

CCS 2023, Copenhagen, Denmark

Interactive Proofs For Differentially Private Counting

Ari Biswas ¹

Graham Cormode ^{1,2}

¹ University Of Warwick

² Meta AI

Motivating Problem: Counting



The local government of Wolvercote, a small village in Oxfordshire want to know if they should change public healthcare policy.

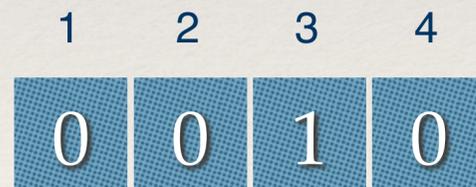
In order to gauge public opinion they conduct a survey over the population of the village.

Survey Question

Each resident is asked to vote for a single policy only



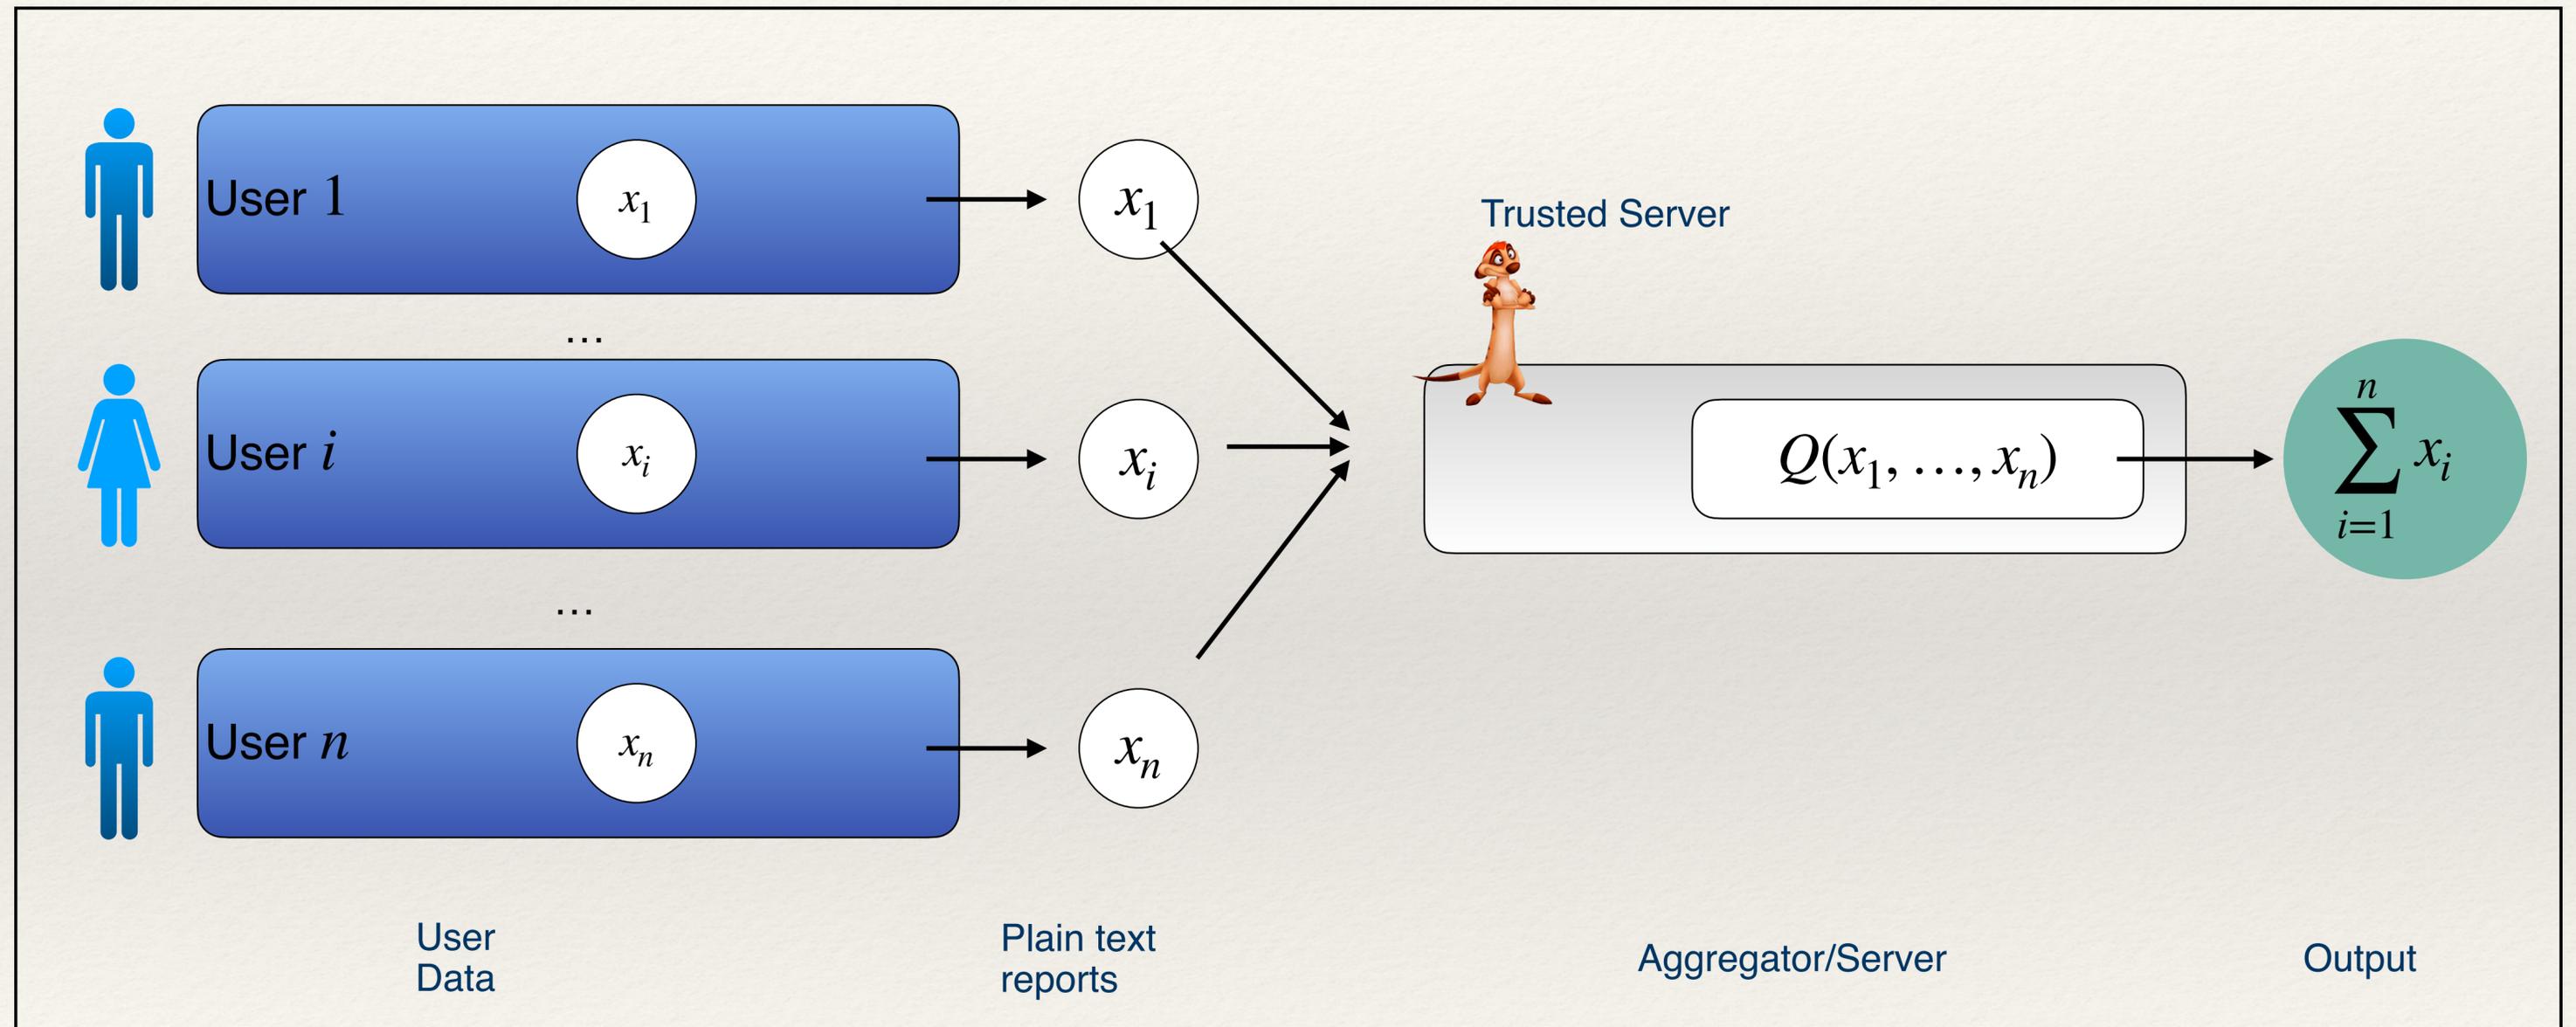
i



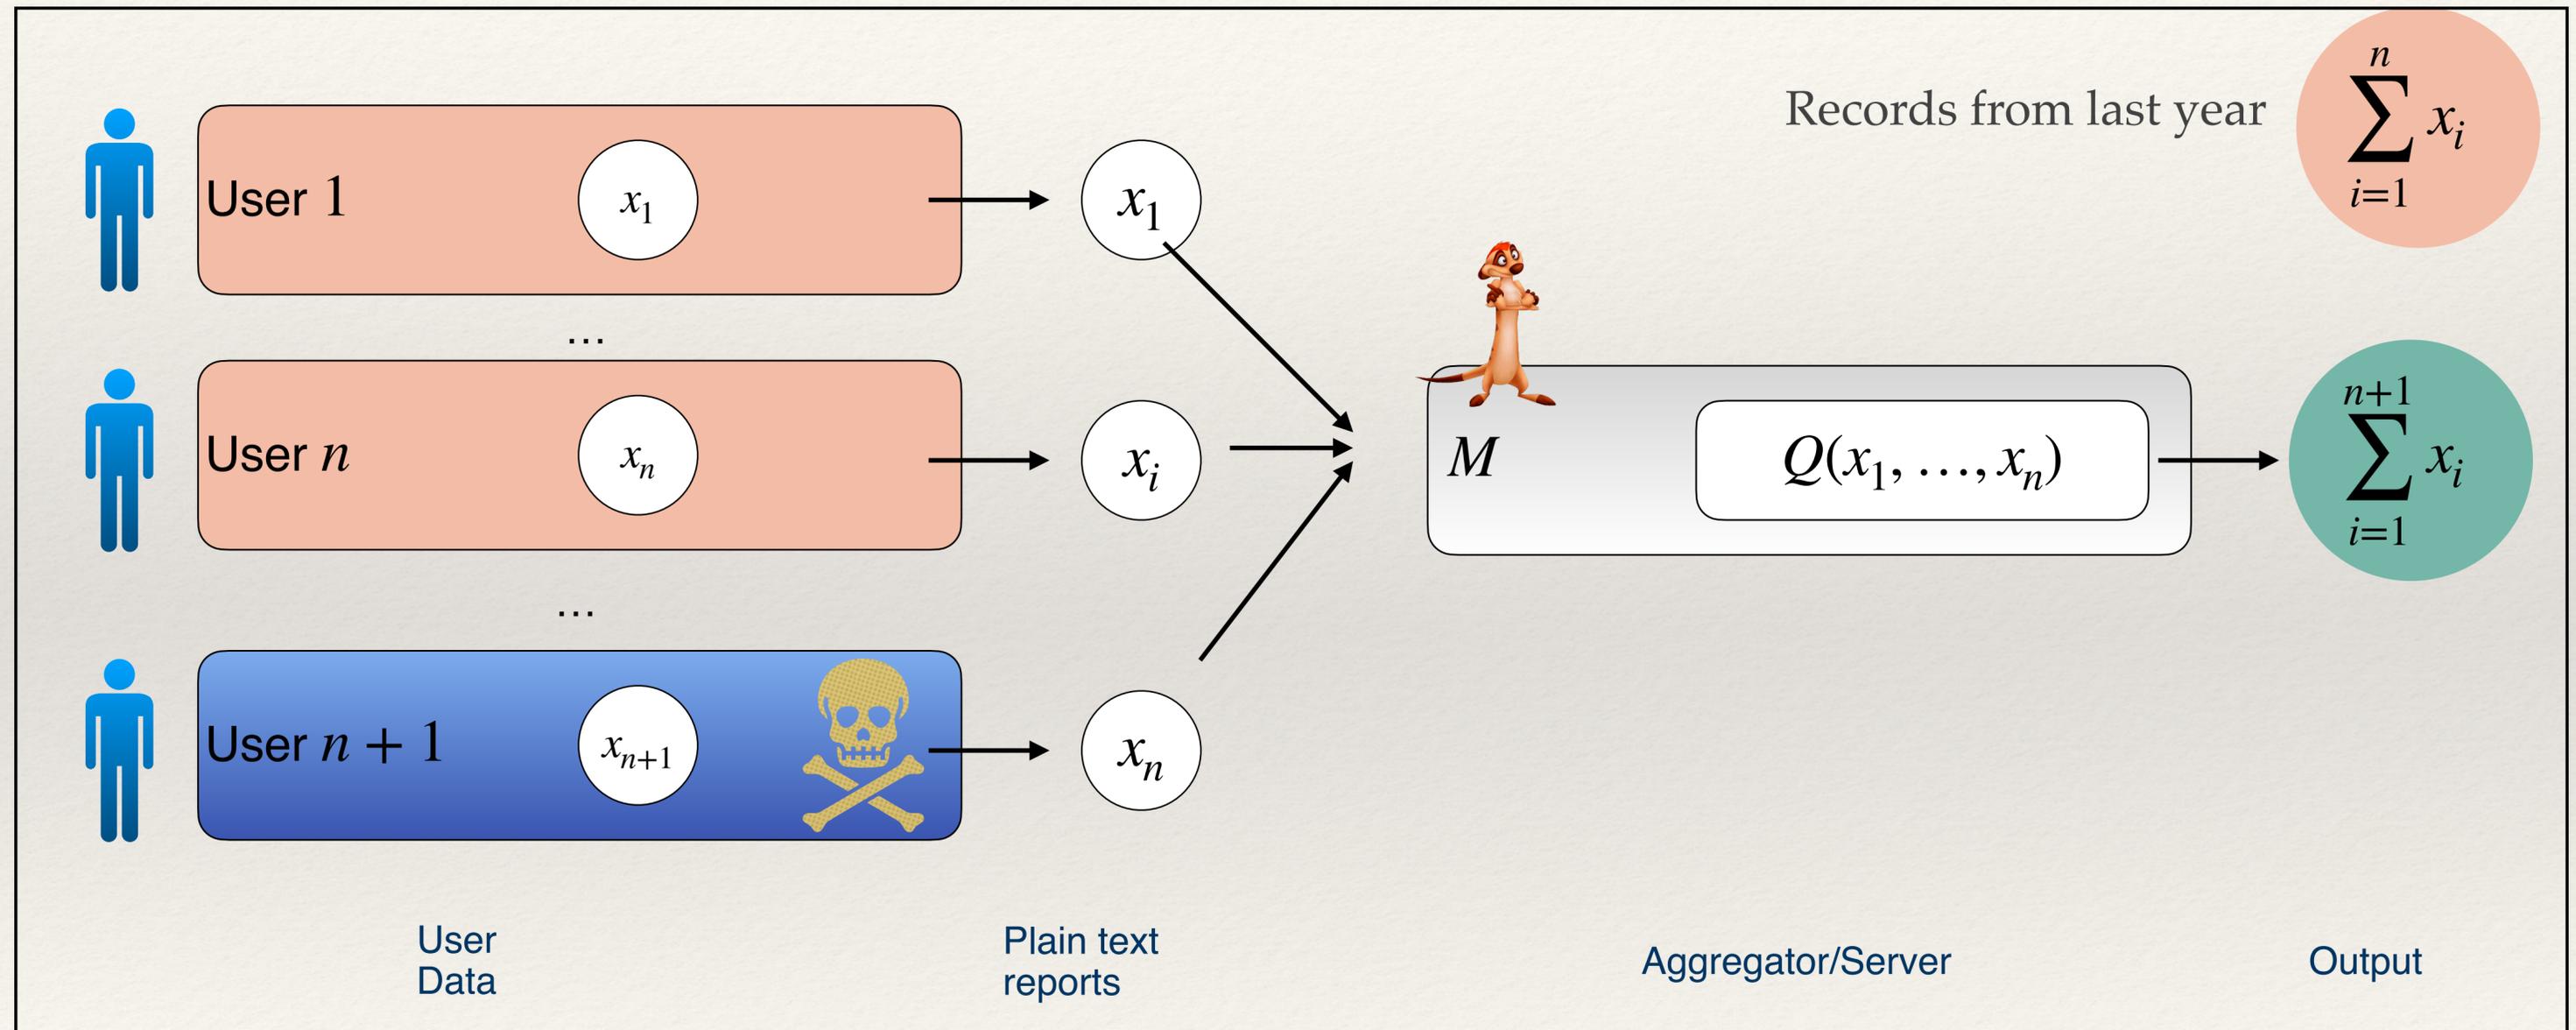
x_i

- 1: Mandatory Vaccination
- 2: Increase Pay Towards Healthcare workers
- 3: Decrease Taxes Towards Healthcare
- 4: Increase Taxes Towards Healthcare

An Ideal Solution



A New Person Moves in

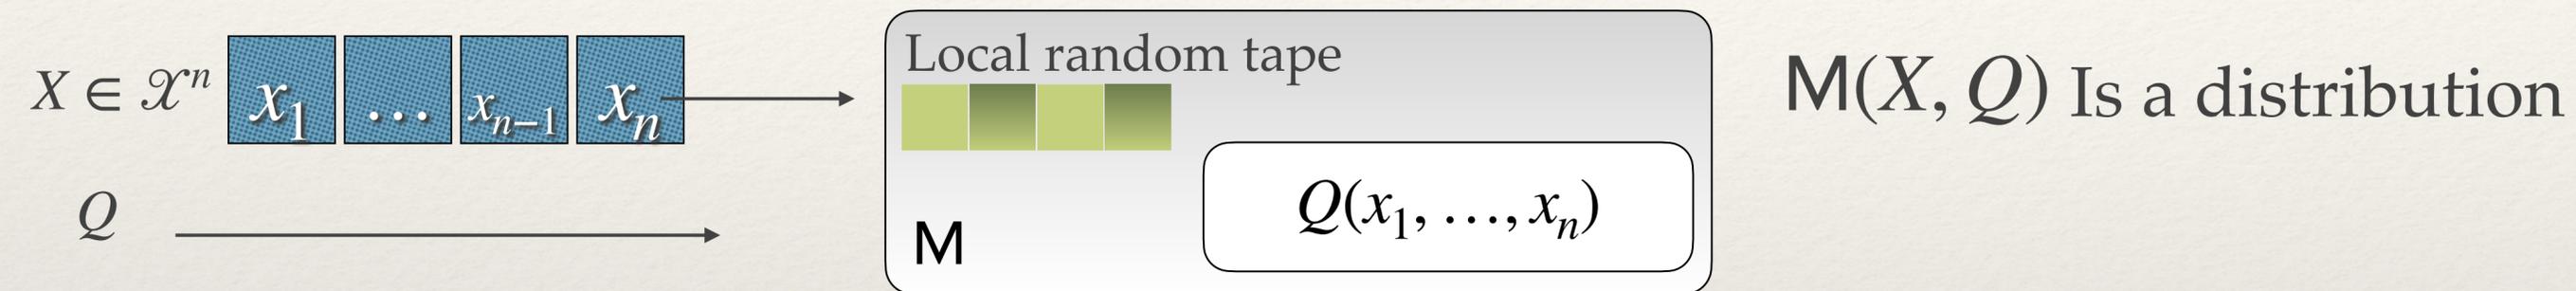


Randomness To The Rescue

- ❖ In this scenario, there is no deterministic algorithm that can help prevent information leakage about the n 'th user's value.
- ❖ Thus we **MUST** use randomness to obfuscate information about the new user.

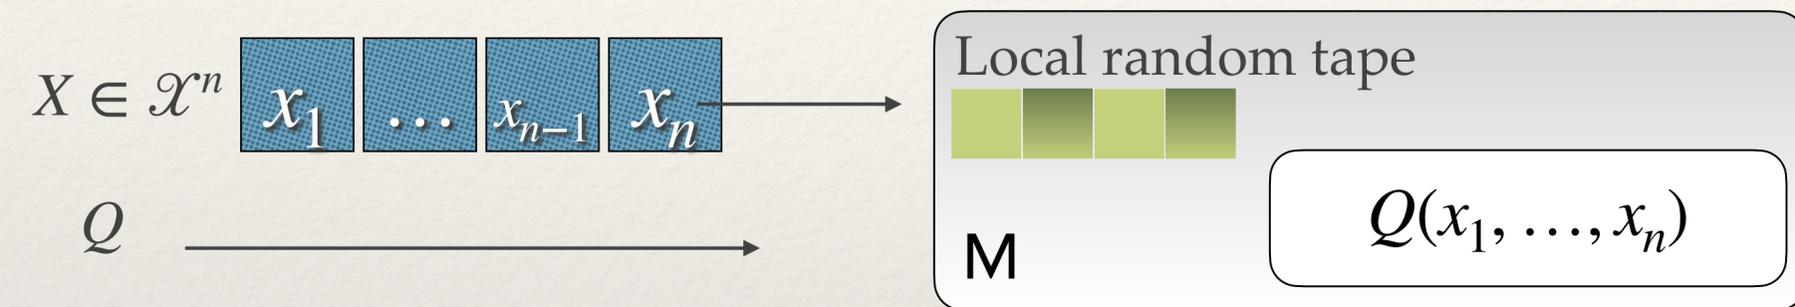
(ϵ, δ) -Differential Privacy (DP)

An algorithm $M : \mathcal{X}^n \times \mathcal{Q} \rightarrow \mathcal{Y}$ for releasing $Q(X)$

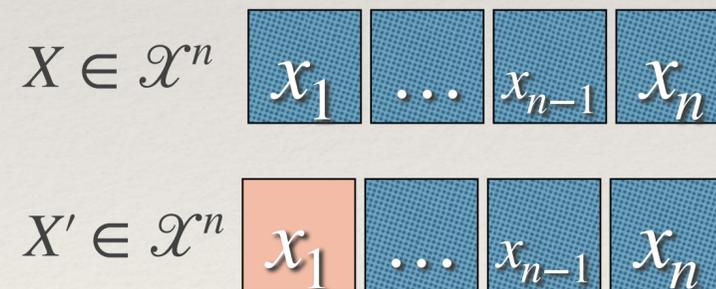


(ϵ, δ) -Differential Privacy (DP)

An algorithm $M : \mathcal{X}^n \times \mathcal{Q} \rightarrow \mathcal{Y}$ for releasing $Q(X)$

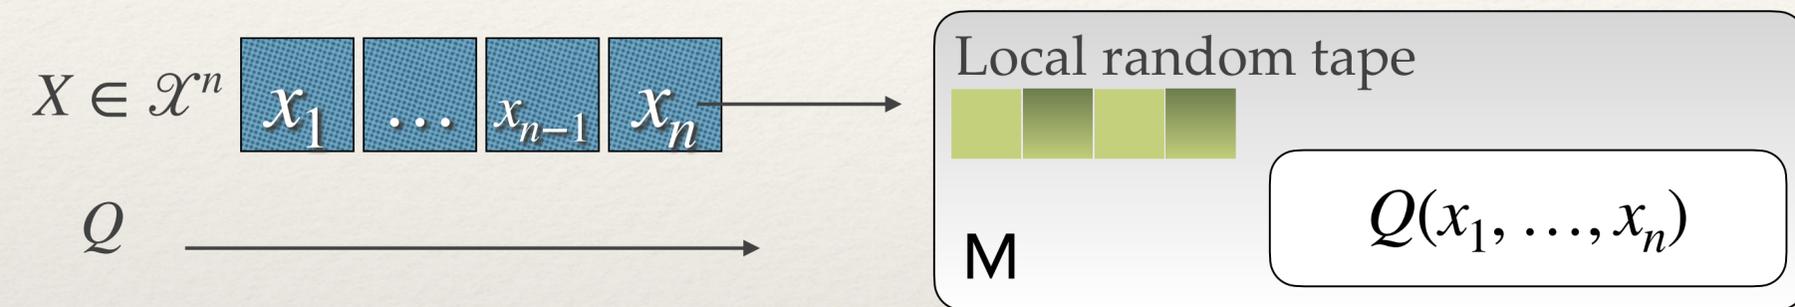


For **any** neighbouring datasets $X \sim X'$ i.e datasets that differ by just one element



(ϵ, δ) -Differential Privacy (DP)

An algorithm $M : \mathcal{X}^n \times \mathcal{Q} \rightarrow \mathcal{Y}$ for releasing $Q(X)$



M is said to be (ϵ, δ) -Differentially Private if for any subset $T \subseteq \mathcal{Y}$

For **any** neighbouring datasets $X \sim X'$ i.e. datasets that differ by just one element



$$\Pr_{y \leftarrow M(X, Q)} [y \in T] \leq e^\epsilon \Pr_{y \leftarrow M(X', Q)} [y \in T] + \delta$$

Utility Of A DP Algorithm

An algorithm $M : \mathcal{X}^n \times Q \rightarrow \mathcal{Y}$ for releasing a DP version of $y = Q(X)$ where (\mathcal{Y}, d) is a metric space we define utility

$$\text{Error} = \mathbb{E}_{\hat{y} \leftarrow M(X, Q)} [d(\hat{y}, y)]$$

Candidate metrics

$$\mathcal{Y} = \mathbb{R}^d \quad d(x, y) = \|x - y\|_1$$

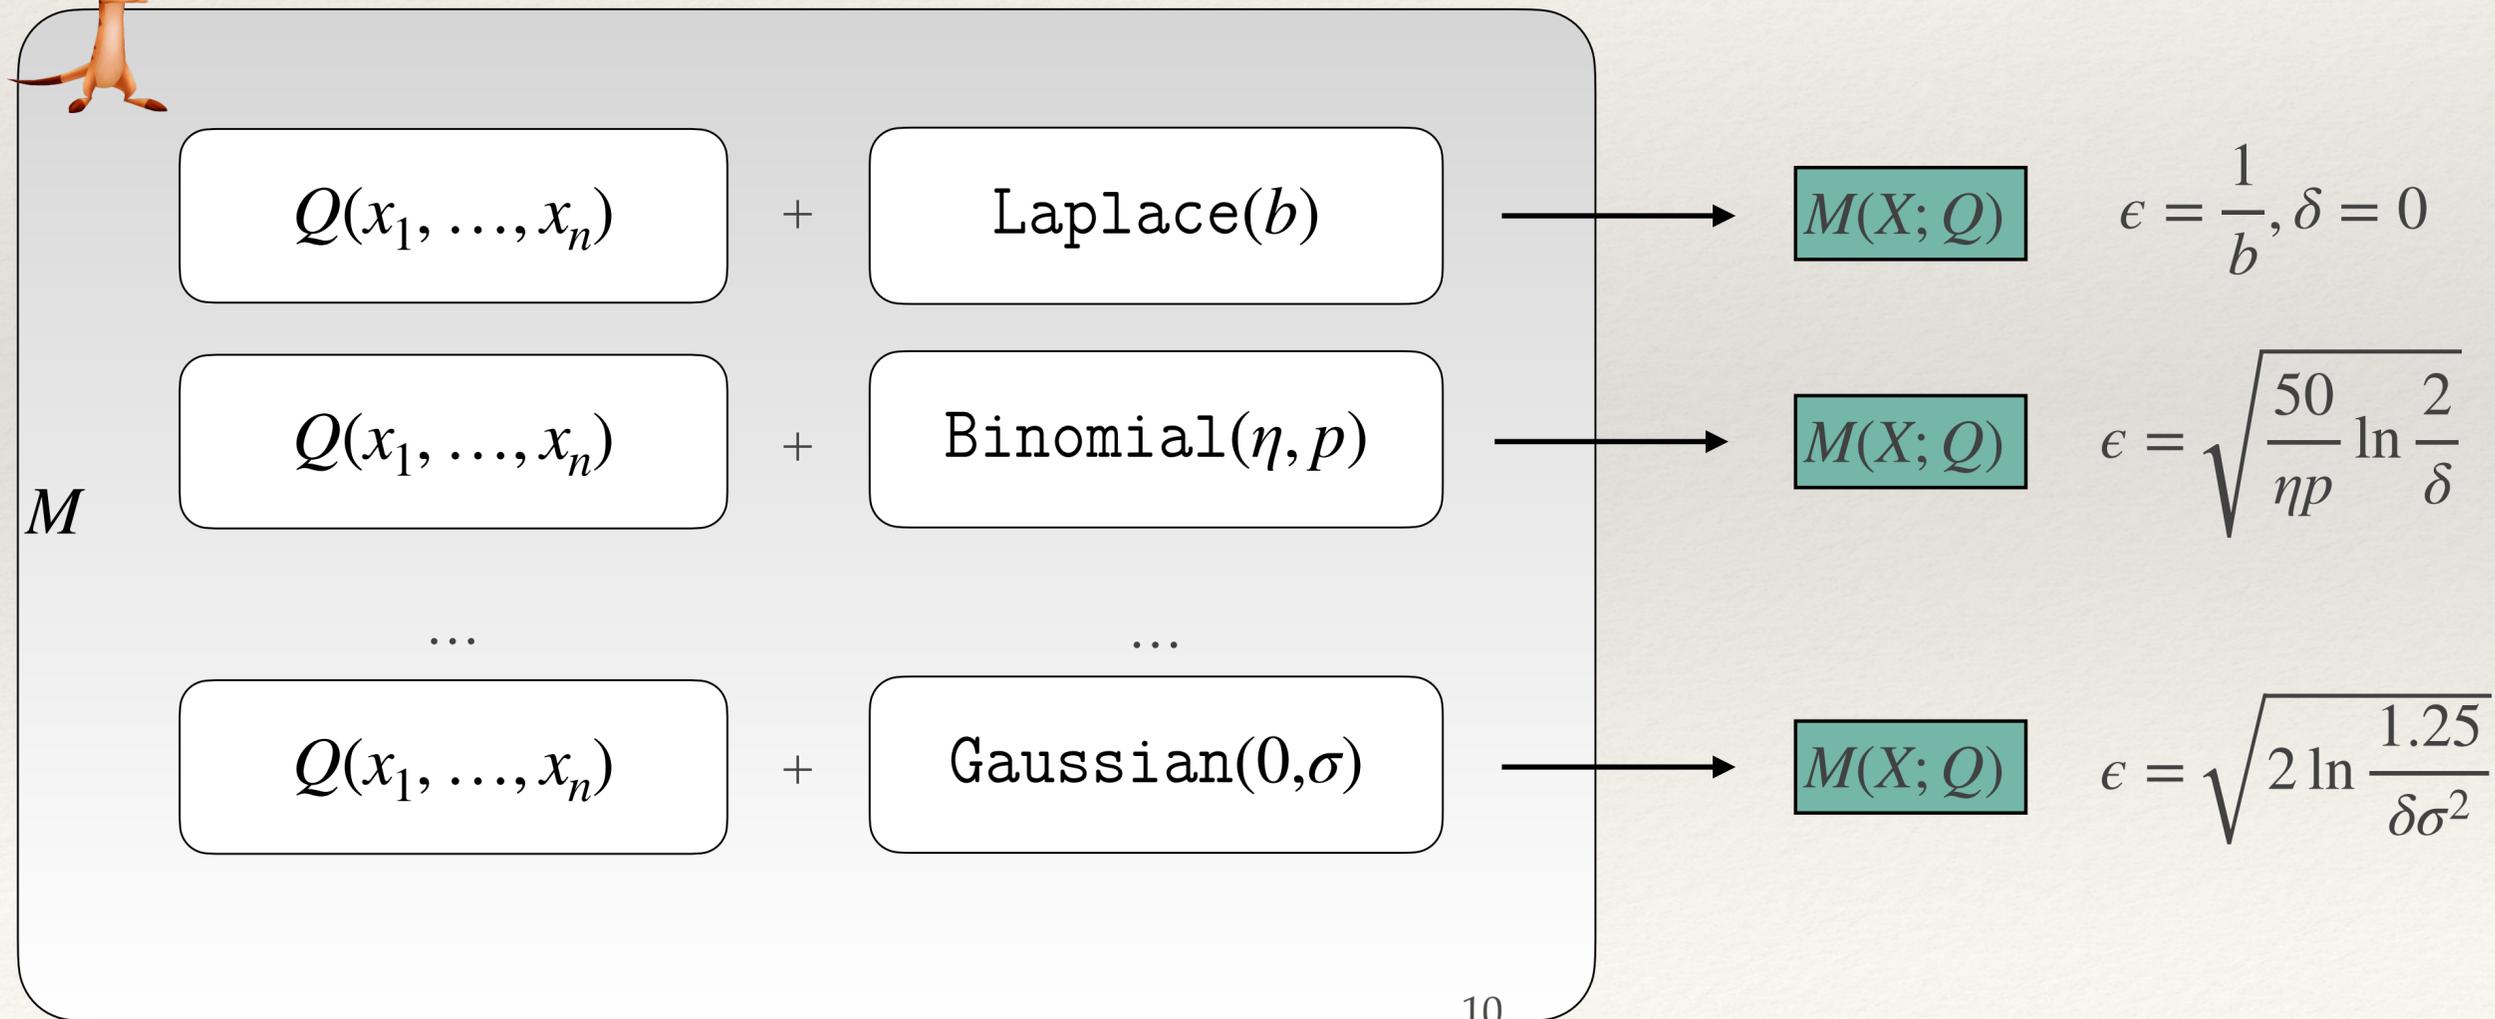
$$\mathcal{Y} = \mathbb{Z}_q^d \quad d(x, y) = \|x - y\|_2$$

$$d(x, y) = \|x - y\|_\infty$$

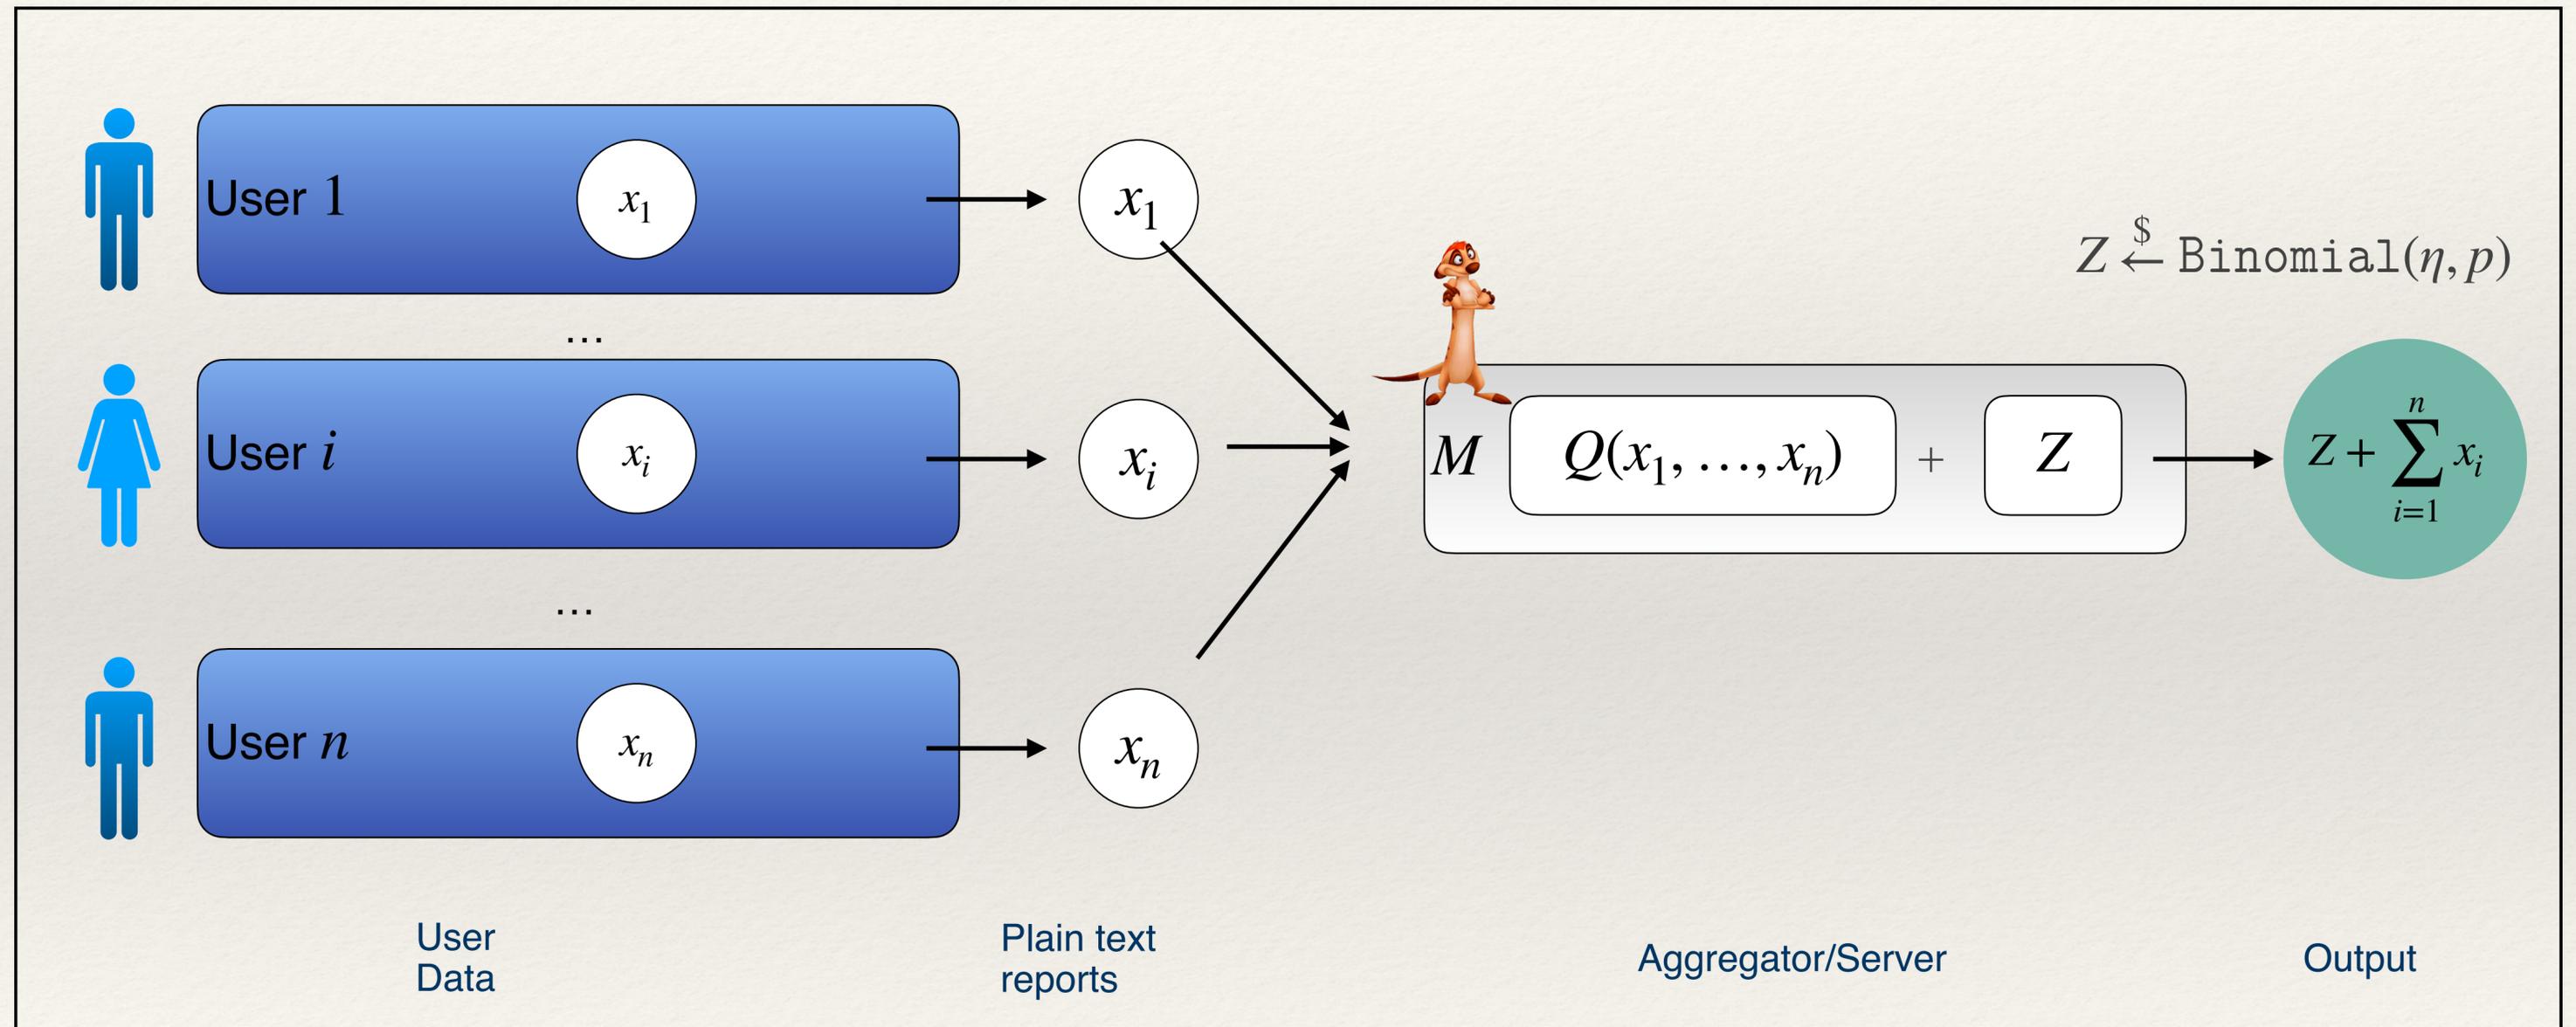
If we draw a sample from $M(X, Q)$, then on average how far is that sample from the true untampered answer.

DP Counting

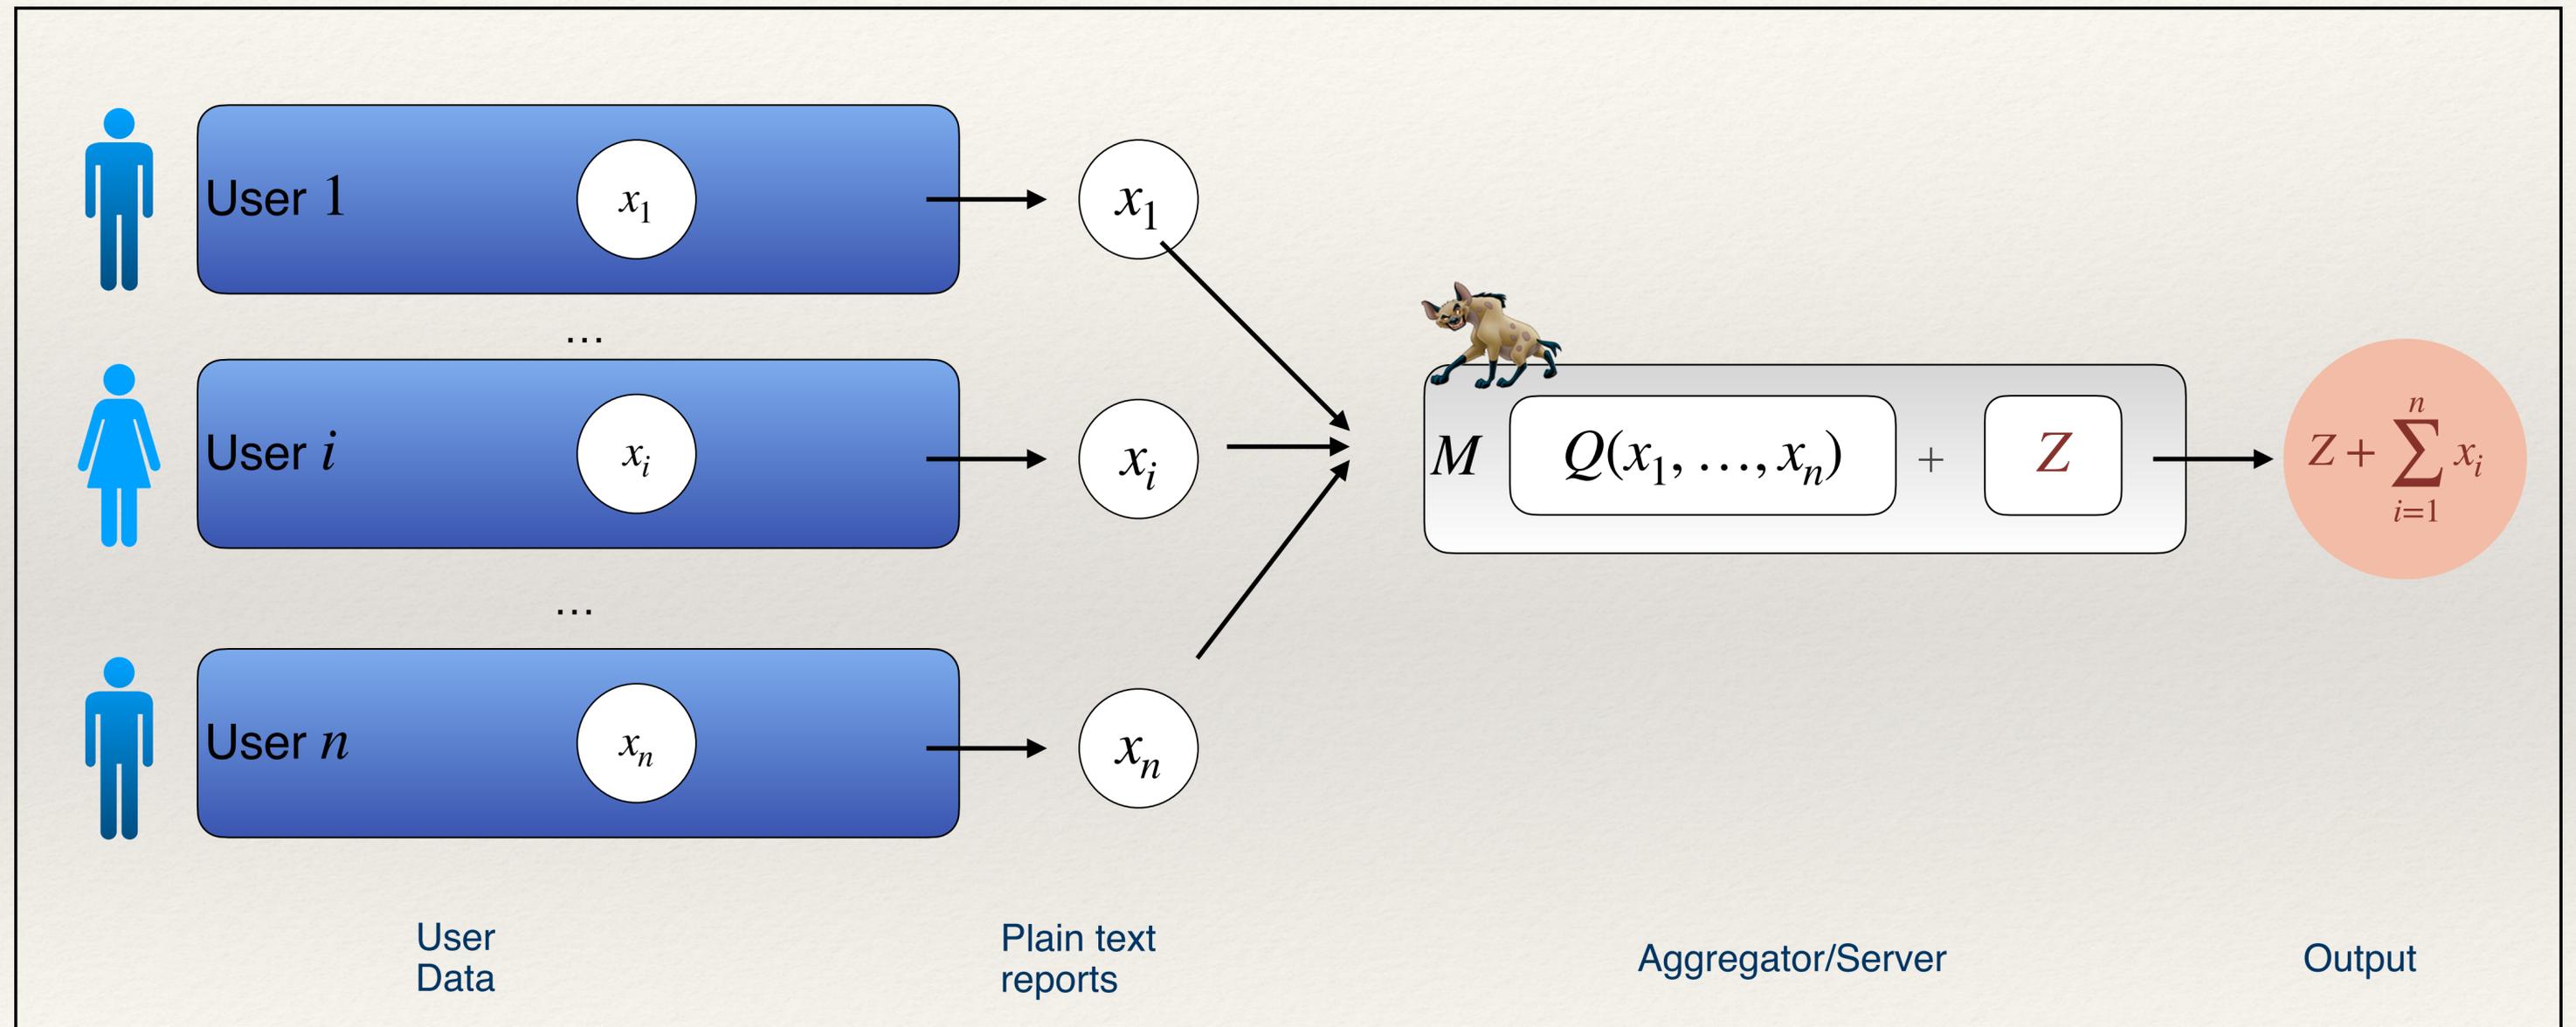
$$Q(x_1, \dots, x_n) = \sum_{i=1}^n x_i$$



Back To Our Ideal World



What If We Cannot Trust The Server ?



What Do We Want

- ❖ We want outputs to be differentially private
- ❖ However, we also want the output to be reliable i.e, by that we mean any error in the output must come as a result of DP noise and that only.

Need Some Crypto

Commitments

Two stage interactive protocol between a Committer and a Receiver



Committer

Commit Phase

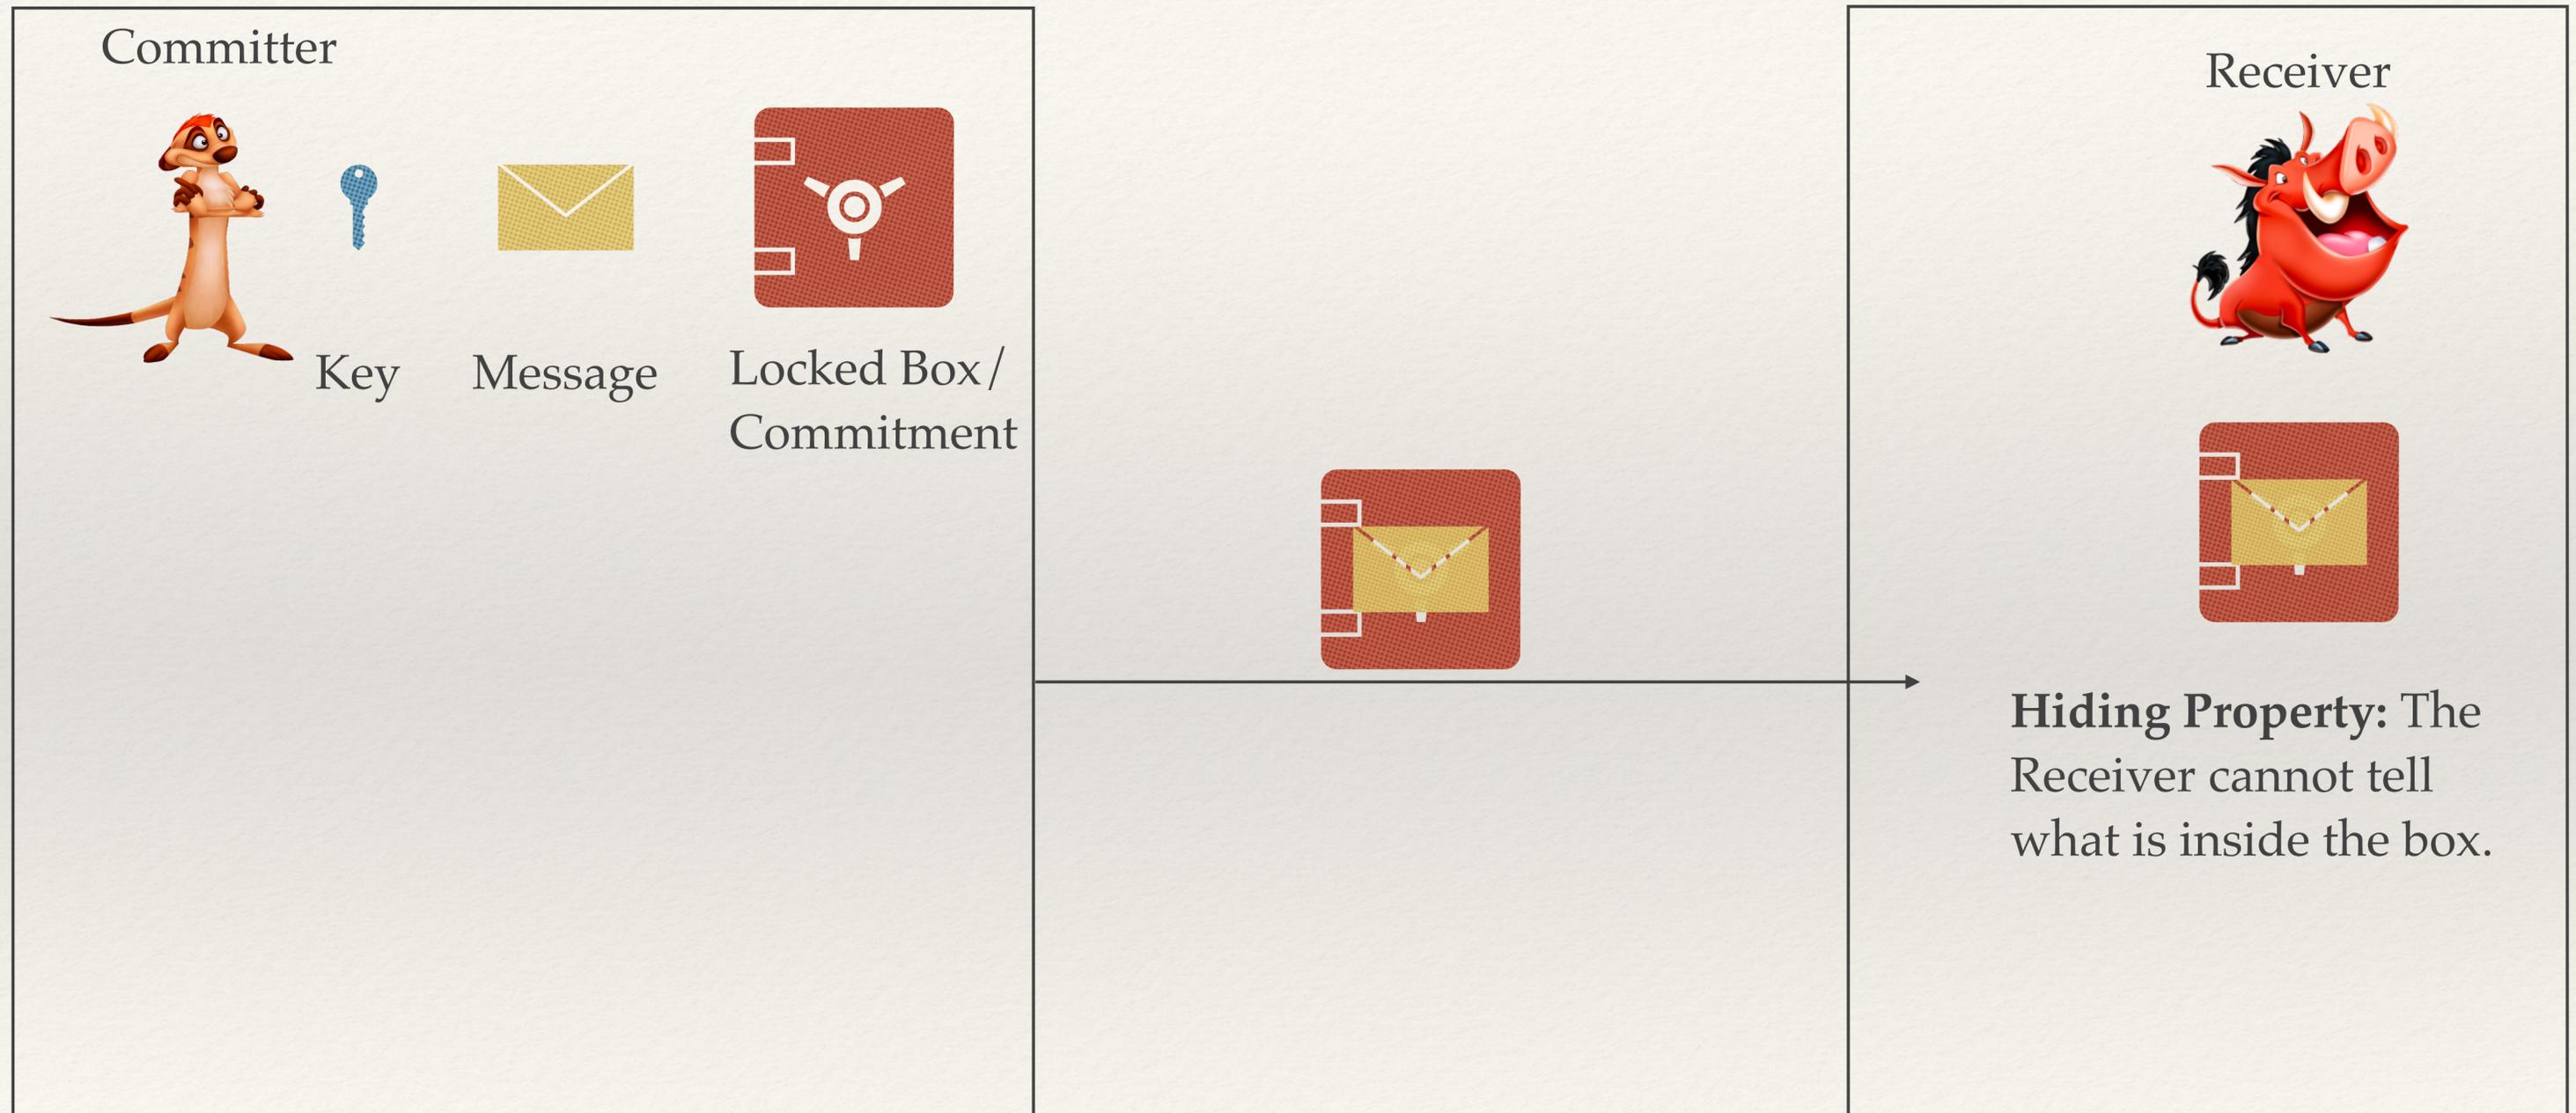


Receiver

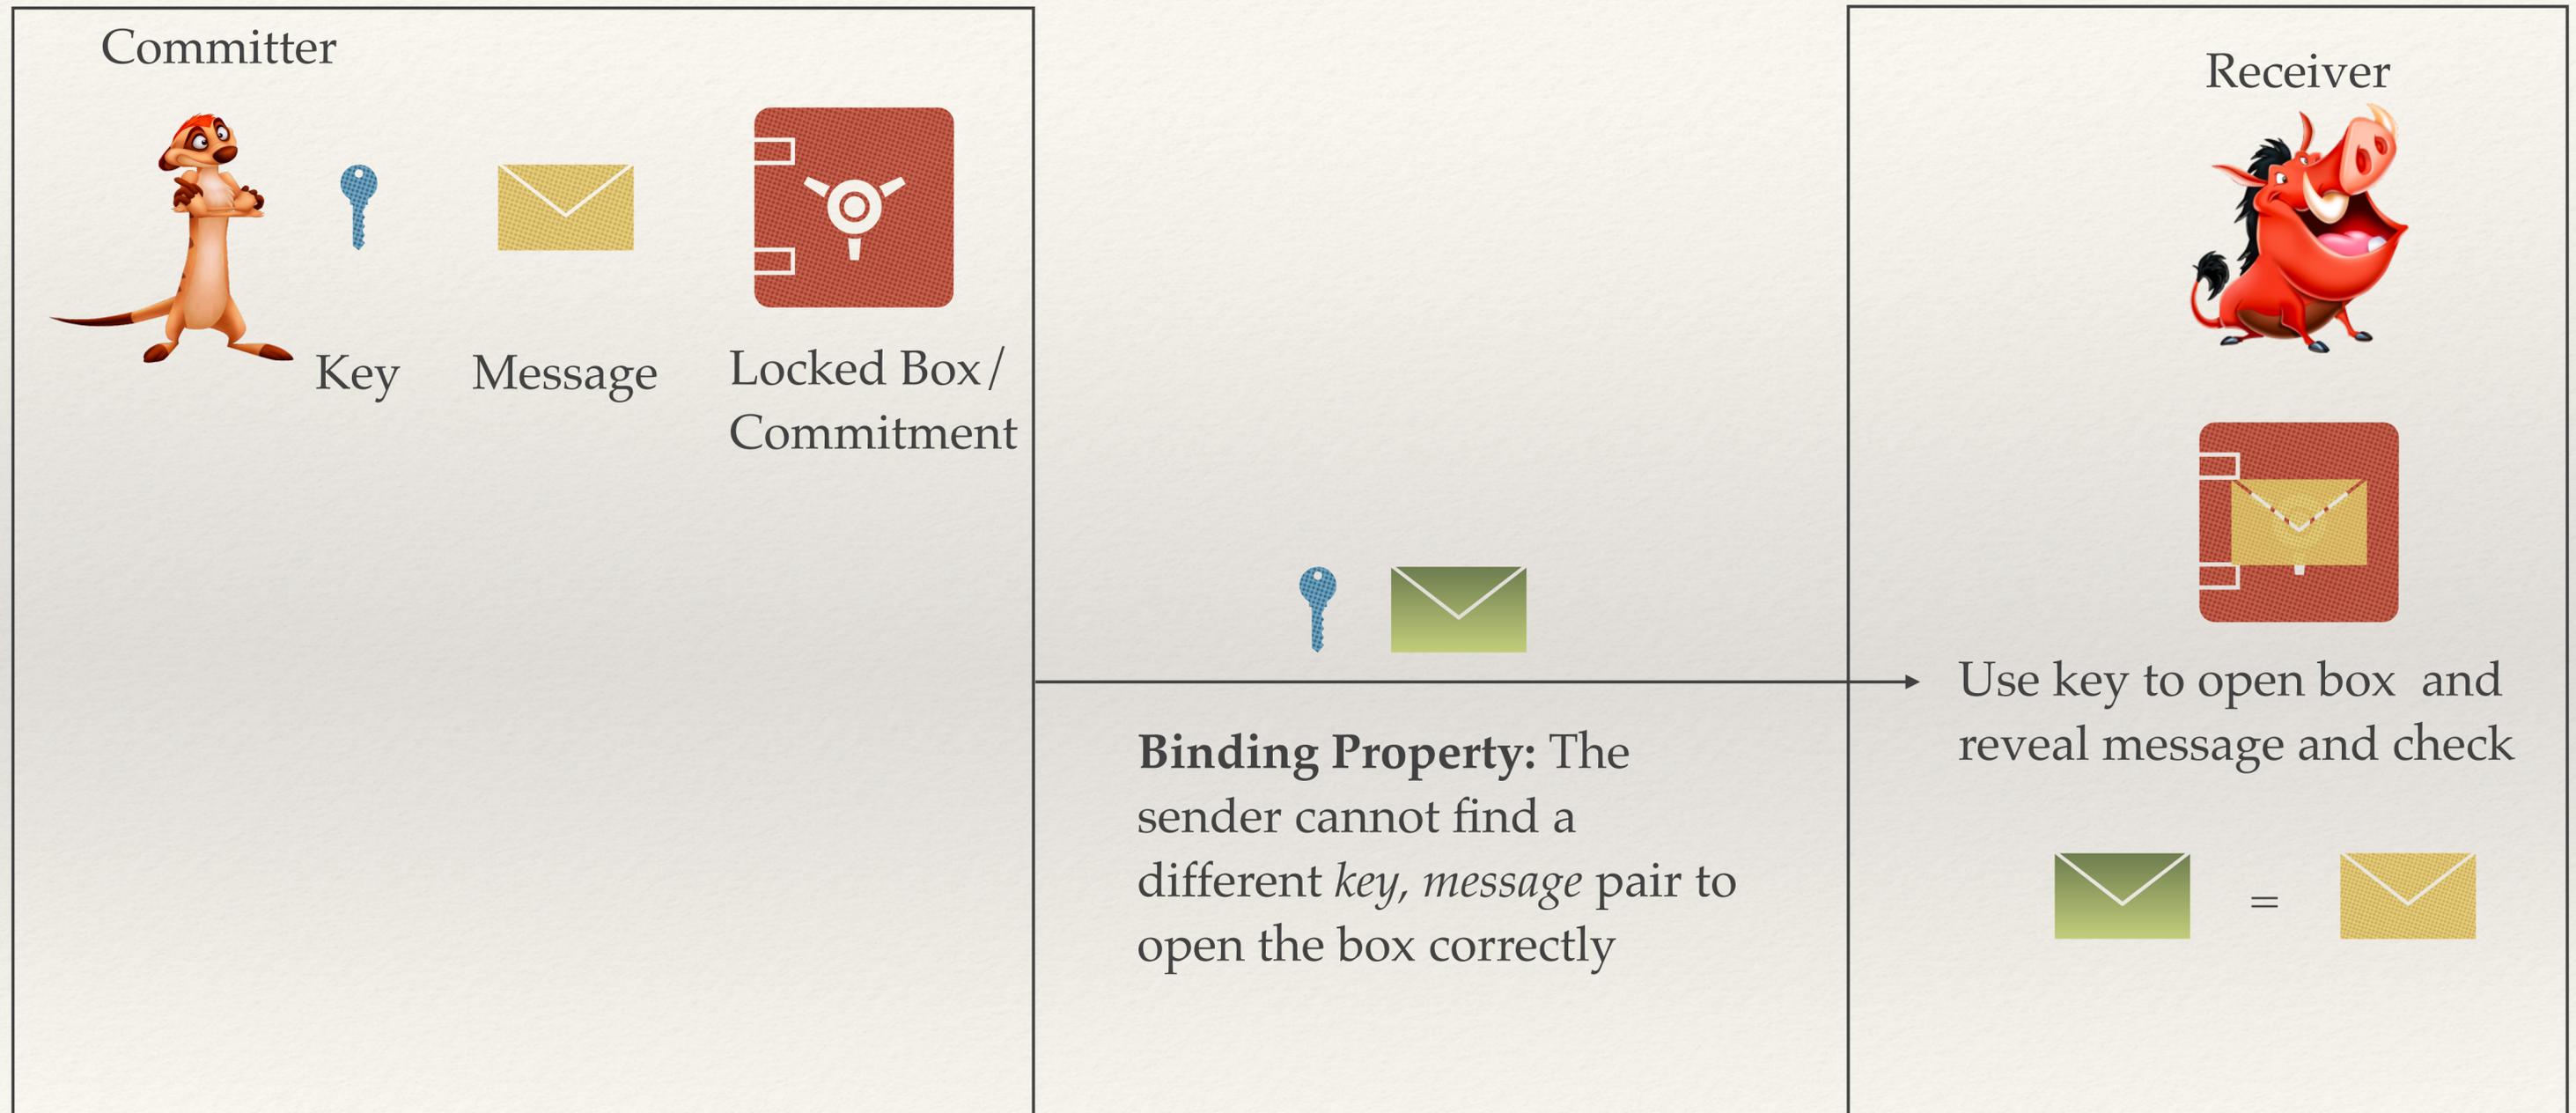
Reveal Phase



Commit Phase



Reveal Phase



Homomorphic Commitments



Key



Message



Locked Box/
Commitment



+



=



Key



Message



Locked Box/
Commitment



+



=



+

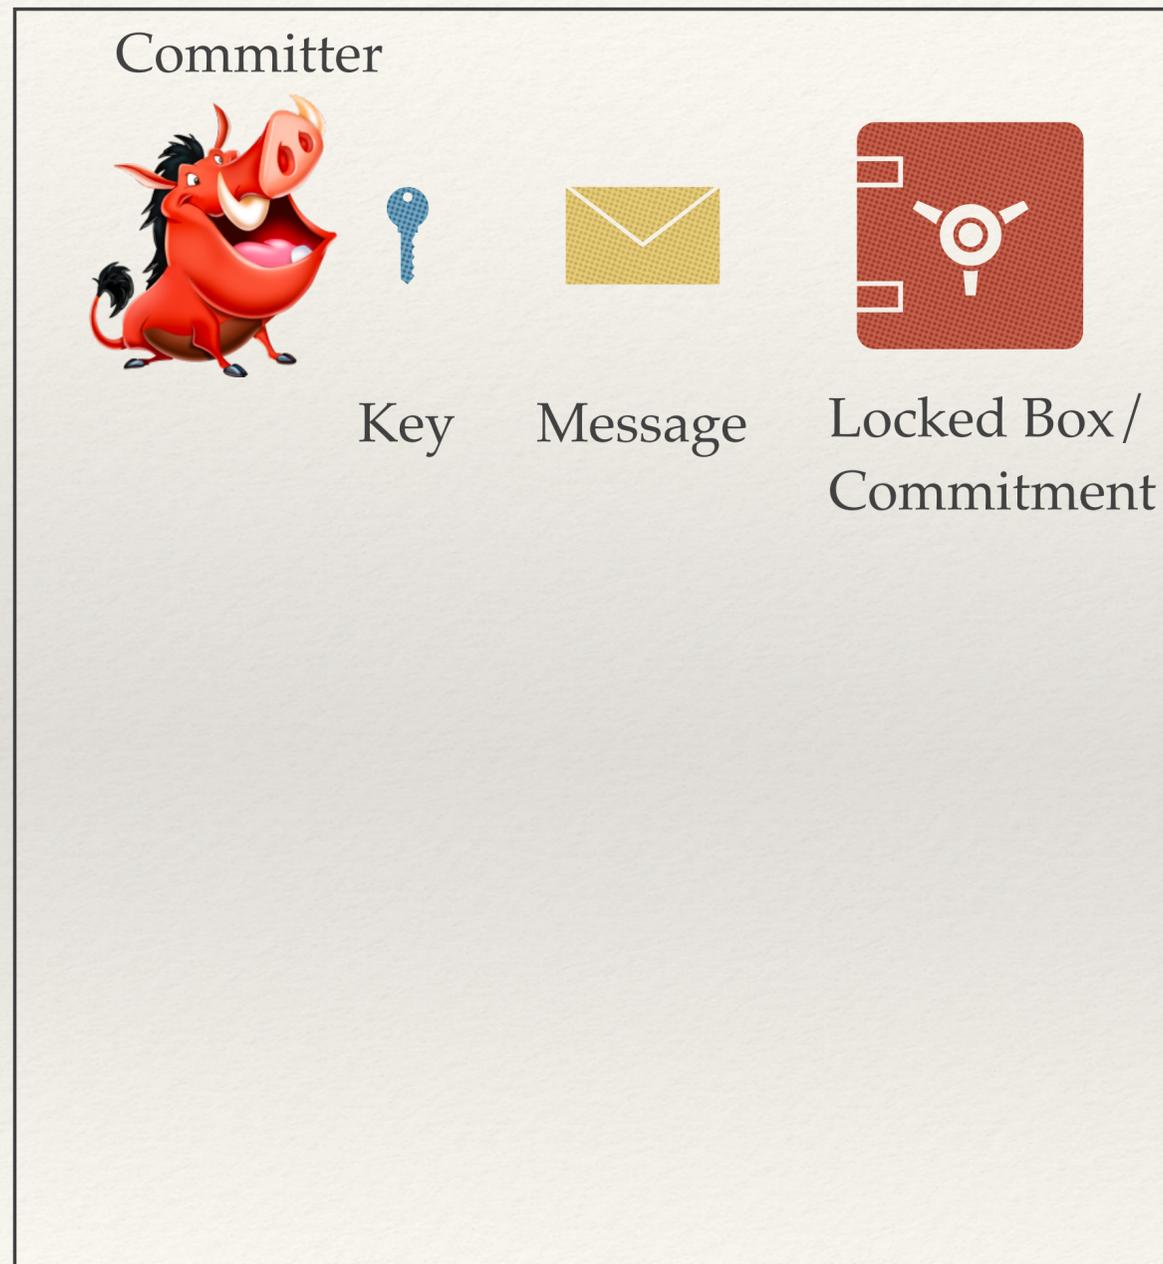


=

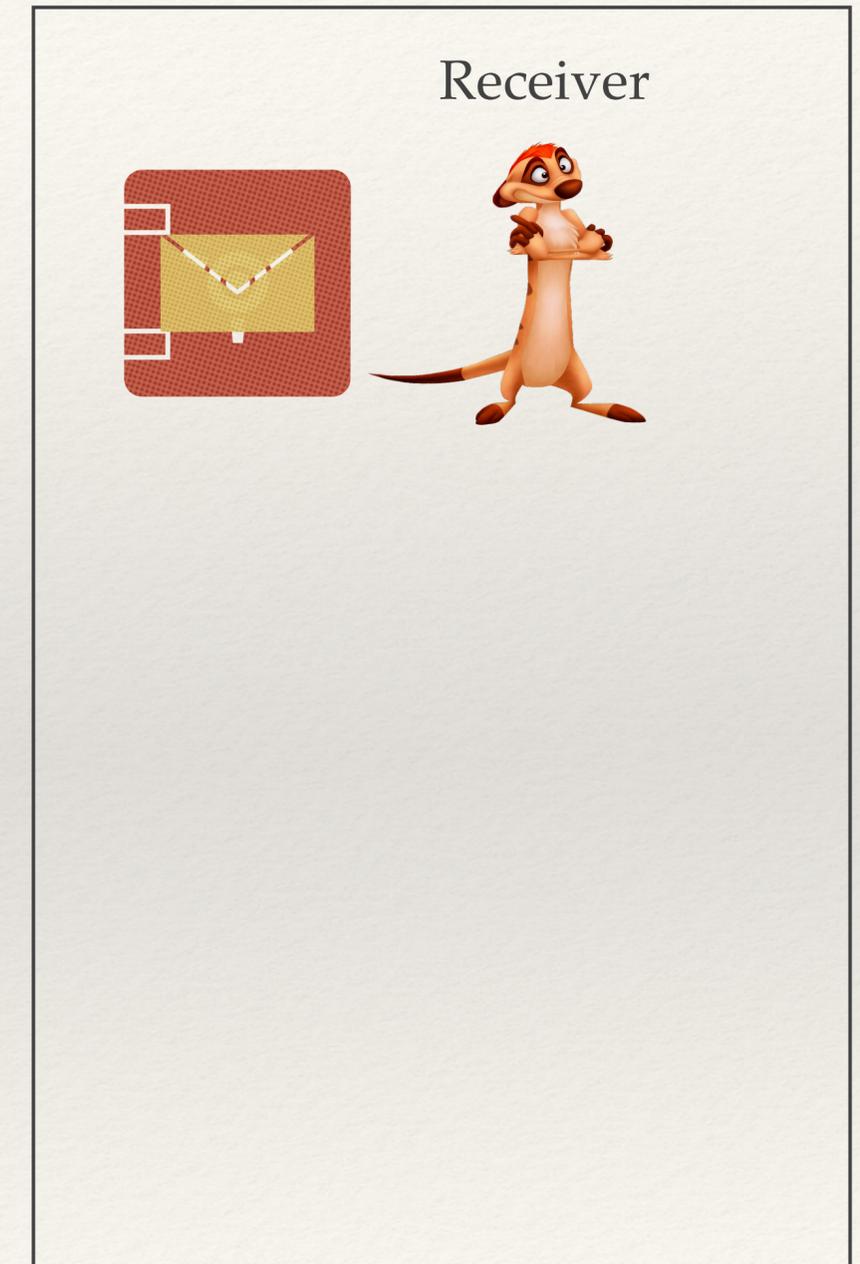


The combined keys open
the combined boxes

Disjunctive OR Arguments



The prover can convince the receiver that the message is either 0 or 1 without revealing which one it is

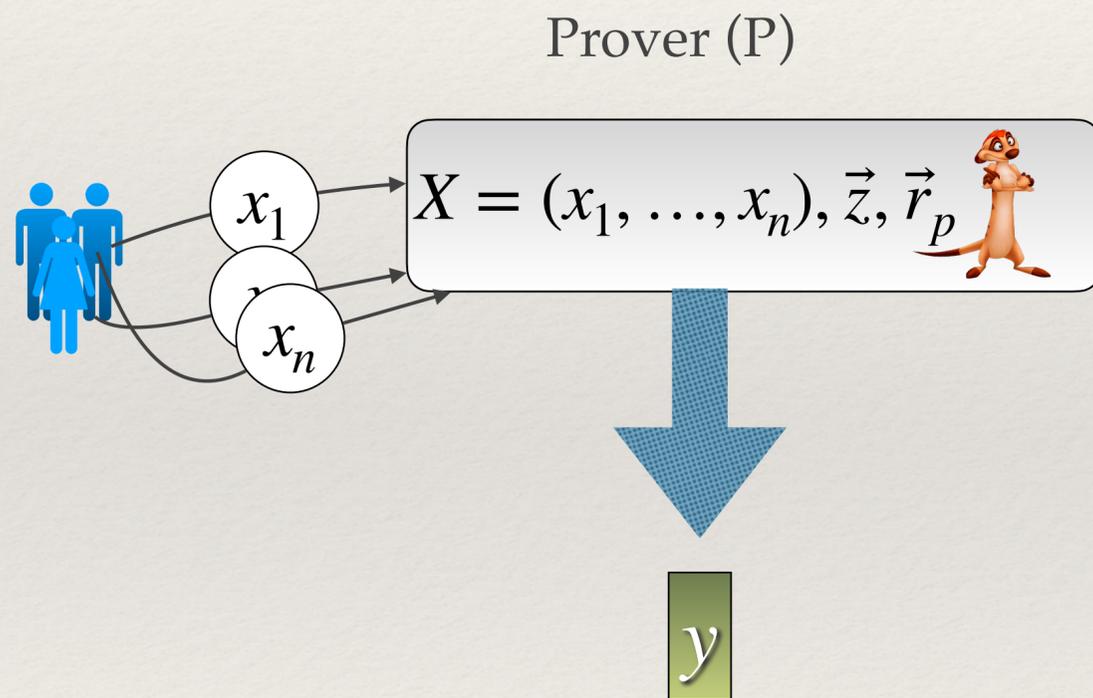


Quick Recap

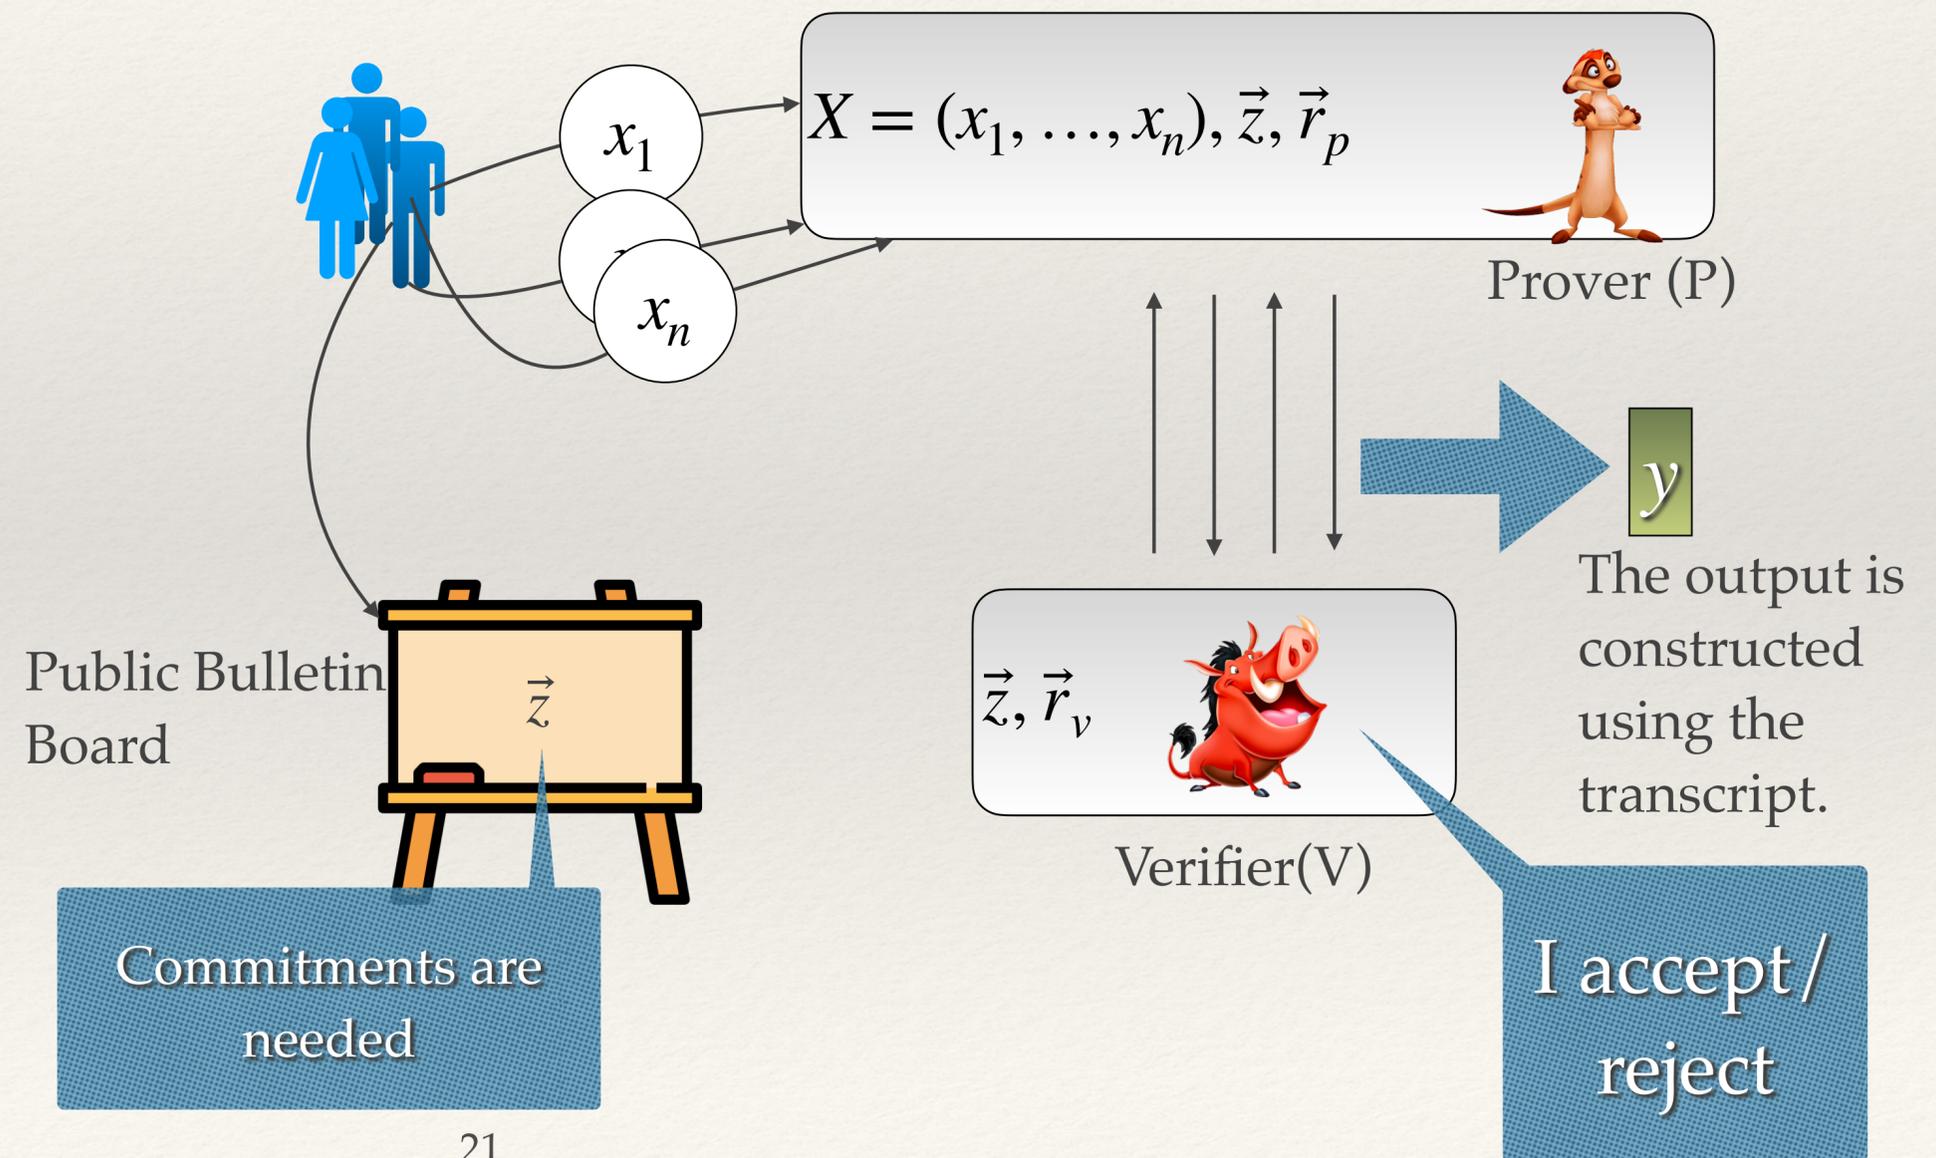
- ❖ We have commitments that are homomorphic and support OR arguments.

Verifiable - The Setting

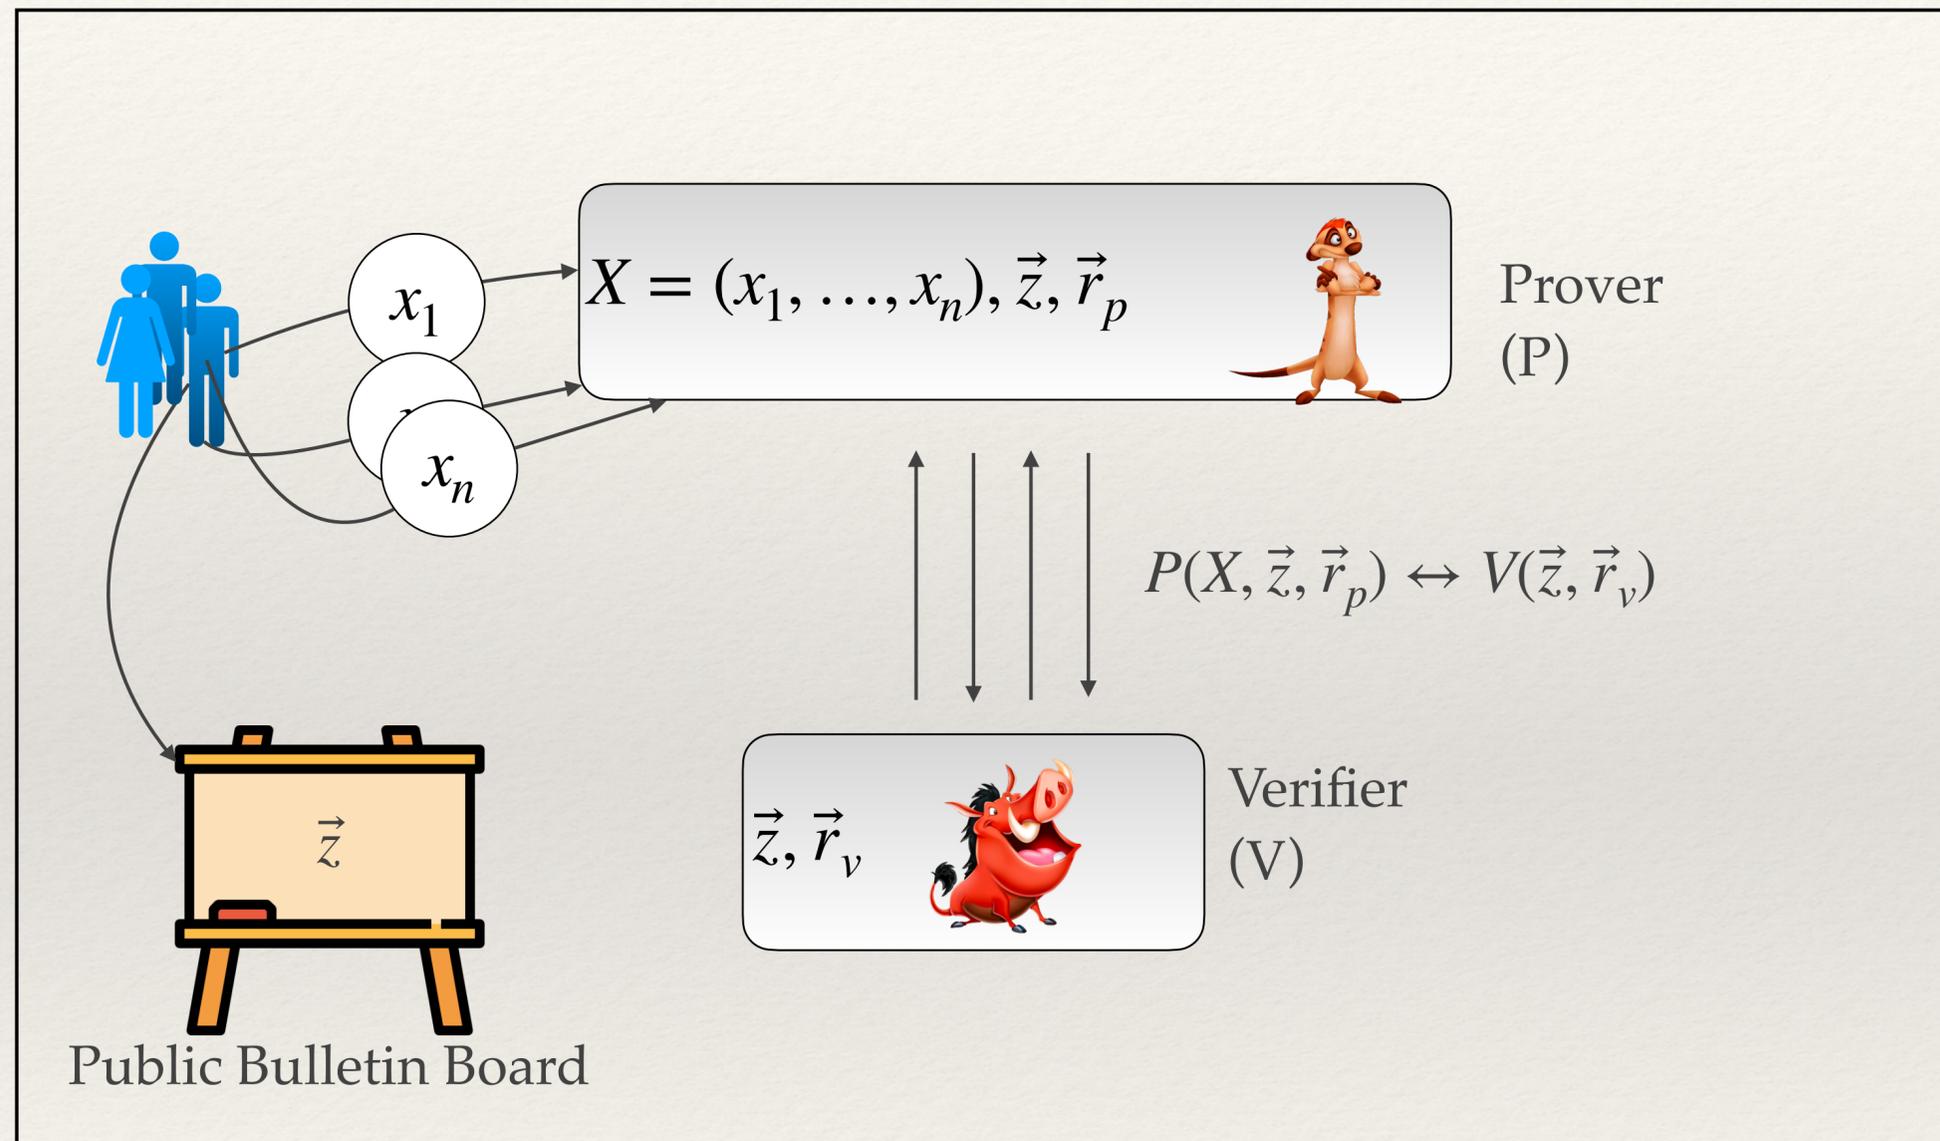
Un-verifiable DP



Verifiable DP



Verifiable DP

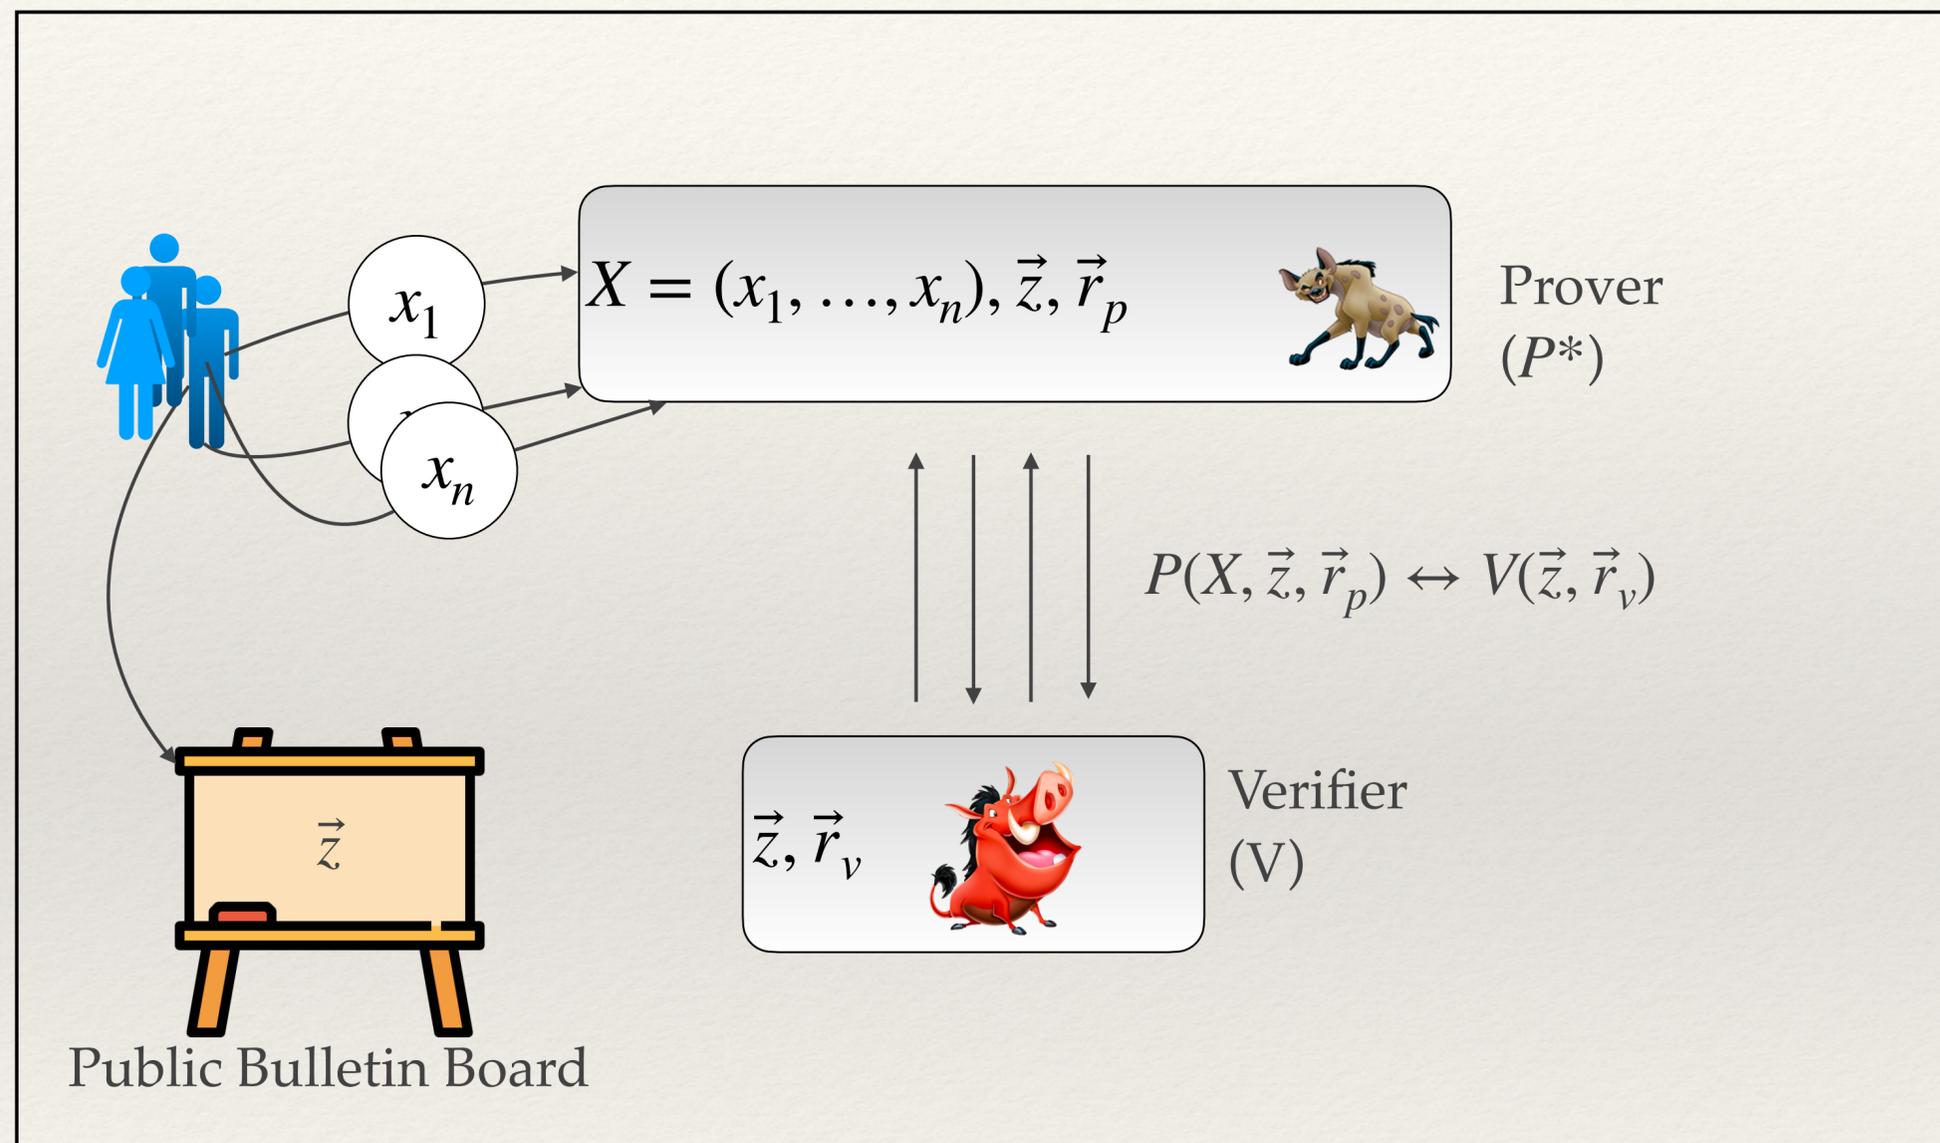


Completeness:

If both the prover and the verifier are honest, then $y \stackrel{\$}{\leftarrow} M(X, Q)$ and

$$\Pr[\text{Verify}(P \leftrightarrow V) = 1] = 1$$

Verifiable DP

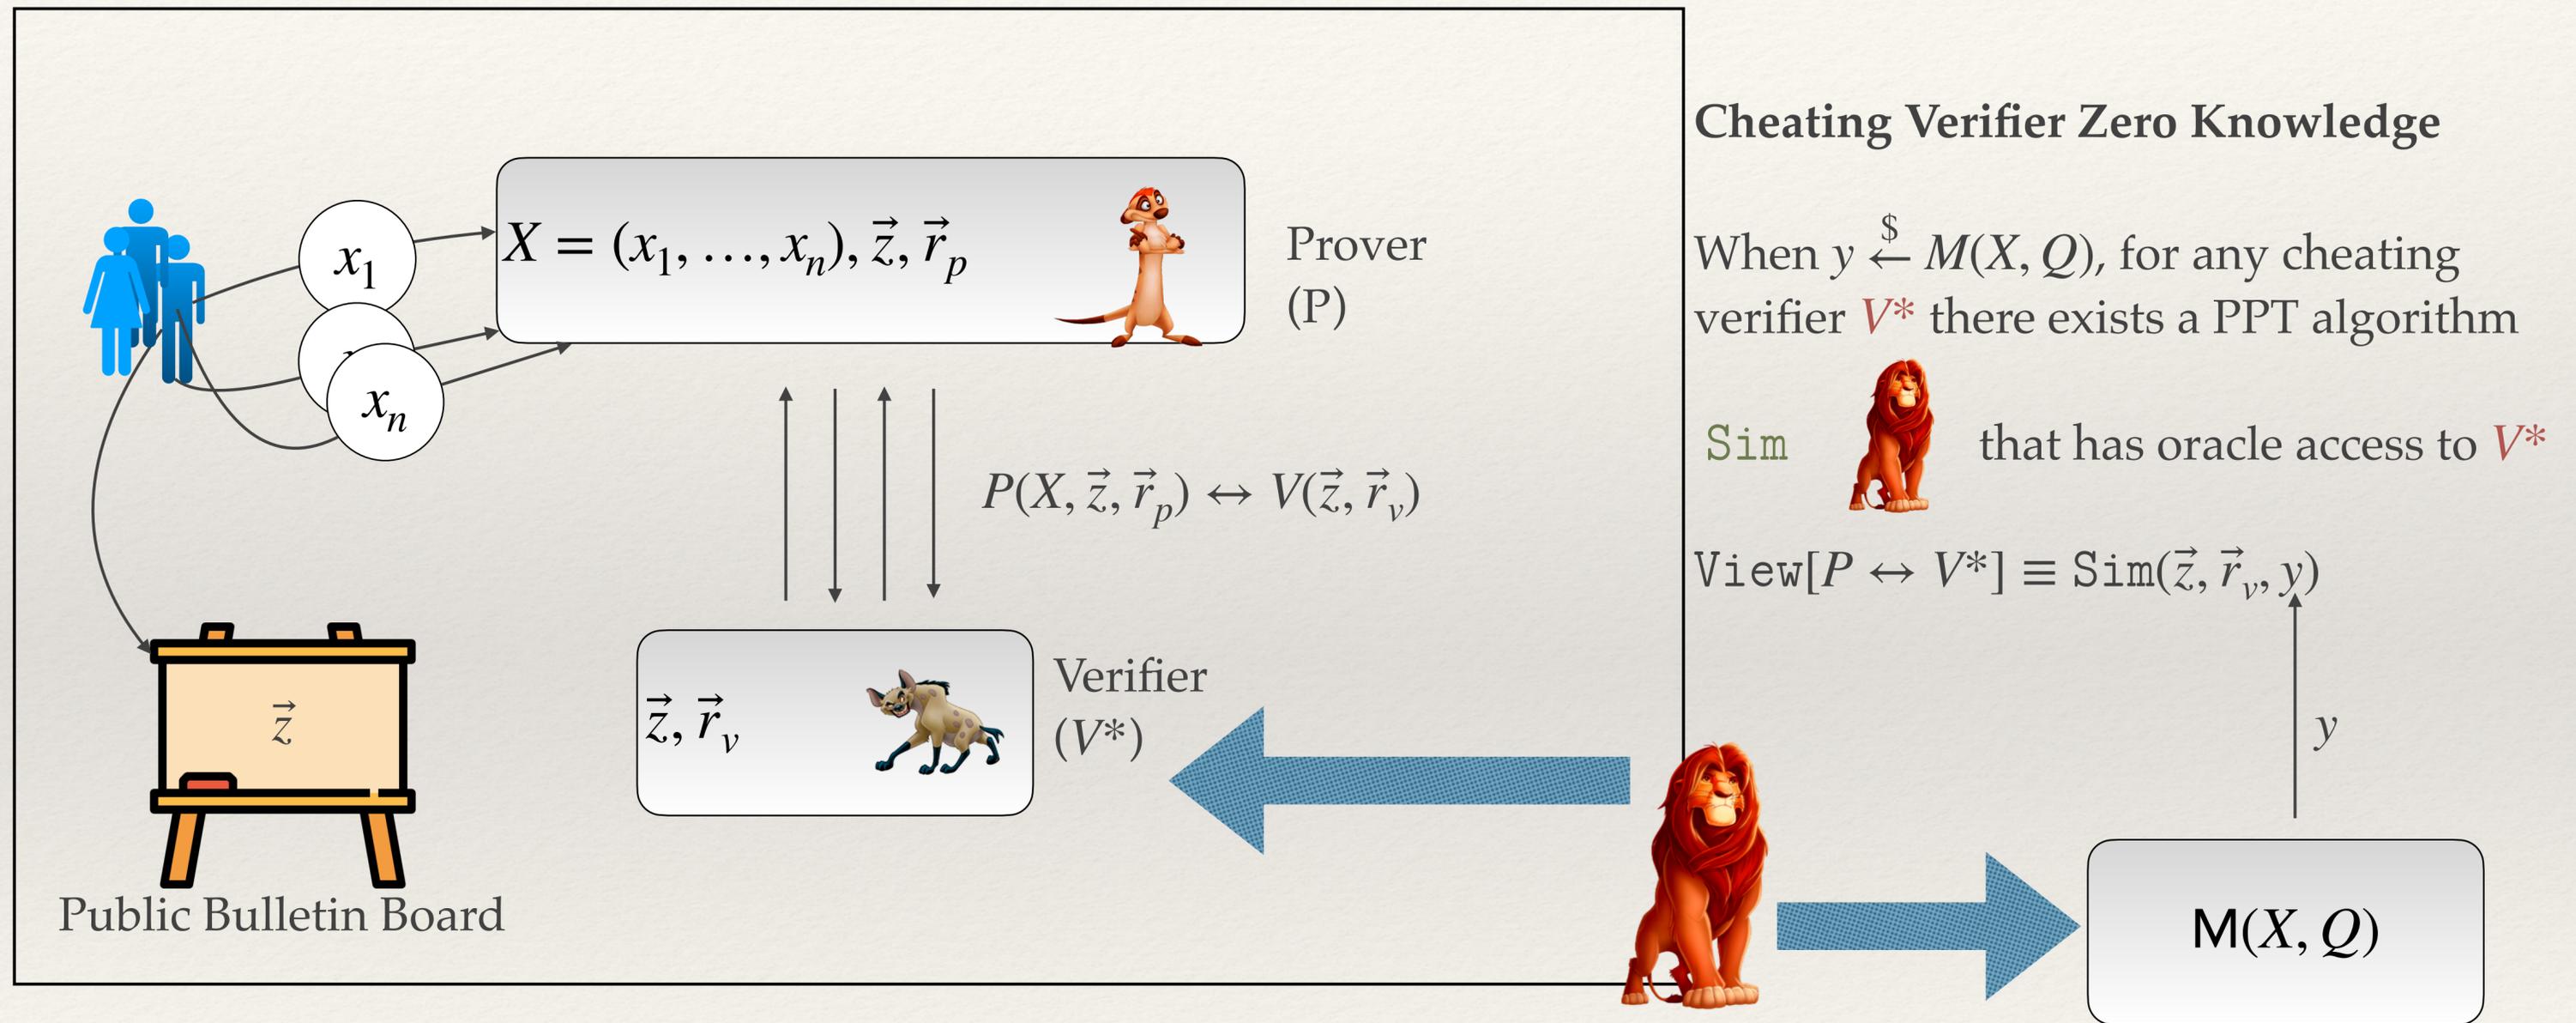


Soundness

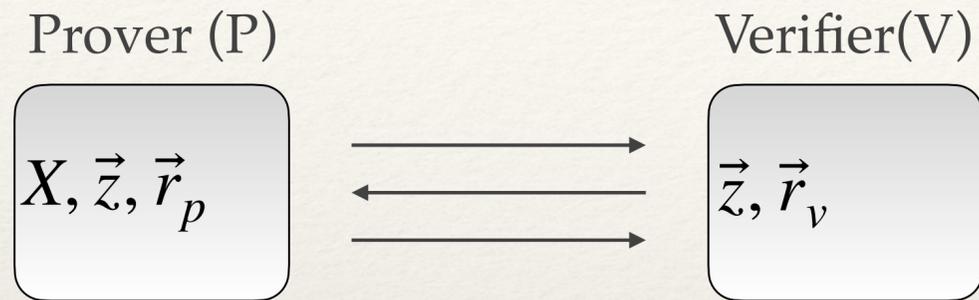
For any cheating prover P^* that samples y from a distribution \mathcal{D} such that $\text{TV}(\mathcal{M}(X, Q), \mathcal{D}) > \mu(\kappa)$

$$\Pr[\text{Verify}(P^* \leftrightarrow V) = 1] \leq 1/3$$

Verifiable DP

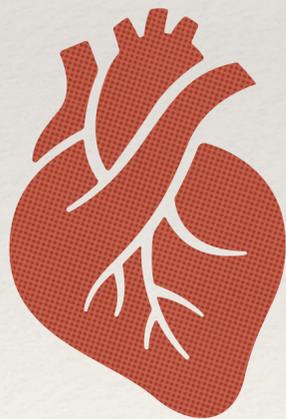


The Soundness/ZK conflict



$$Z \stackrel{\$}{\leftarrow} \text{Binomial}(\eta, \frac{1}{2})$$
$$y = Q(x_1, \dots, x_n) + Z \longrightarrow M(X; Q)$$

THE HEART OF THE PROBLEM



The output is a function of the provers local randomness. However the prover cannot ever reveal this randomness to the verifier as it would compromise DP.

The prover must find a way to prove that Z was sampled from the right distribution without ever revealing any information about Z .

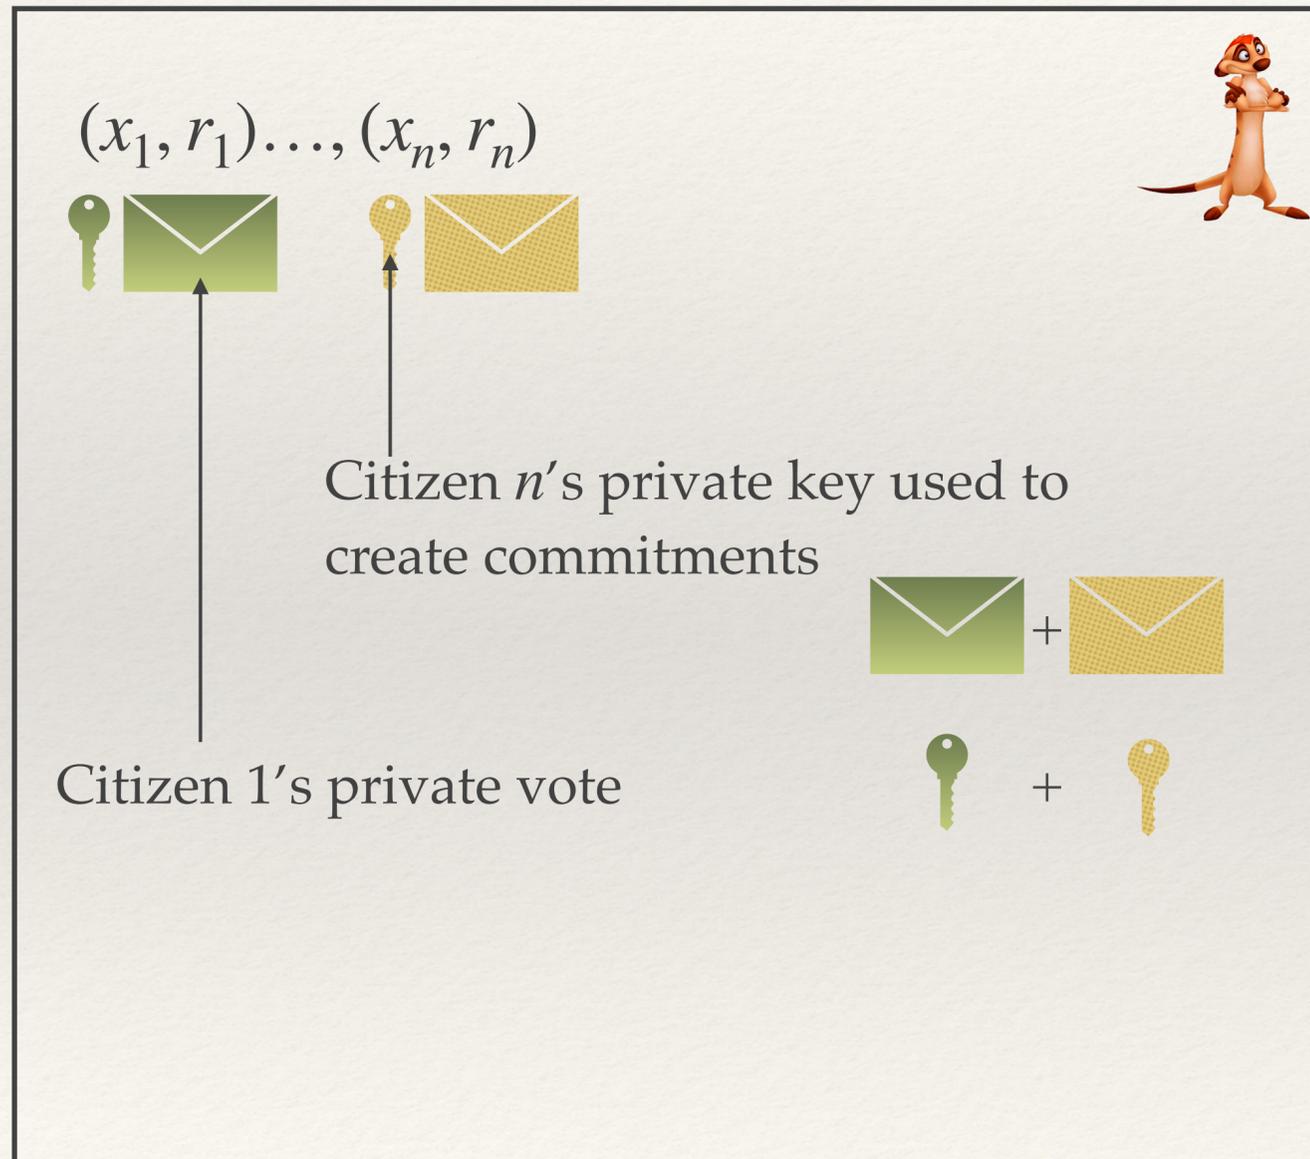
However, we also need some shared information (like say public randomness) for the verifier to be able to confident that Z is sampled correctly.

*Not to be confused with Proof Of Knowledge

** The noise used is not pseudorandom noise either

Non Private Counting

Server / Prover

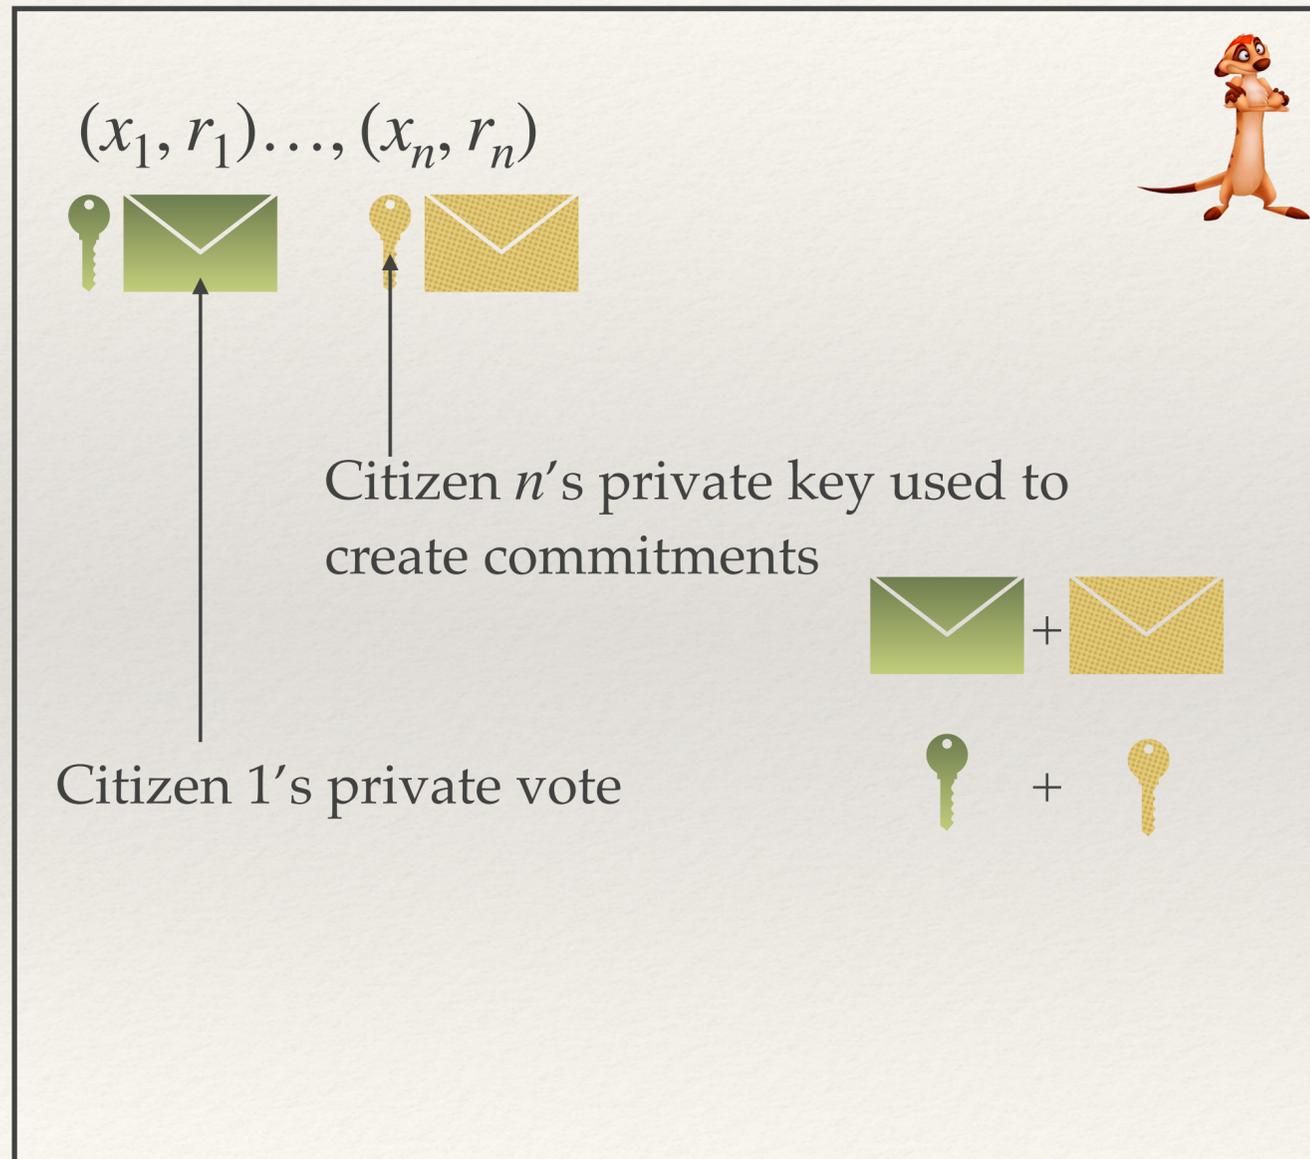


$\text{Com}(x_1, r_1), \dots, \text{Com}(x_n, r_n)$

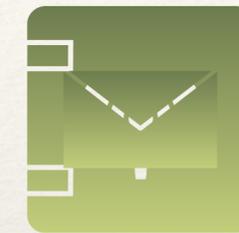


Non Private Counting

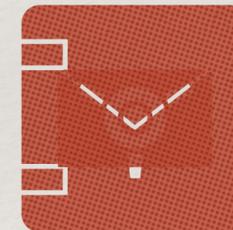
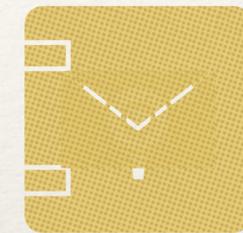
Server / Prover



$\text{Com}(x_1, r_1), \dots, \text{Com}(x_n, r_n)$

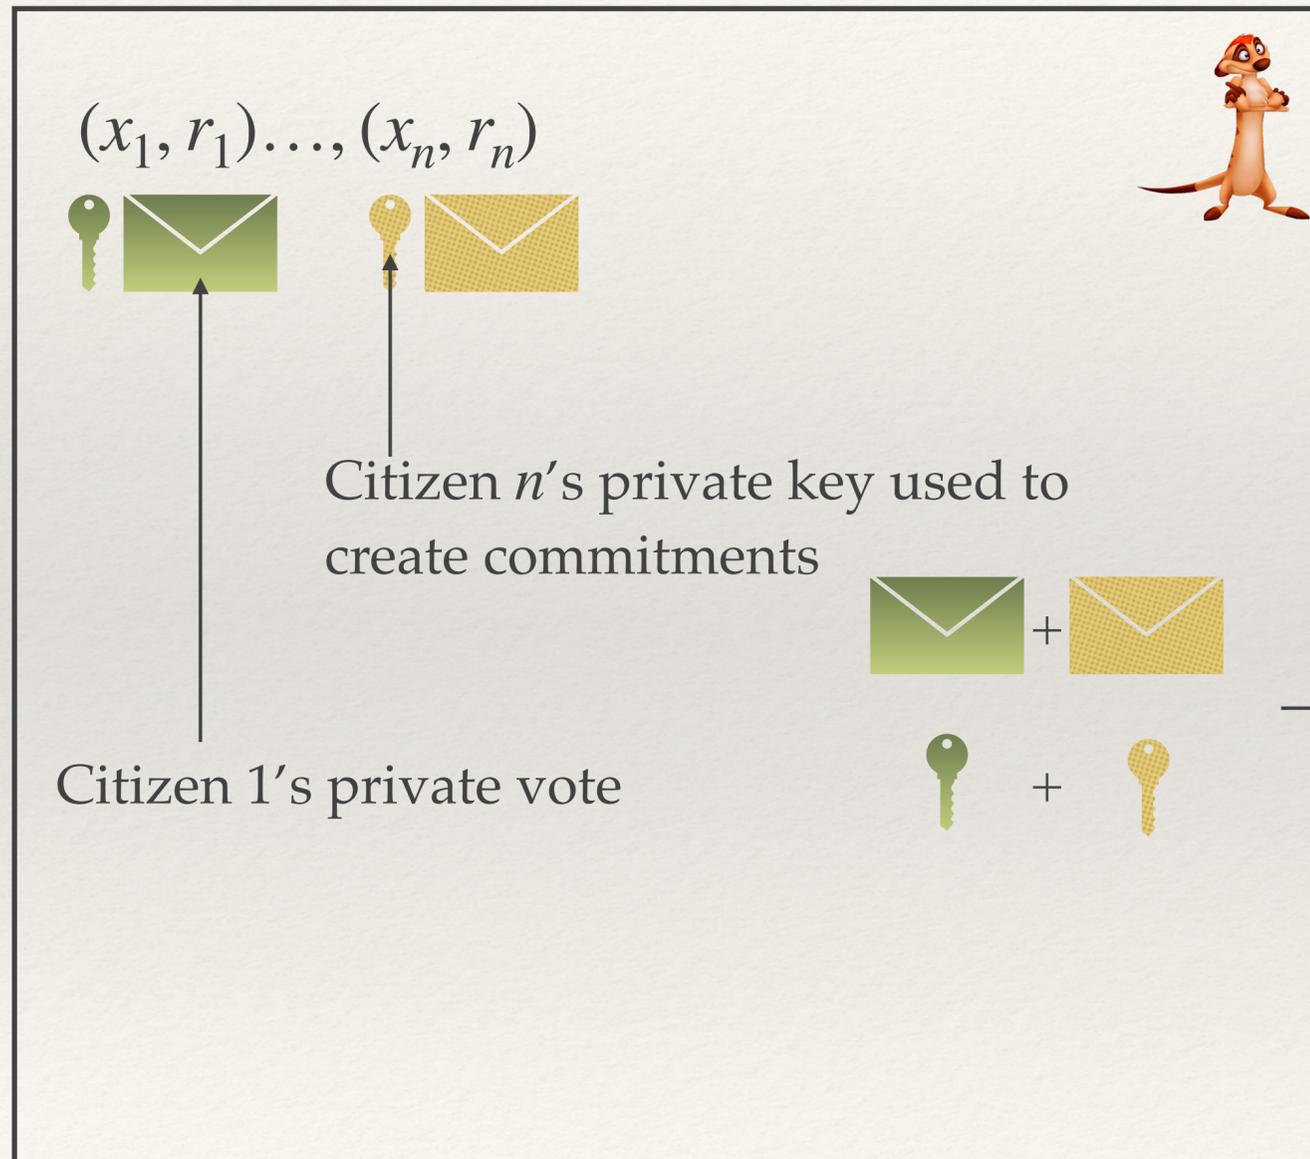


+



Non Private Counting

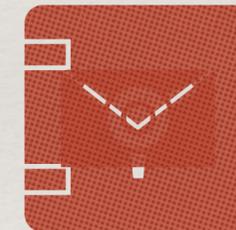
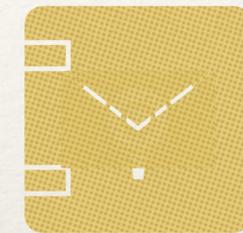
Server / Prover



$\text{Com}(x_1, r_1), \dots, \text{Com}(x_n, r_n)$



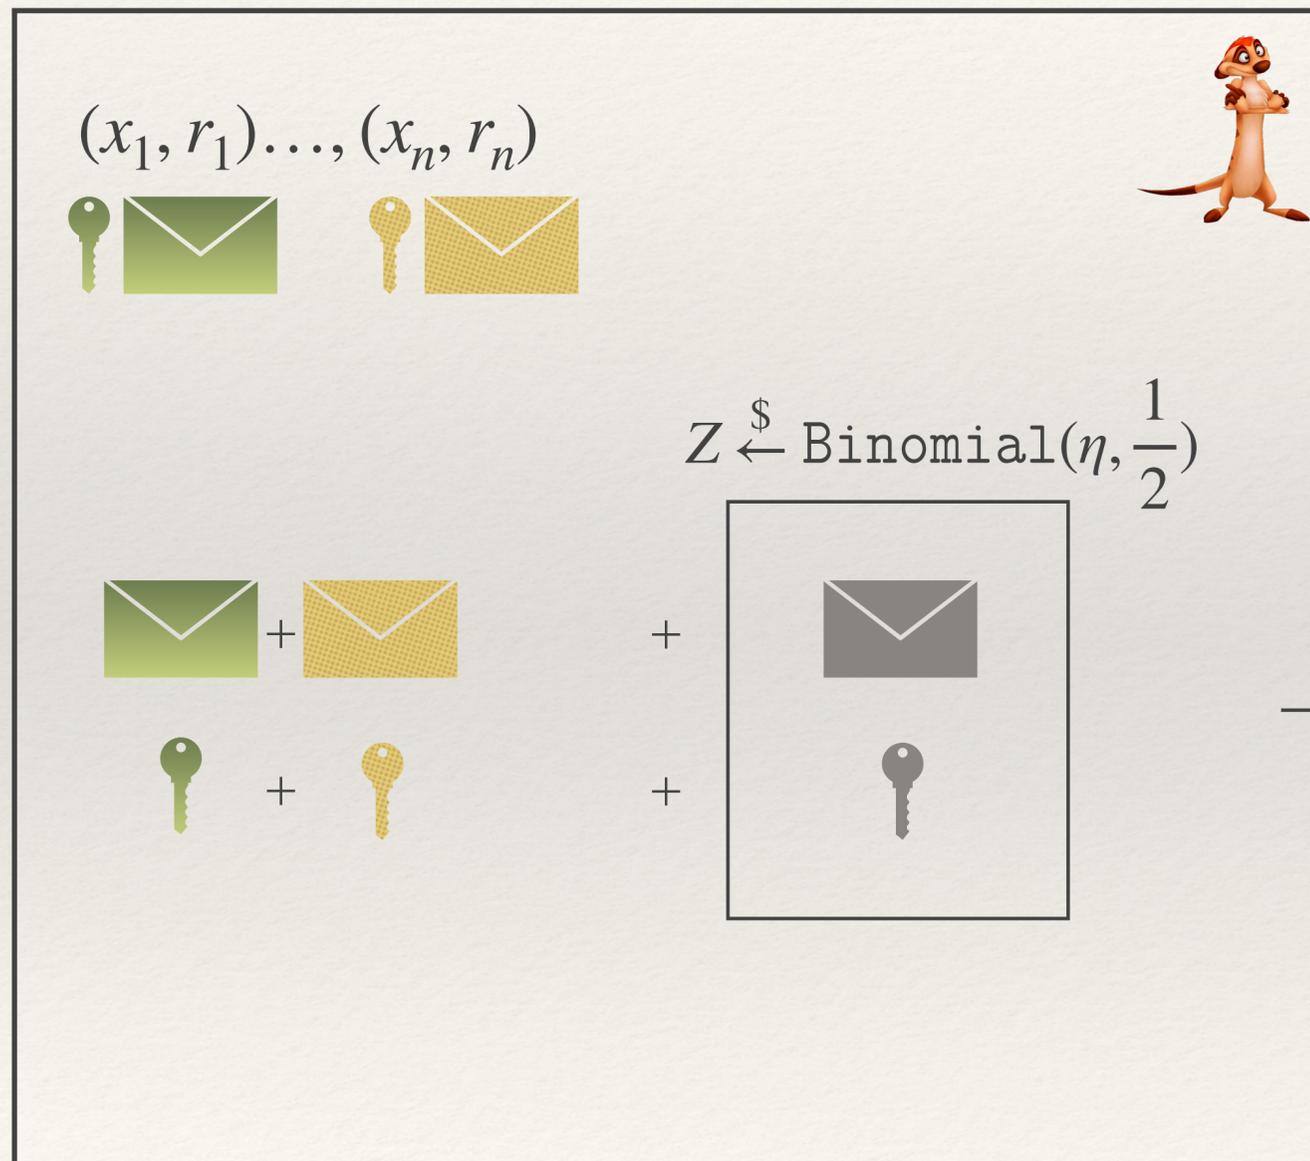
+



Check if key opens locked box properly.

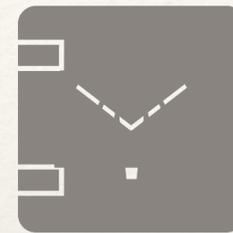
Verifiable DP counting - Essence

Server / Prover



Somehow need to
create public
commitment to Z

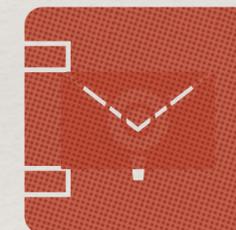
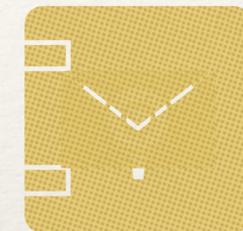
$\text{Com}(x_1, r_1), \dots, \text{Com}(x_n, r_n)$



+



+



Check if key opens locked box properly.

A Simple Trick

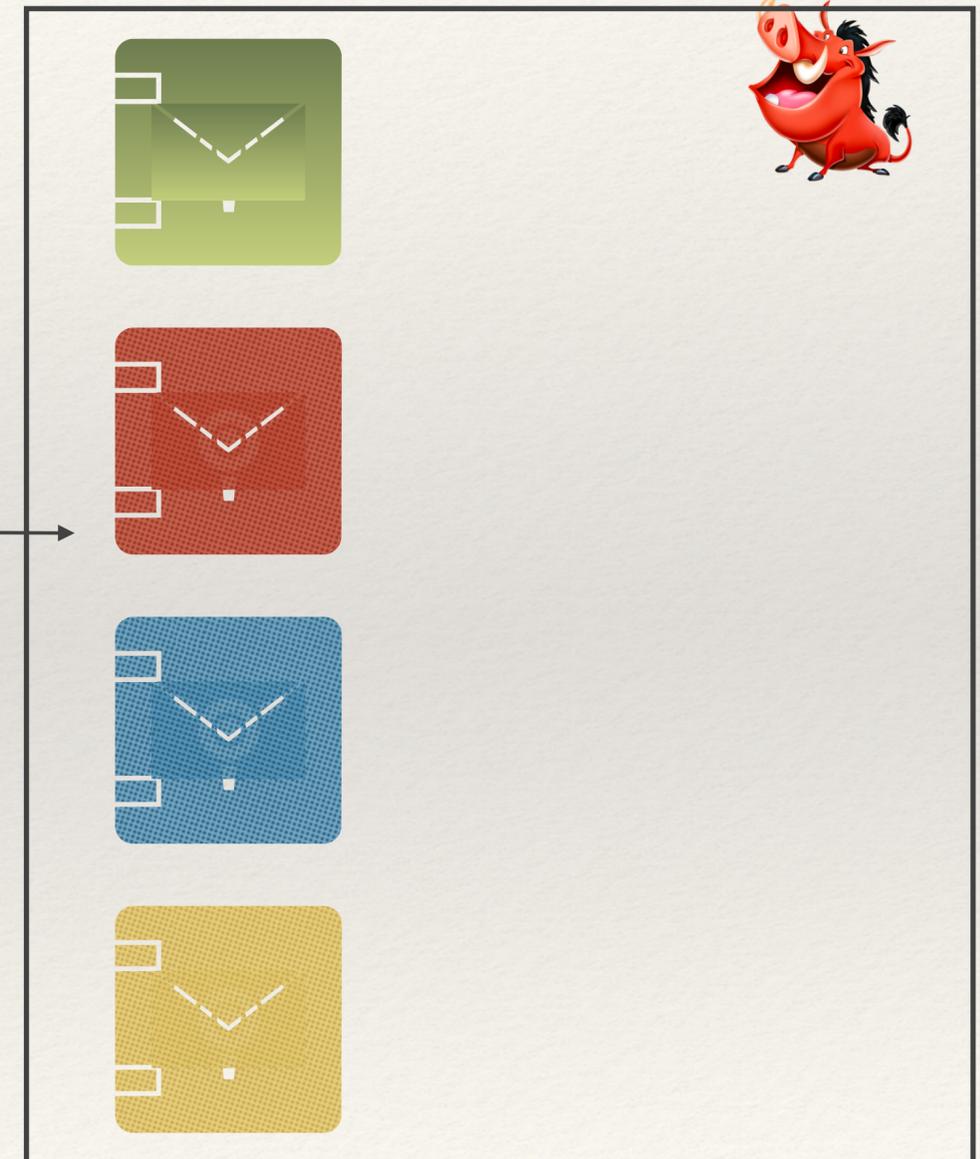
Server / Prover



Note we cannot say anything about the distribution from which these bits are being sampled.

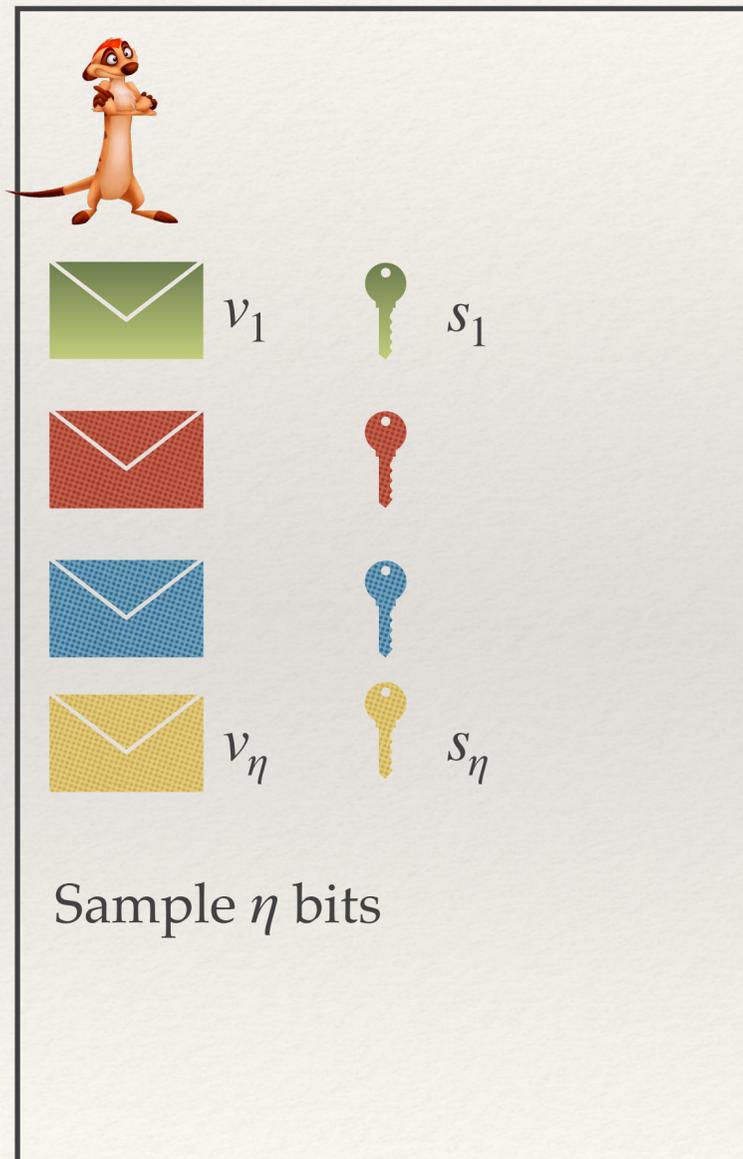
All the verifier knows is that these boxes are a commitment to a bit.

Verifier

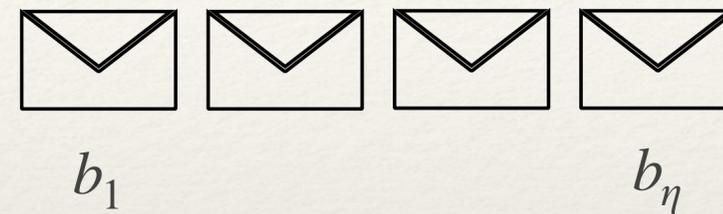


A Simple Trick

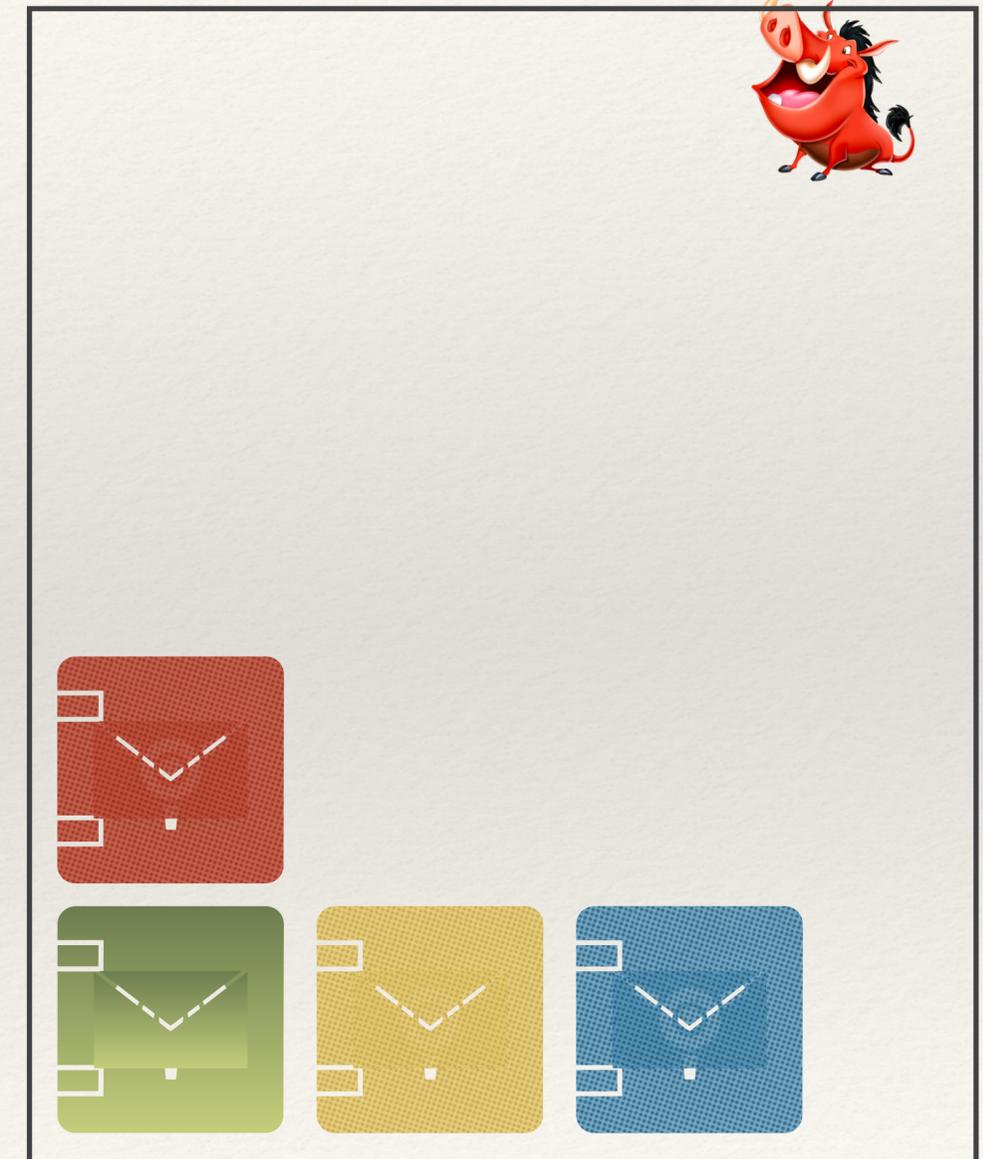
Server / Prover



Verifier generates η public unbiased coins

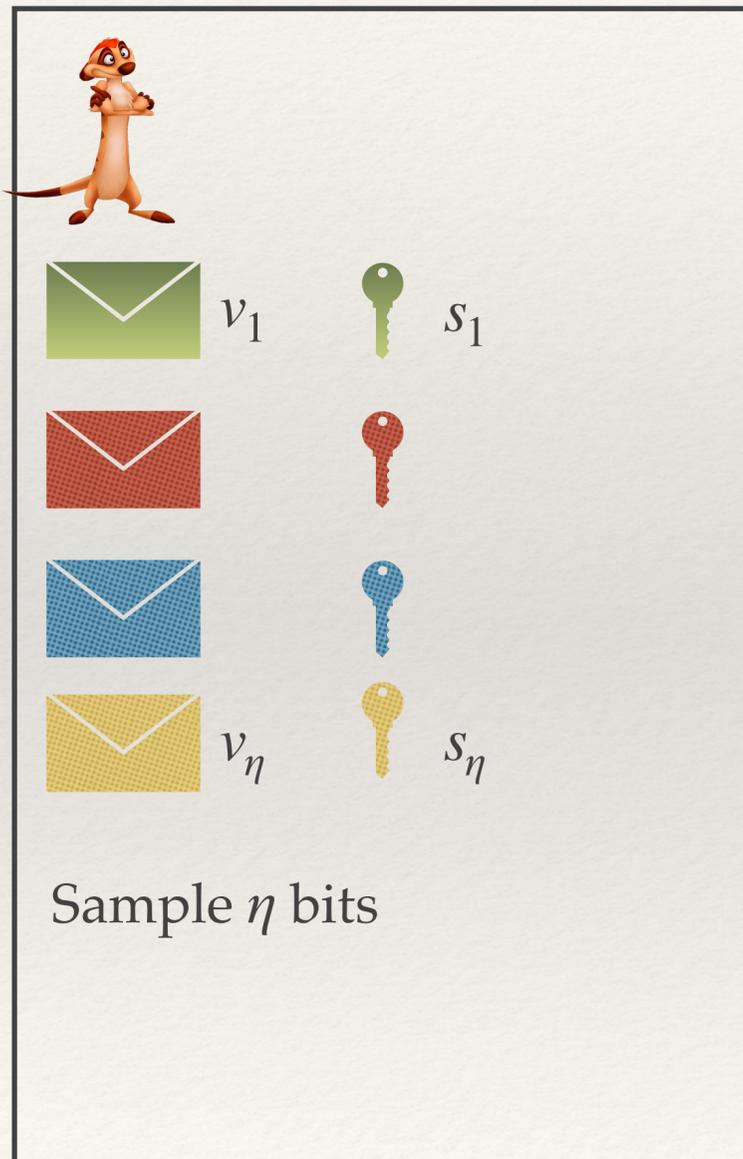


Verifier



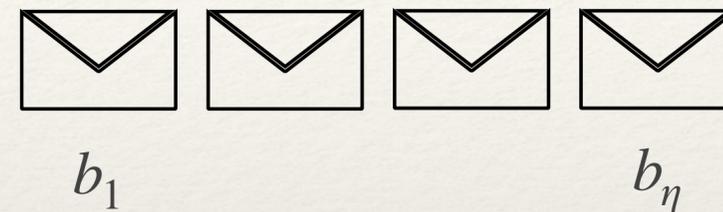
The Final Trick

Server / Prover



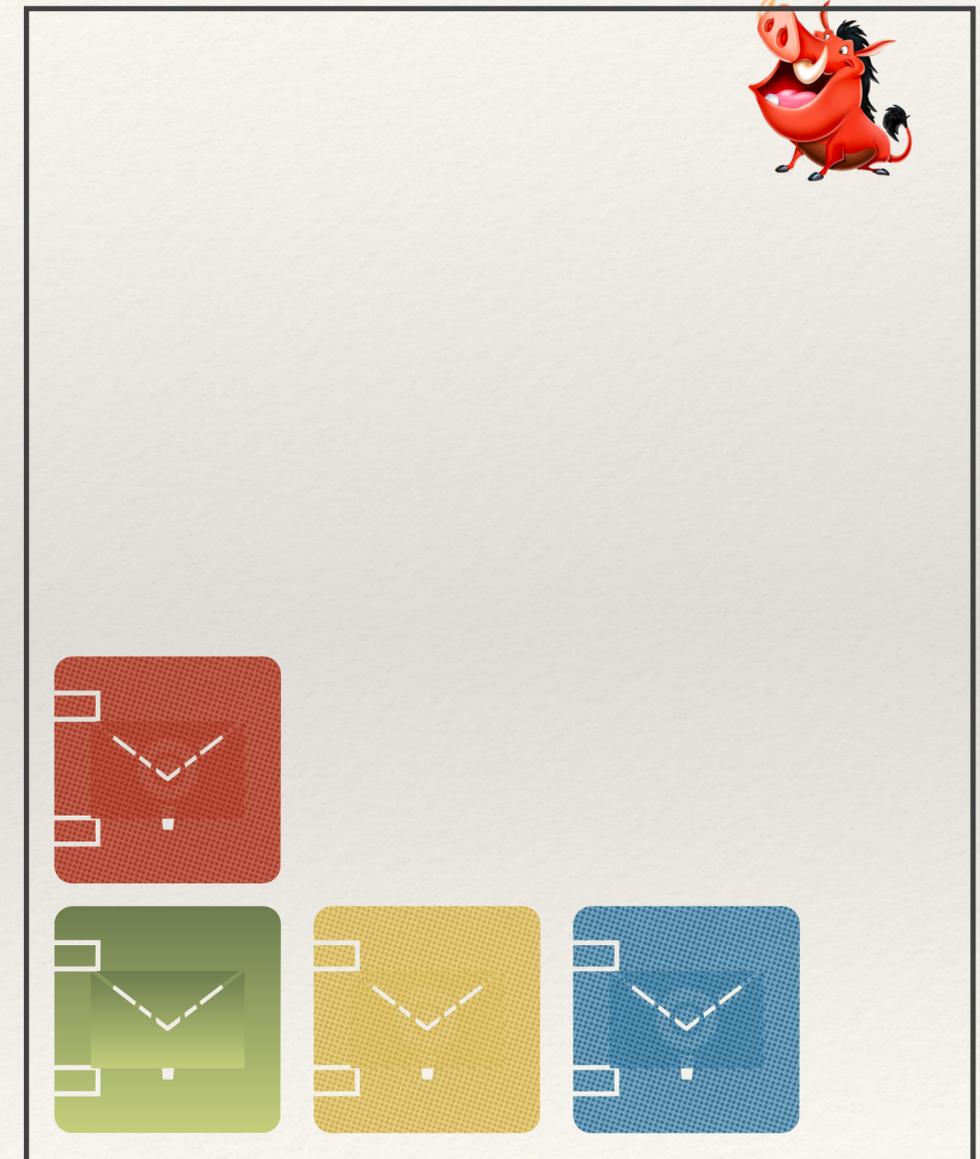
The Server/Prover side of the diagram is enclosed in a black-bordered box. At the top left is a cartoon dog icon. Below it are four rows of items, each consisting of an envelope icon and a key icon. The first row has a green envelope labeled v_1 and a green key labeled s_1 . The second row has a red envelope and a red key. The third row has a blue envelope and a blue key. The fourth row has a yellow envelope labeled v_η and a yellow key labeled s_η . At the bottom left of the box, the text "Sample η bits" is written.

Verifier generates η public unbiased coins



If $b_i = 1$ then set $v_i = 1 - v_i$
Otherwise, leave v_i unchanged.

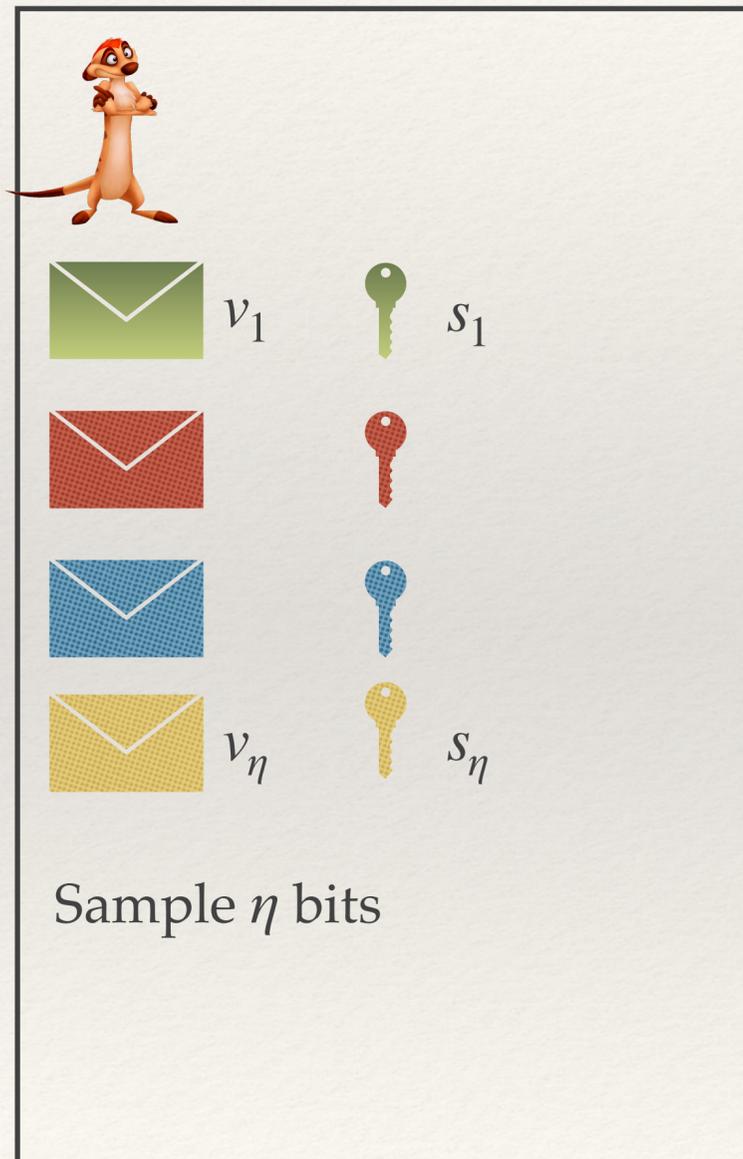
Verifier



The Verifier side of the diagram is enclosed in a black-bordered box. At the top right is a cartoon pig icon. Below it are four rows of items, each consisting of a square icon with a dashed envelope shape inside. The first row has a red square icon. The second row has a green square icon, a yellow square icon, and a blue square icon. The third row has a yellow square icon and a blue square icon. The fourth row has a blue square icon.

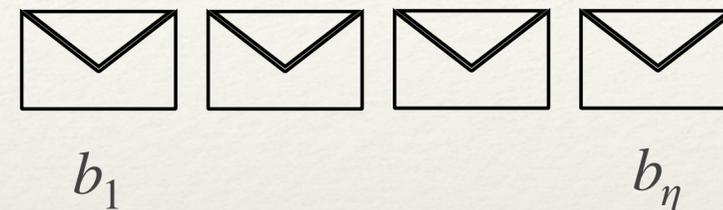
The Final Trick

Server / Prover



The Server/Prover side of the diagram features a cartoon dog character at the top left. Below it, there are four rows of items. Each row consists of an envelope icon on the left and a key icon on the right. The first row has a green envelope labeled v_1 and a green key labeled s_1 . The second row has a red envelope and a red key. The third row has a blue envelope and a blue key. The fourth row has a yellow envelope labeled v_η and a yellow key labeled s_η . At the bottom of this section, the text "Sample η bits" is written.

Verifier generates η public unbiased coins



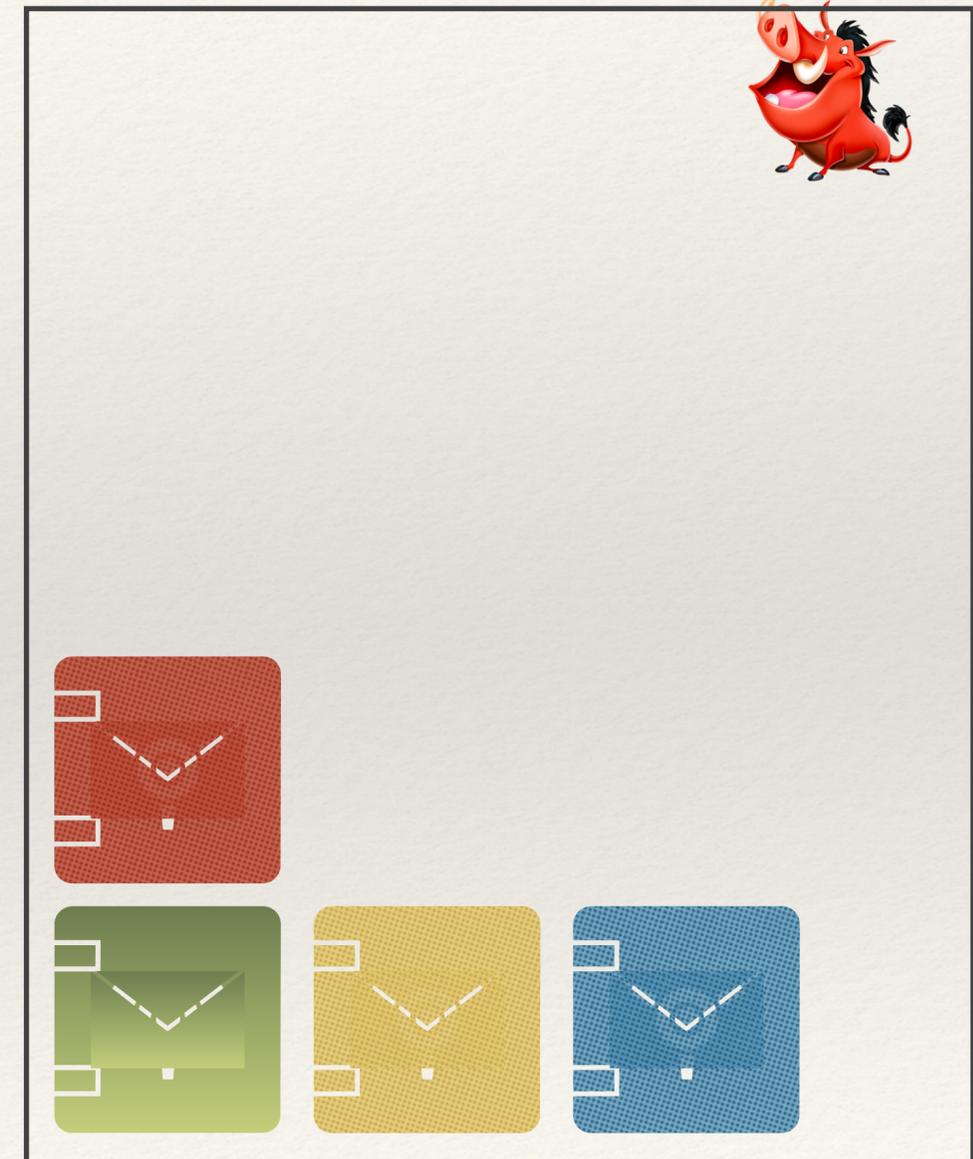
If $b_i = 1$ then set $v_i = 1 - v_i$ and $s_i = 1 - s_i$

Otherwise, leave v_i and s_i unchanged.

Observation 1:
The updates are LINEAR conditioned on b_i

Without ever seeing v_i the verifier
can update
 $\text{Com}(v_i, s_i) = \text{Com}(1, 1) - \text{Com}(v_i, s_i)$

Verifier



The Verifier side of the diagram features a cartoon pig character at the top right. Below it, there are three rows of items. The first row has a red envelope icon with a dashed white outline, representing an update. The second row has three envelope icons: a green one, a yellow one, and a blue one, all with dashed white outlines, representing the original commitments. The text "Sample η bits" is written at the bottom of this section.

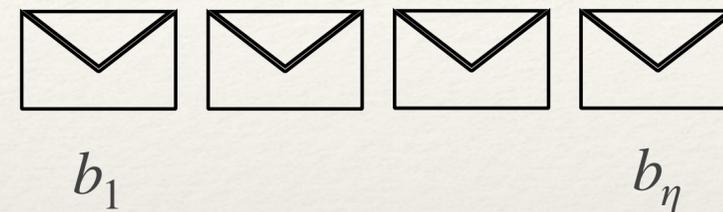
The Final Trick

Server / Prover



The Server/Prover side of the diagram features a cartoon dog character at the top left. Below it, there are four rows of icons. Each row consists of an envelope icon on the left and a key icon on the right. The first row has a green envelope labeled v_1 and a green key labeled s_1 . The second row has a red envelope and a red key. The third row has a blue envelope and a blue key. The fourth row has a yellow envelope labeled v_η and a yellow key labeled s_η . At the bottom of this section, the text "Sample η bits" is written.

Verifier generates η public unbiased coins



If $b_i = 1$ then set $v_i = 1 - v_i$

Otherwise, leave v_i unchanged.

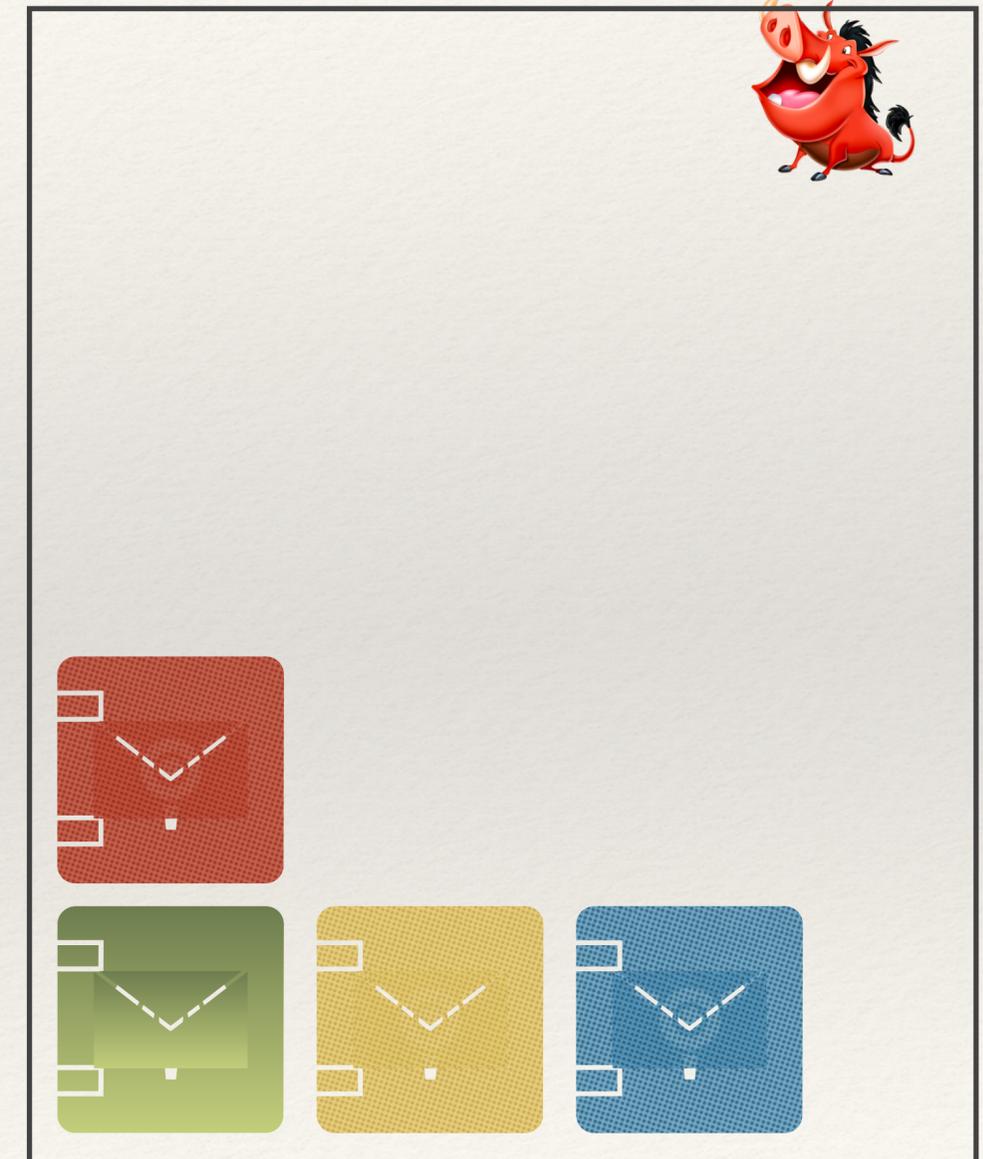
Observation 2:

The above conditional statement is equivalent to

$$v_i = v_i \oplus b_i$$

This forces the provers bit to have the correct distribution.

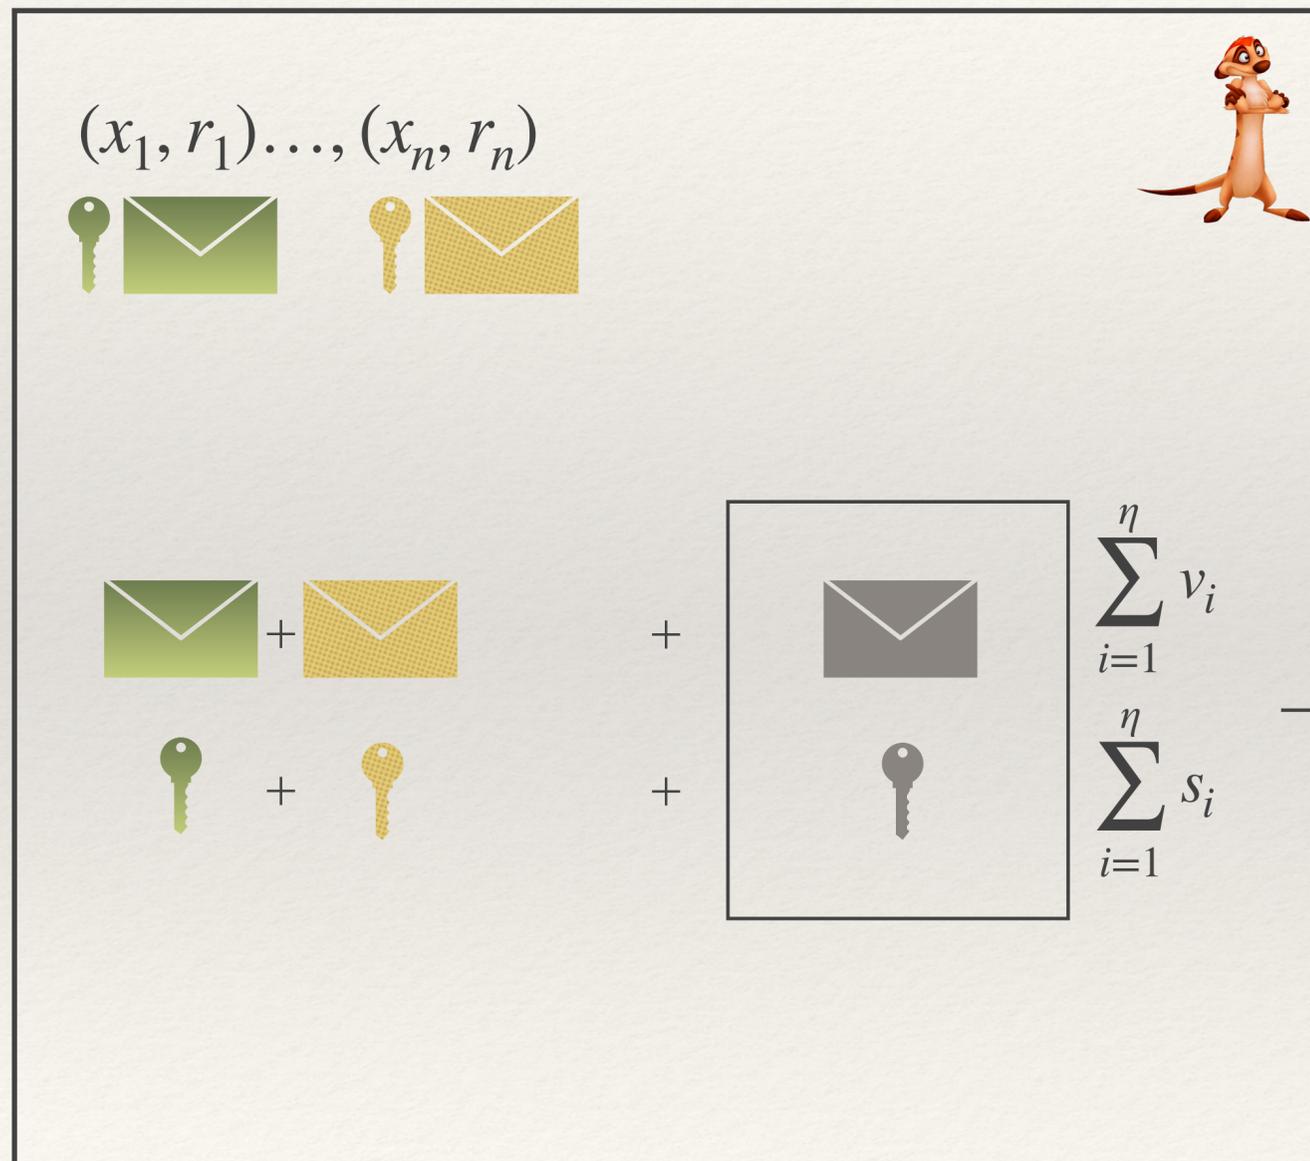
Verifier



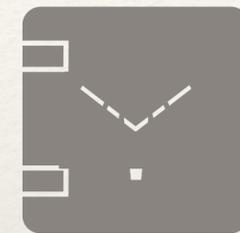
The Verifier side of the diagram features a cartoon red devil character at the top right. Below it, there are three rows of icons. The first row has a red envelope with a dashed white outline. The second row has a green envelope with a dashed white outline. The third row has a yellow envelope with a dashed white outline. The fourth row has a blue envelope with a dashed white outline.

Final Check

Server / Prover



$$\text{Com}\left(\sum_{i=1}^{\eta} v_i, \sum_{i=1}^{\eta} s_i\right)$$



+

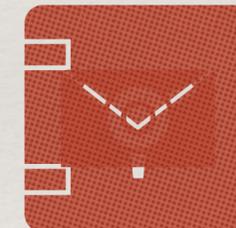
$$\text{Com}(x_1, r_1), \dots, \text{Com}(x_n, r_n)$$



+



Verifier



Check if key opens locked box properly.