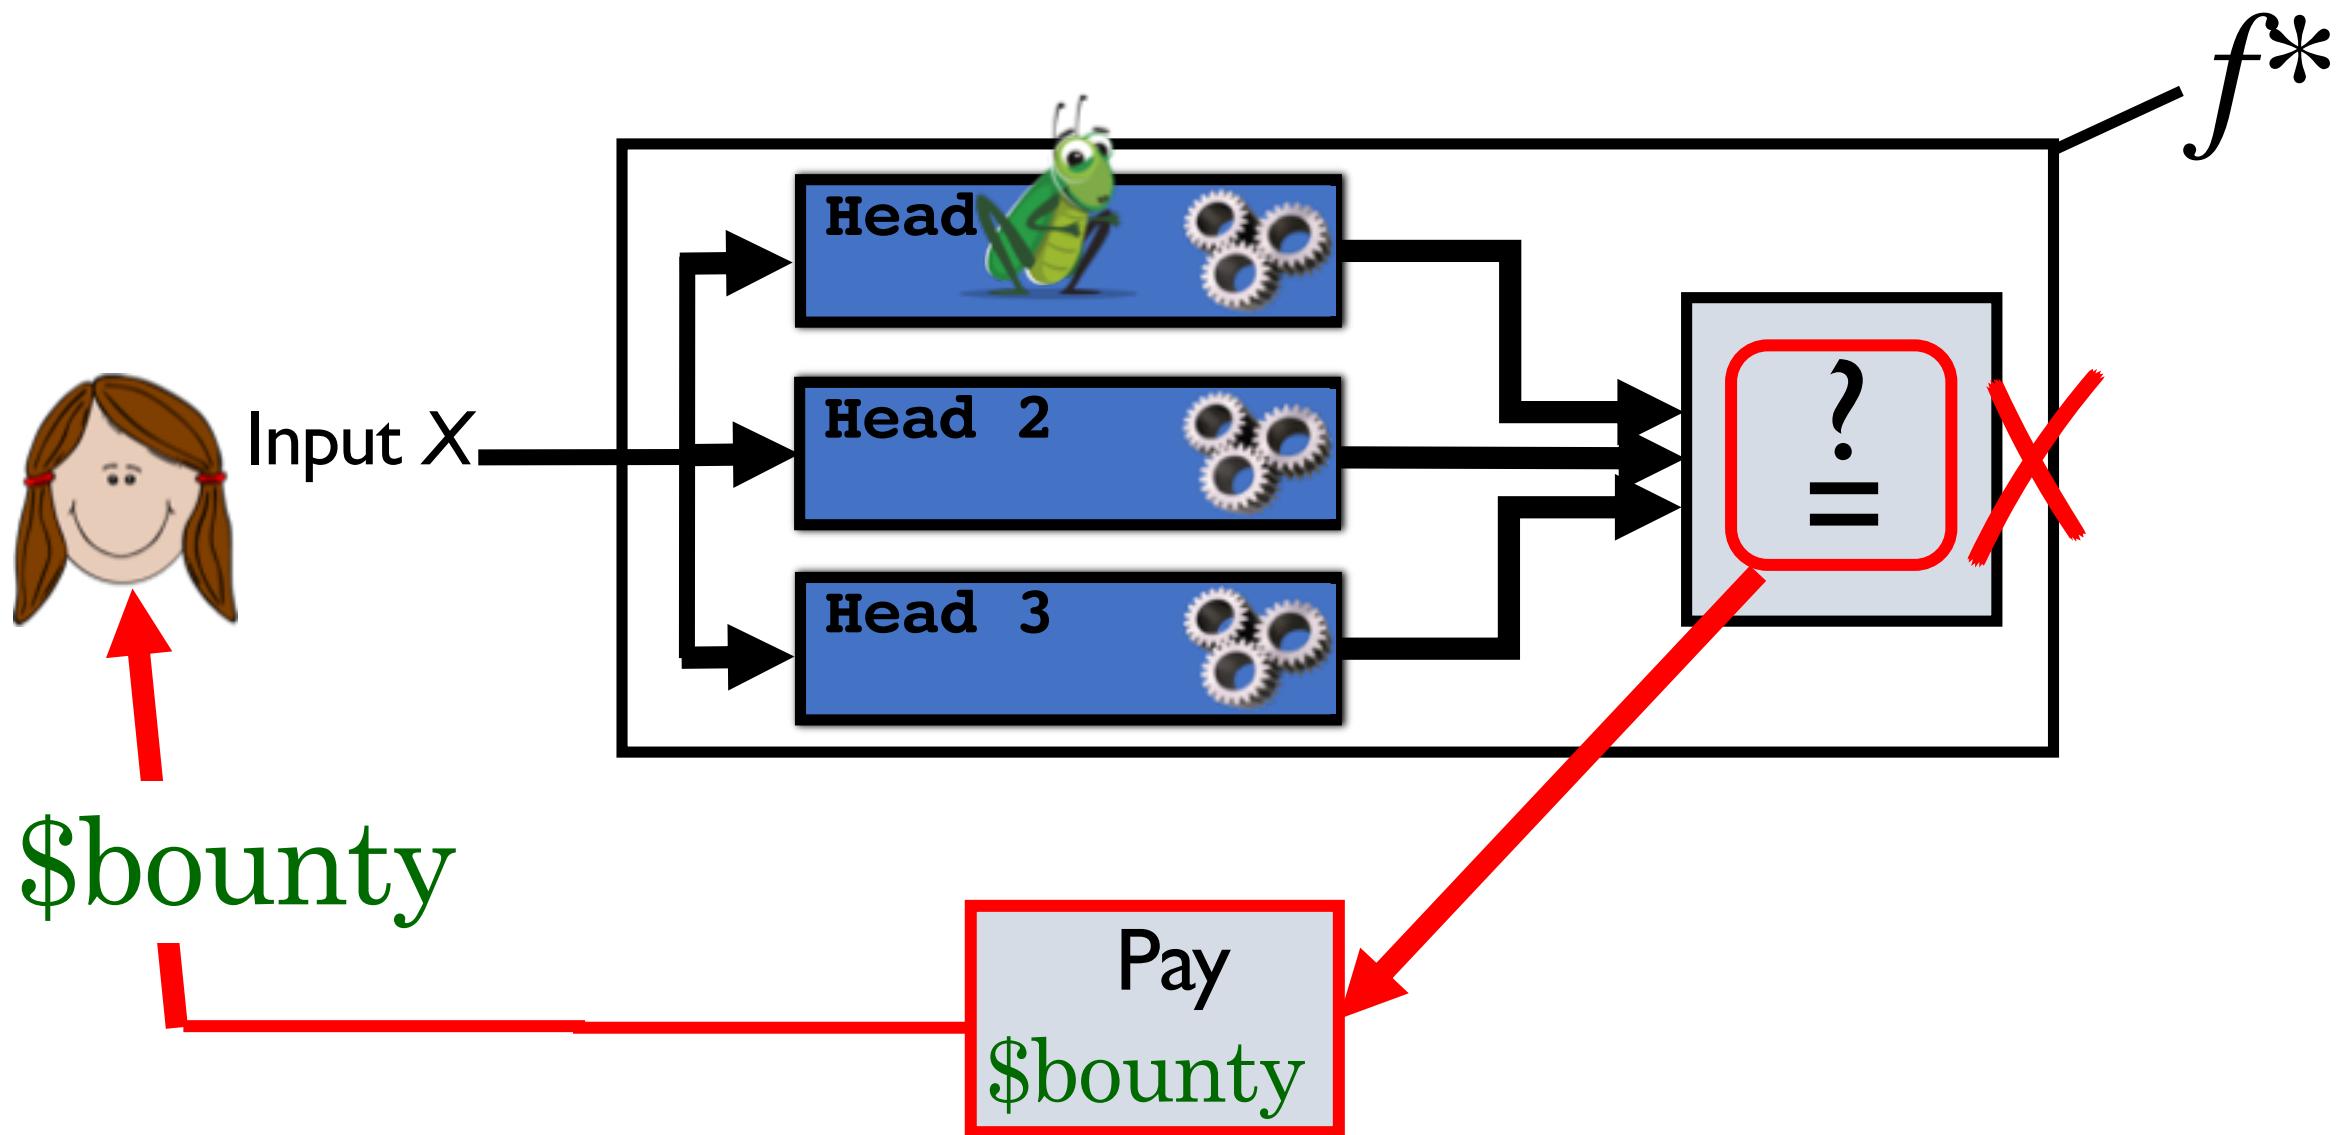
- Recall: NASA experiments imply:
 - gap = 4,409 for $N = 3$ heads
 - gap = 34,546 for $N = 4$ heads
- So...
 - **Approx \$1 billion** contract (e.g., OmiseGo)
 - $N = 4$
 - **\$30k bounty** incentivizes adversary to disclose!

The perfect bug bounty



1. **High leverage:** Small bounty incentivizes disclosure for valuable program
2. **Automatic payout:** Bounty hunter need not trust bounty administrator to pay
 - Censorship-resistant, verifiable
3. **Automatic remediation:** Immediate intervention in affected software

It's a smart contract! It's automatically automatic!



The perfect bug bounty

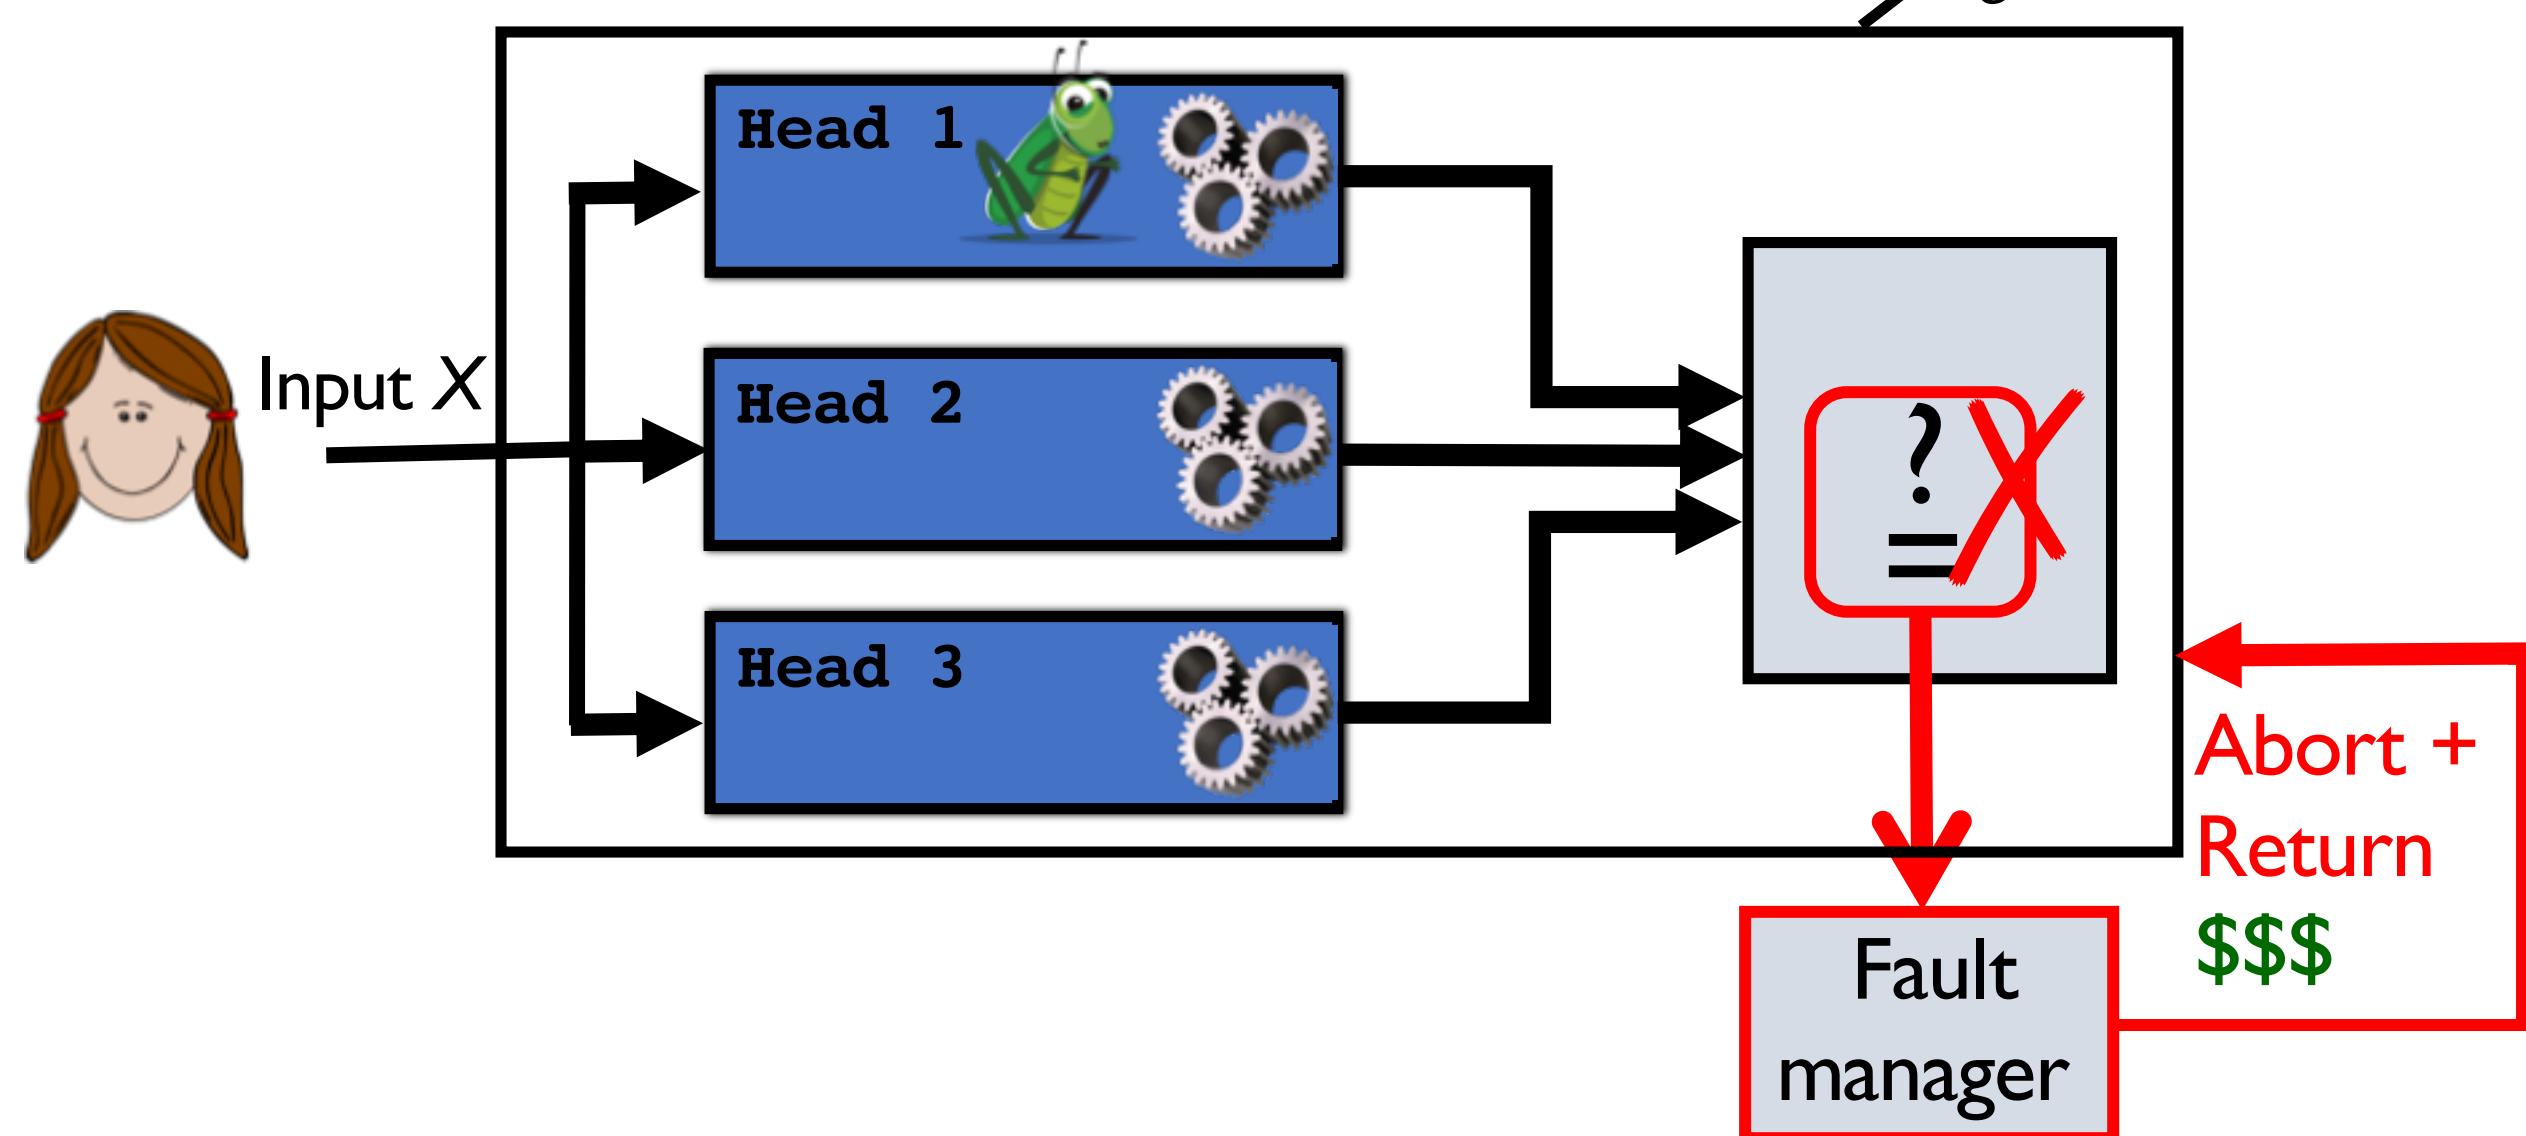


1. **High leverage:** Small bounty incentivizes disclosure for valuable program
2. **Automatic payout:** Bounty hunter need not trust bounty administrator to pay
 - Censorship-resistant, verifiable
3. **Automatic remediation** Immediate intervention in affected software

How to remediate if contract fails?

- The DAO (\$50+ million stolen)
 - **Remedy:** Fork *returned money (in ETH-land) to victims*
- Parity multisig hack (\$30 million stolen)
 - **(Partial) Remedy:** White hats “stole” \$78 mil.; *returned money to victims*
 - (Two co-authors of Hydra paper among these hackers...)
- Parity multisig hack—Redux! (\$150 million frozen)
 - **(Proposed) Remedy:** Unfreeze funds and return to victims

The Hydra Framework for Bug Bounties



The perfect bug bounty



- ✓ 1. **“Strong exploit gap”**: Small bounty incentivizes disclosure for valuable program
- ✓ 2. **Automatic payout**: Bounty hunter need not trust bounty administrator to pay
 - Censorship-resistant, verifiable
- ✓ 3. **Automatic remediation**: Immediate intervention in affected software

Smart contracts: Perfect bug-bounty targets

- Vulnerable:
 - Bug-prone / hard to code correctly
 - Many \$\$\$ per line of code
- But promising:
 - Hydra-friendly
 - Support (1) High leverage; (2) Automated payout; and (3) Reasonable remediation
 - **Bonus:** Automatic value-at-risk assessment
 - First opportunity to reason about bounty amounts in principled way!



Implementation

- ERC20
 - Standard token-management contract
 - $N = 3$
 - $\$bounty = 3\text{ETH} \approx \1500
 - Deployed @ [0xf4ee935a3879ff07362514da69c64df80fa28622](https://etherscan.io/address/0xf4ee935a3879ff07362514da69c64df80fa28622)
- Generalized Monty-Hall game
 - Extension of Monty Hall game to K out of M doors
 - In progress

Reveal



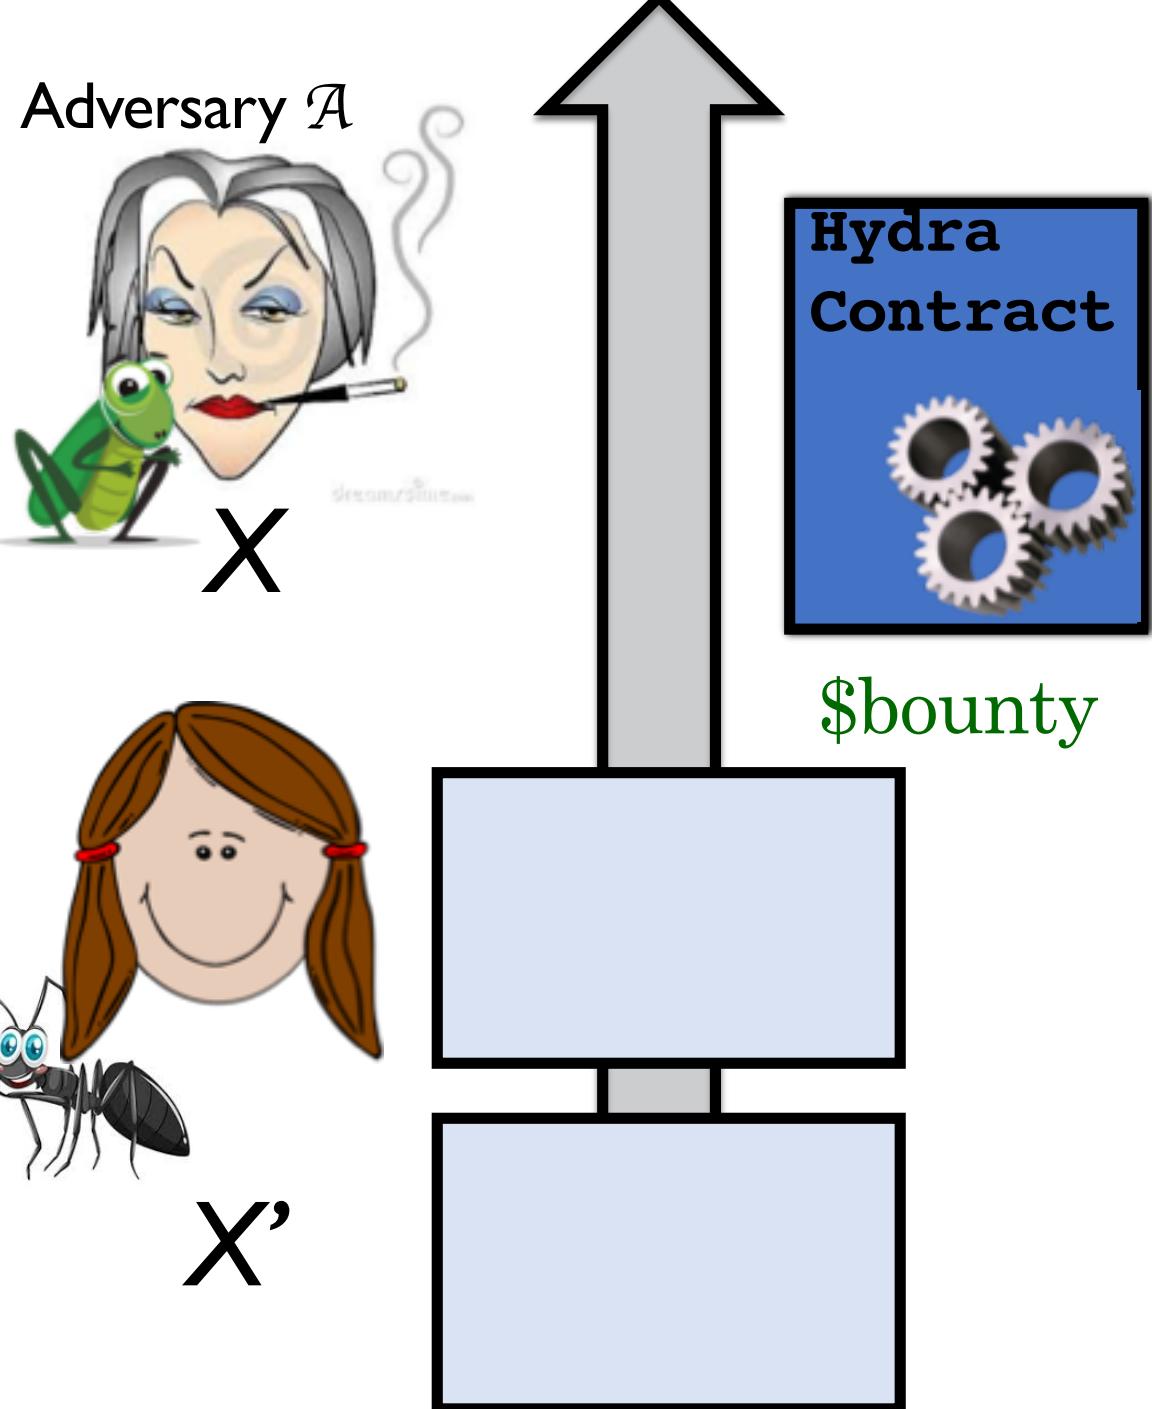
Submarine Commitments

Commit



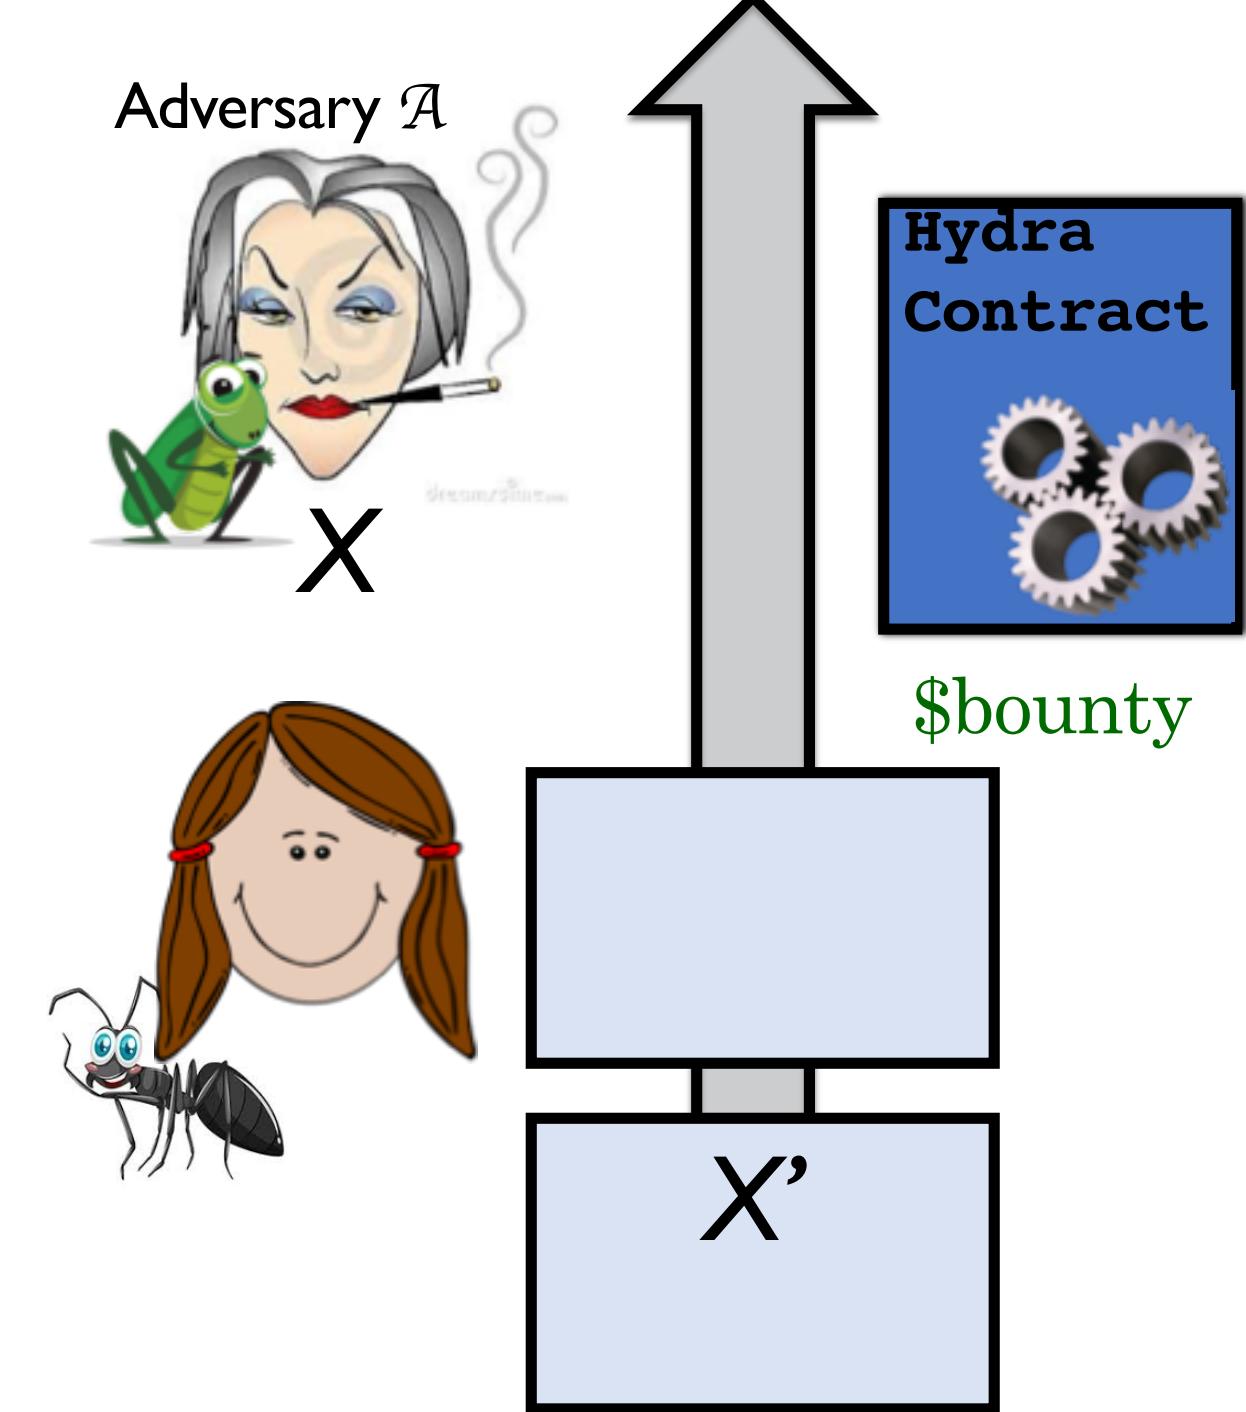
Bug withholding

- Suppose adversary \mathcal{A} discovers bug X
- \mathcal{A} should disclose fast to prevent honest user claiming \$bounty



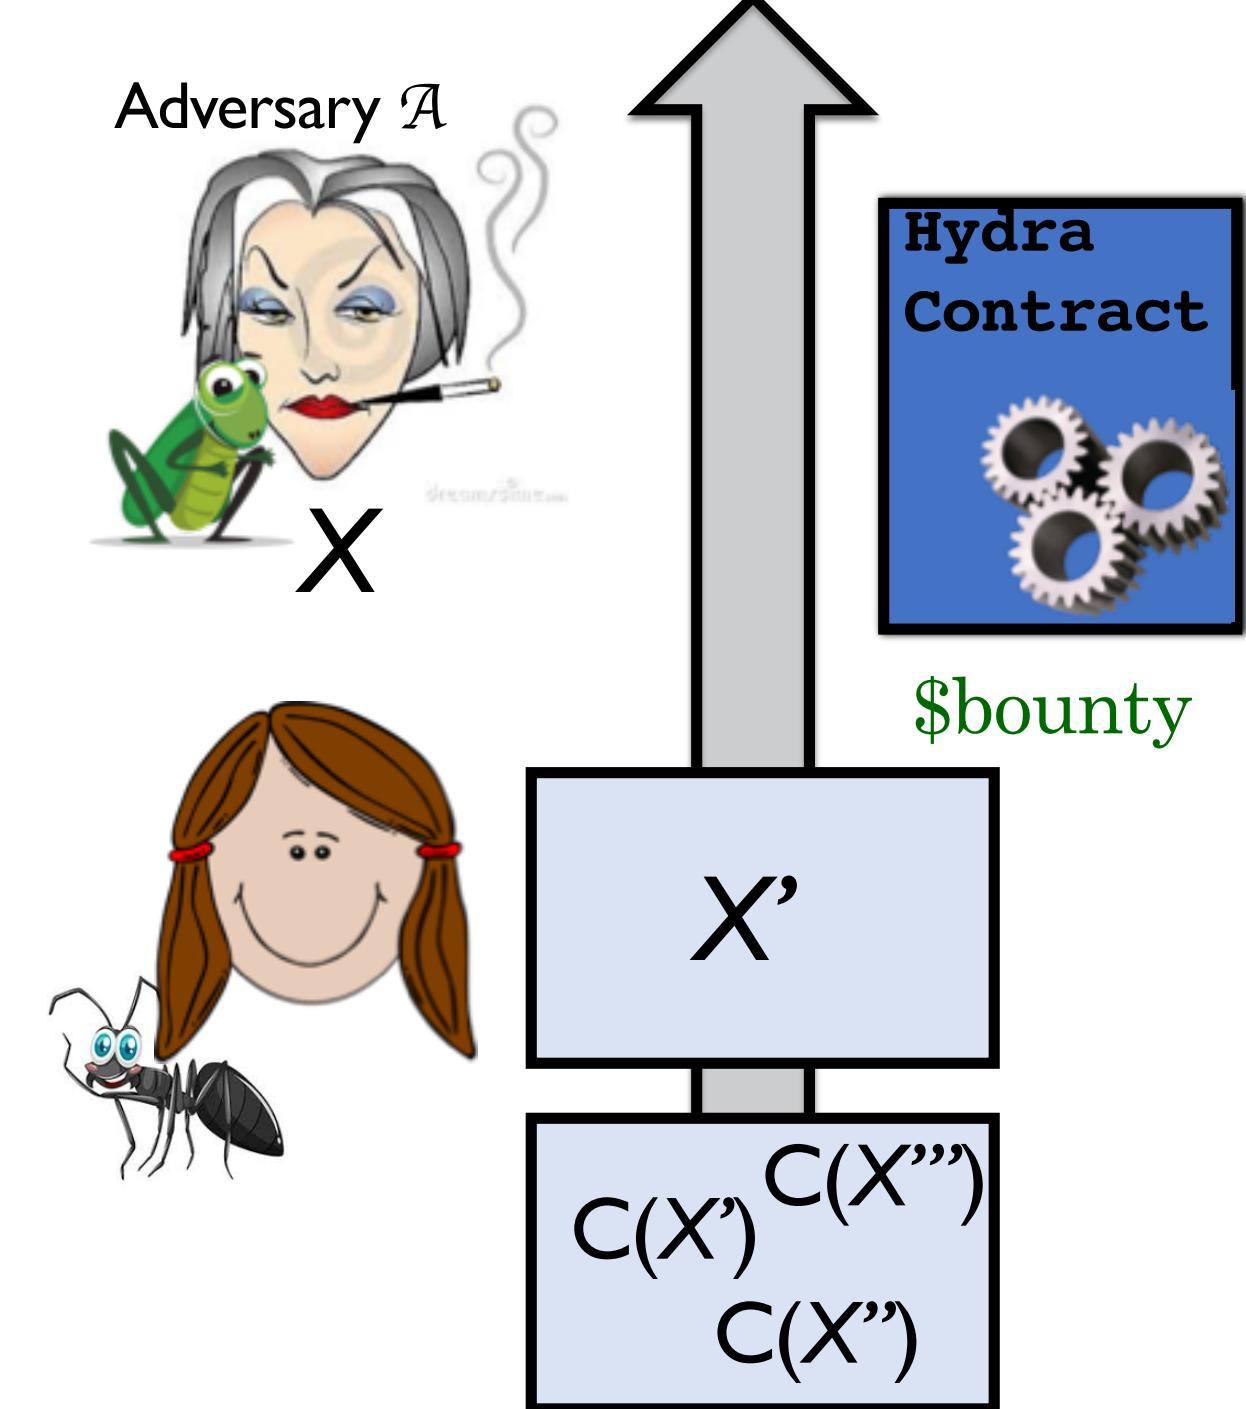
Bug withholding

- Unfortunately, blockchains are messy...
- \mathcal{A} can *front-run* honest user!
- So \mathcal{A} can *withhold* X and *keep looking for full exploit of**
• Ruins our whole bounty analysis!
 - No immediate incentive to disclose compromise of individual heads!



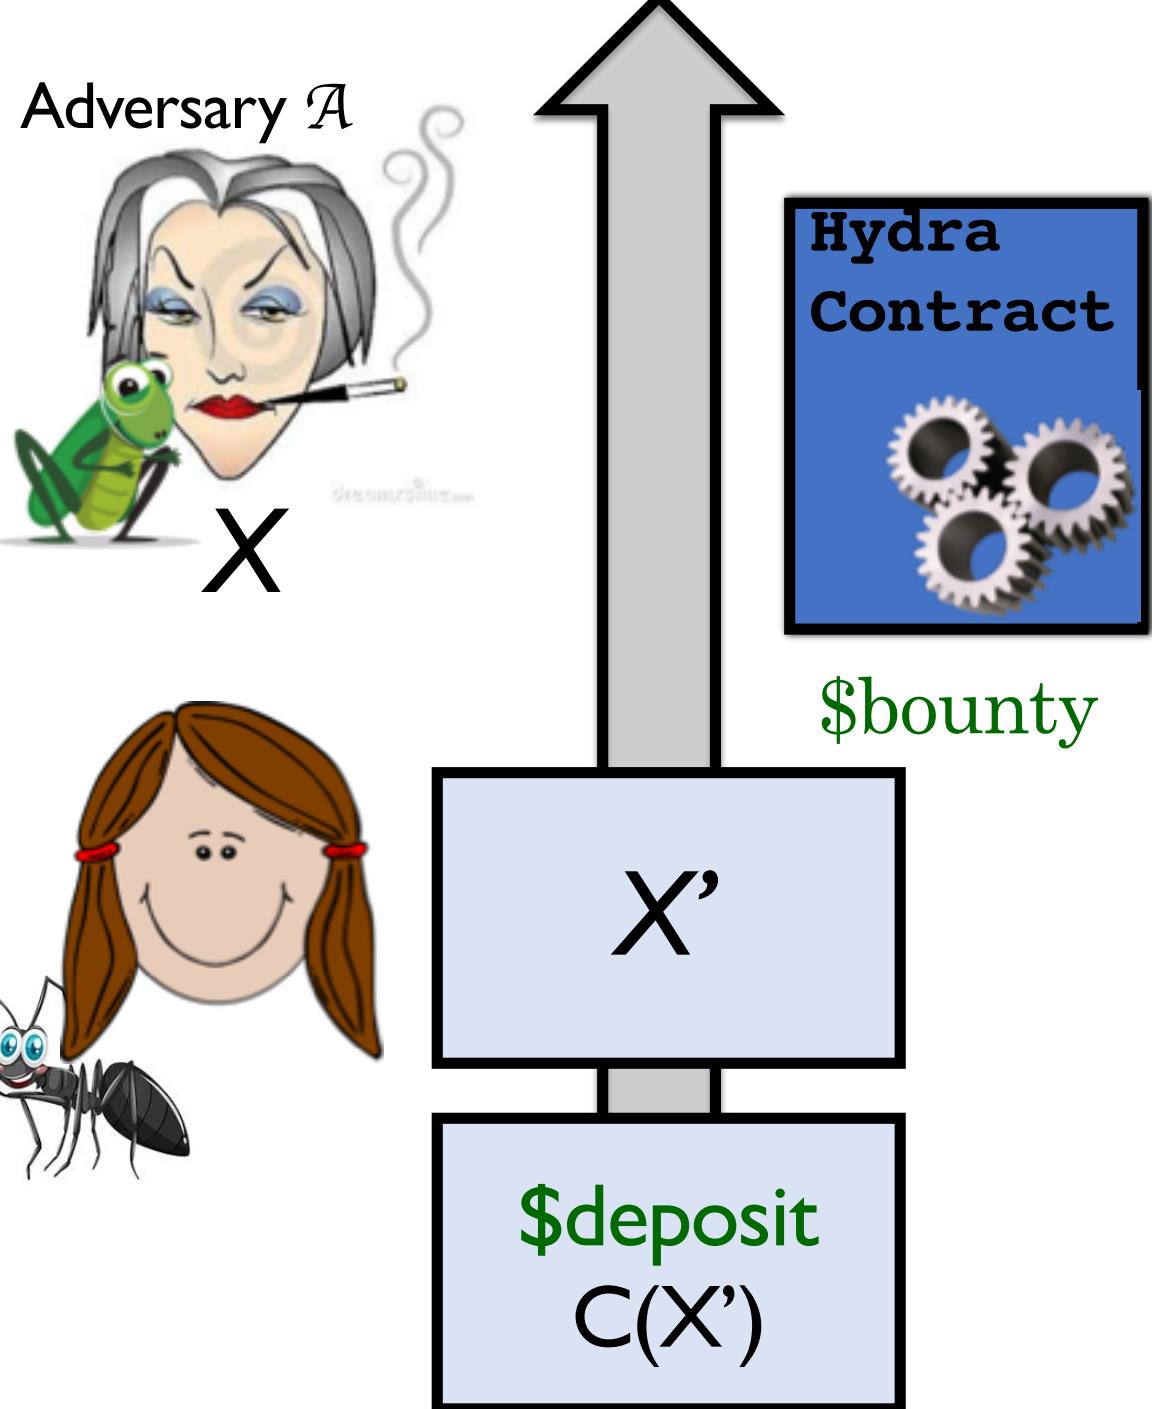
Solution?

- Idea 1: Must commit in block $t-1$ to reveal claim in block t
+ Lots of cover traffic
- Problem: \mathcal{A} commits in every round and front-runs reveal!



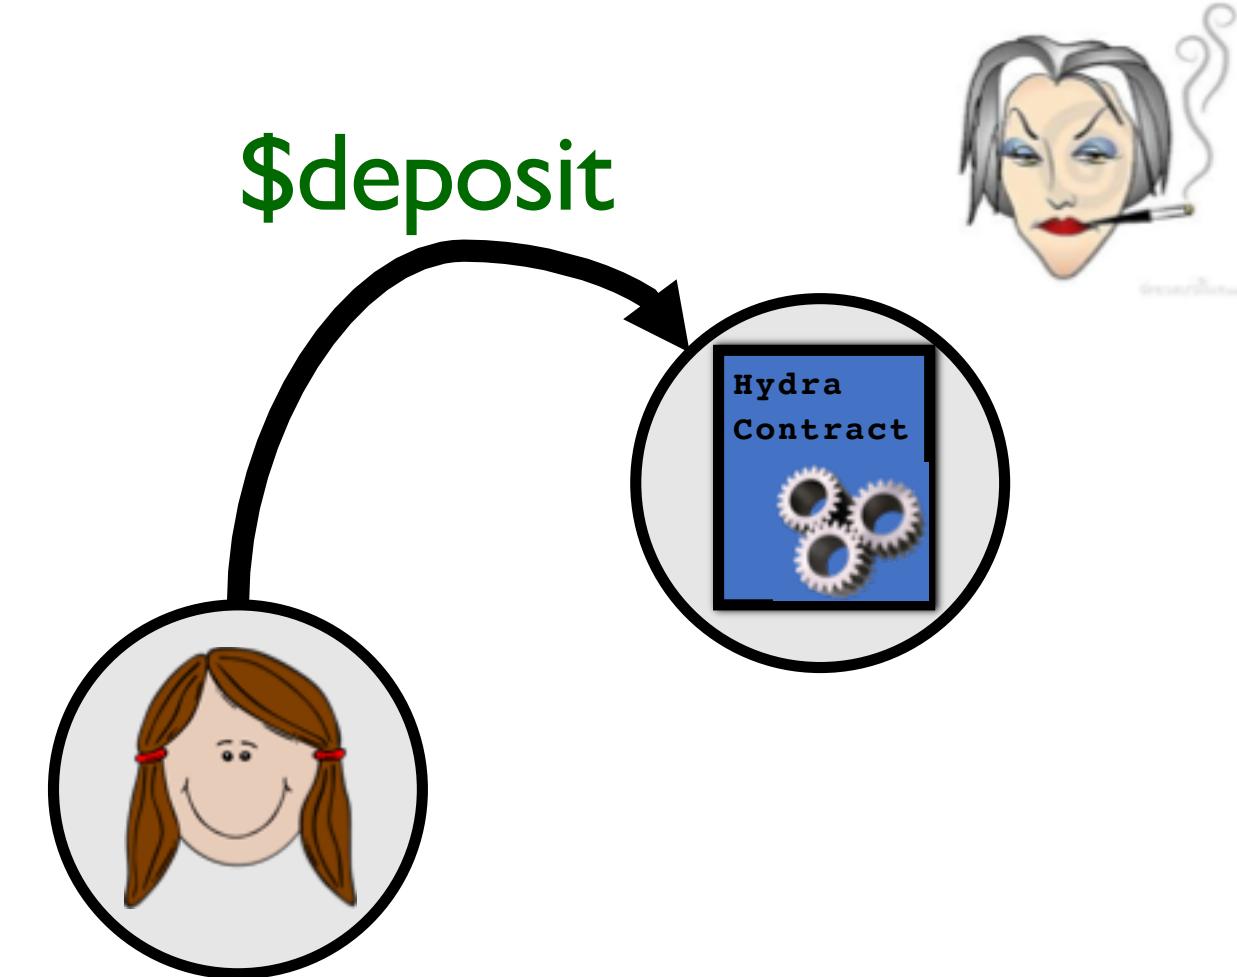
Solution?

- Idea 2: Must commit $\$deposit$ in block $t-1$ to reveal claim in block t



Solution?

- Idea 2: Must commit **\$deposit** in block $t-1$ to reveal claim in block t
- Problem: **\$deposit sent to Hydra Contract is publicly visible**
 - So \mathcal{A} can front-run commit!

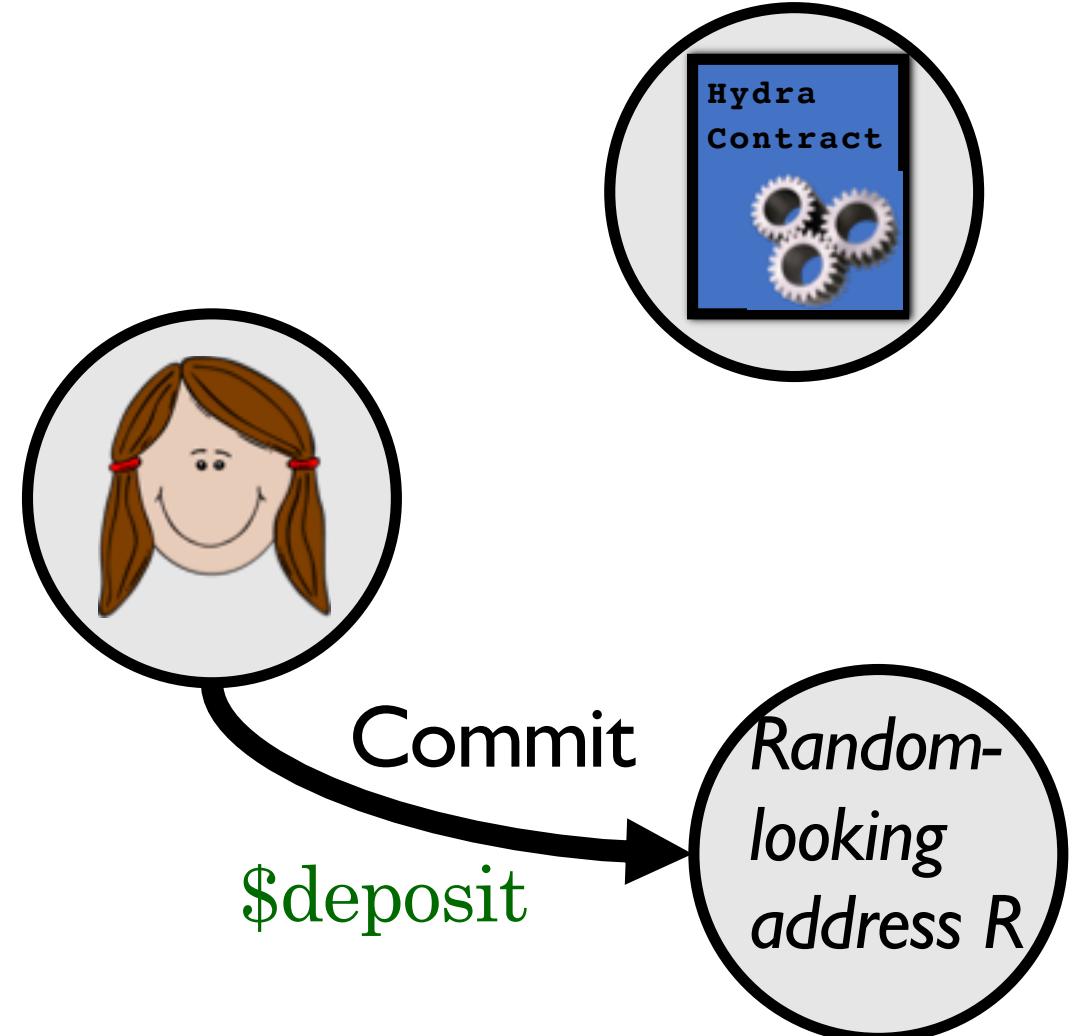


In general, if \mathcal{A} can observe honest users' behavior, she can front-run them!

Solution: Submarine Commitment



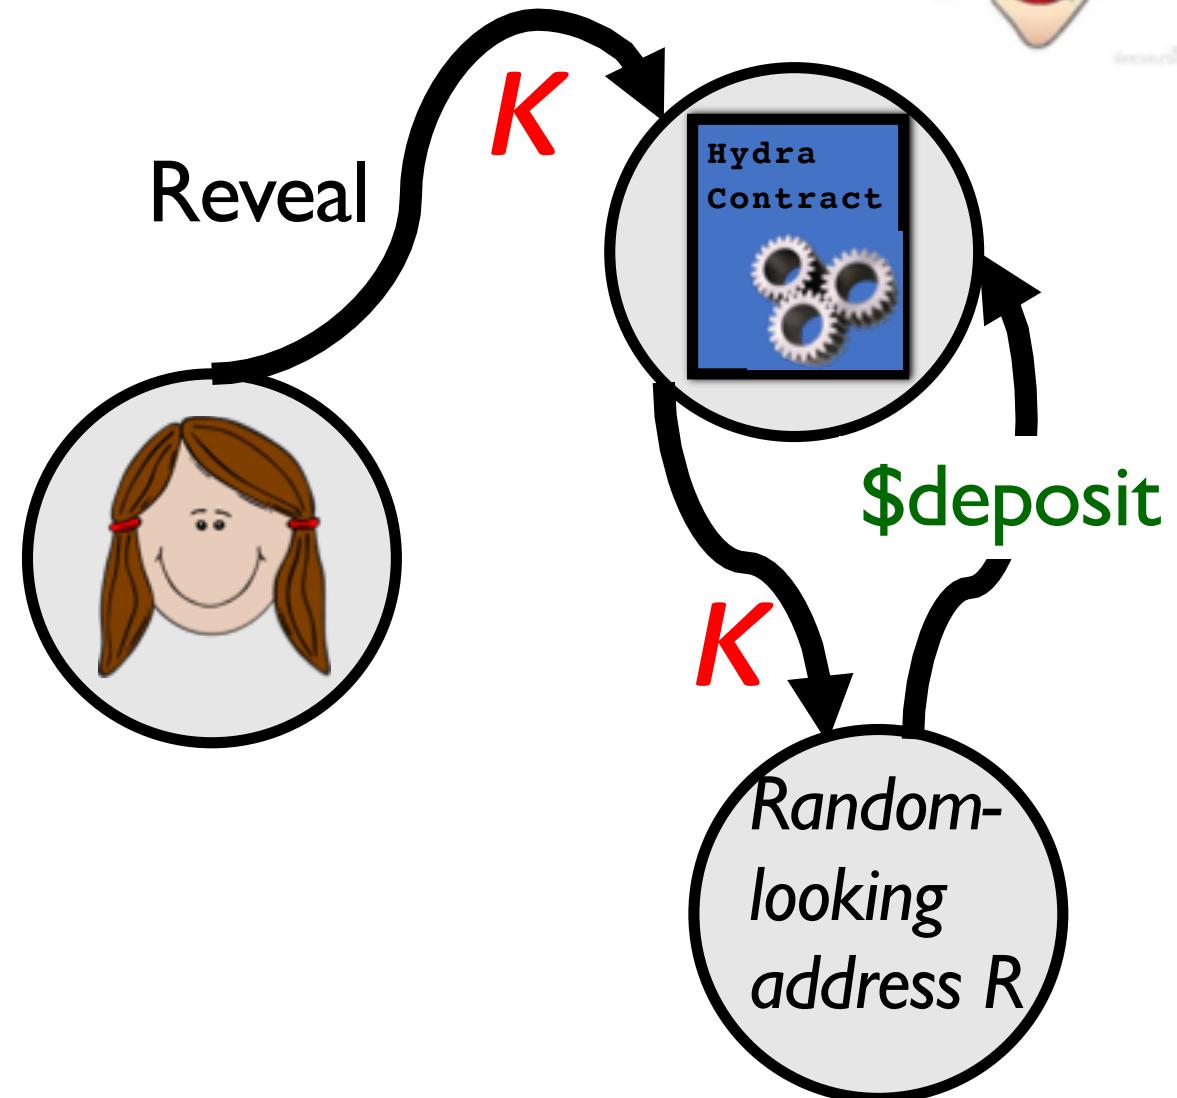
- **Commit** sends **\$deposit** to random address
- People send money to fresh addresses all the time!
- So **Commit** looks like ordinary traffic...
 - No visible association with Hydra Contract



Solution: Submarine Commitment



- But actually, R is specially constructed
- Only HydraContract can recover money from R , with key K
- **Reveal** sends key K
- Key K allows fund recovery by HydraContract
- Thus we can:
 - Commit $\$deposit$ stealthily and
 - Prevent front-running!



Submarine Commitments

- Security analysis a bit involved:
 - New, strong adversarial model introduced for blockchains

```
 $\mathcal{F}_{\text{withhold}}$  with  $\mathcal{P} = \{P_0, P_1, \dots, P_m\}$ ,  $(\delta, \rho)$ -adversary  $\mathcal{A}$ , blocksize  $s$ , target height  $n$ 
```

Init: $\mathcal{B} \leftarrow \emptyset$, $\mathcal{B}.\text{Height} \leftarrow 0$, $\text{MaxHeight} \leftarrow 0$, $\text{Mempool} \leftarrow \emptyset$

On receive ("post", τ) from P_i : // P_i submits tx
assert $\text{ValidTx}(\tau; \mathcal{B}, \text{Mempool})$
 $\text{tag}(\tau) \leftarrow (\mathcal{B}.\text{Height}, P_i)$ // Label tx with current chain height and sender
 $\text{Mempool} \leftarrow \text{Mempool} \cup \tau$
send Mempool to \mathcal{A}

On receive ("add block", B) from \mathcal{A} : // \mathcal{A} extends blockchain
if $\mathcal{B}.\text{Height} = n$ then
 output \mathcal{B} ; halt // To complete chain, \mathcal{A} adds arbitrary $n + 1^{\text{th}}$ block
assert $(|B| = s) \wedge (B \subseteq \text{Mempool})$
assert $\exists \tau \in \text{Mempool} - B$ s.t. $(\text{tag}(\tau) = (h, P_0)) \wedge (h \leq \mathcal{B}.\text{Height} - \delta)$
 // Ensure delay at most δ for P_0 's transactions
 $\mathcal{B}.\text{Height} \leftarrow \mathcal{B}.\text{Height} + 1$
 $B_{\mathcal{B}.\text{Height}} \leftarrow B$ // Add new block to chain
 $\text{Mempool} \leftarrow \text{Mempool} - B$ // Remove processed txs from Mempool
 $\text{MaxHeight} \leftarrow \max(\mathcal{B}.\text{Height}, \text{MaxHeight})$
send \mathcal{B} to P_0

On receive ("rewind", r) from \mathcal{A} // \mathcal{A} rewinds by r blocks
assert $\text{MaxHeight} - (\mathcal{B}.\text{Height} - r) \leq \rho$
 // Ensure that \mathcal{A} rewinds by no more than ρ
 $\text{Mempool} \leftarrow \text{Mempool} \cup \{B_i\}_{i \in [\mathcal{B}.\text{Height} - r + 1, \mathcal{B}.\text{Height}]}$
 // Return rewound transactions to Mempool
 $\mathcal{B}.\text{Height} \leftarrow \mathcal{B}.\text{Height} - r$

Figure 2: Ideal functionality $\mathcal{F}_{\text{withhold}}$ for (δ, ρ) -adversary \mathcal{A}

Submarine Commitments

- Security analysis a bit involved:
 - New, strong adversarial model introduced for blockchains
 - Standard cryptographic modeling of adversaries... but with money

```
Experiment  $\text{Exp}_{\mathcal{A}}^{\text{bnyrace}}(n', \delta, \rho, s; \Delta, \text{\$deposit}, \text{\$bounty})$   
Init:  $n \leftarrow n' - \Delta$ ,  $\text{\$cost} \leftarrow 0$ ,  $\text{commblock}_{P^*} \leftarrow \mathbb{S}[1, n]$   
 $\mathcal{A}^{\{B \leftarrow \mathcal{F}_{\text{withhold}}(\{P_0 = P^*, P_1\}, n, \delta, \rho, s)\}}$  //  $\mathcal{A}$  interacts with  $\mathcal{F}_{\text{withhold}}$   
for  $i = 1$  to  $n$   
    if ("commit",  $\text{\$deposit}) \in B_i$  then  
         $\text{\$cost} \leftarrow \text{\$cost} + \text{\$deposit}$  // Every commit costs  $\text{\$deposit}$   
    if  $(\exists (1 \leq i \leq \text{commblock}_{P^*} \wedge i \leq j \leq \min(i + \Delta, n))$  s.t.  
         $\exists (\tau = \text{"commit"}) \in B_i$  s.t.  $\text{tag}(\tau) = (i, P_1) \wedge$   
         $\exists (\tau = \text{"reveal"}) \in B_j$  s.t.  $\text{tag}(\tau) = (j, P_1)$  then  
            output(TRUE,  $\text{\$payoff} := \text{\$bounty} - \text{\$cost}$ ) //  $\mathcal{A}$  wins  
    output(FALSE,  $\text{\$payoff} := -\text{\$cost}$ )
```

Figure 4: Adversarial game $\text{Exp}_{\mathcal{A}}^{\text{bnyrace}}$

Submarine Commitments

- We prove tight bounds on adversary's front-running ability
- E.g., to protect \$100,000 bounty with reasonable parameters in Ethereum, need $\$deposit = \278
- New, practical Ethereum implementation *not in paper*
 - We're implementing it...

The Hydra Project [alpha]

Hydra is a cutting-edge **Ethereum** contract development framework for:

- decentralized security and bug bounties
- rigorous cryptoeconomic security guarantees
- mitigating programmer and compiler error

[READ THE PAPER](#)

[TRY THE ALPHA](#)

[CHAT ON RIOT](#)

www.thehydra.io

Initiative for CryptoCurrencies and Contracts (IC3)

IC3

The Initiative For
CryptoCurrencies & Contracts

HOME ABOUT PEOPLE PARTNERS PROJECTS PUBLICATIONS BLOGS PRESS EVENTS

IC3: ADVANCING THE SCIENCE AND APPLICATIONS OF BLOCKCHAINS

Latest on Blog



[Paralysis Proofs: How to Prevent Your Bitcoin From Vanishing](#)

by Fan Zhang , Phil Daian , Iddo Bentov , and Ari Juels on Thursday January 18, 2018 at 09:30 AM

Suppose that N players share cryptocurrency using an M -of- N multisig scheme. If $N-M+1$ players disappear, the remaining ones have a problem: They've permanently lost their funds. In this blog, we propose a solution to this critical problem using the power of the trusted hardware.



[The Social Workings of Contract](#)

by Karen Levy on Wednesday January 17, 2018 at 01:00 PM

Guest blogger Prof. Karen Levy describes how contracts often include terms that are unenforceable, purposefully vague, or never meant to be enforced, how this helps set expectations, and what this means for smart contracts.

News & Events

 May 10-11, 2018

[IC3 Spring Retreat in NYC](#) ▾

IC3 faculty, students and industry members gather twice per year to discuss the major technical challenges and innovative solutions to widespread blockchain adoption.

 February 26, 2018 -March 2, 2018

[Financial Cryptography and Data Security 2018 and the 5th Workshop on Bitcoin and Blockchain Research](#) ▾

Prof. Sarah Meiklejohn is co-Program Chair for FC18 and Prof. Ittay Eyal is co-Program Chair for the 5th Workshop on Bitcoin and Blockchain Research.

www.initc3.org