SECURITY-PRESERVING DISTRIBUTED SAMPLERS

HOW TO GENERATE ANY CRS IN ONE ROUND WITHOUT RANDOM ORACLES

DAMIANO ABRAM

AARHUS UNIVERSITY

BRENT WATERS

UNIVERSITY OF TEXAS AT AUSTIN

—

NTT RESEARCH

MARK ZHANDRY

NTT RESEARCH

TRUSTED SETUP

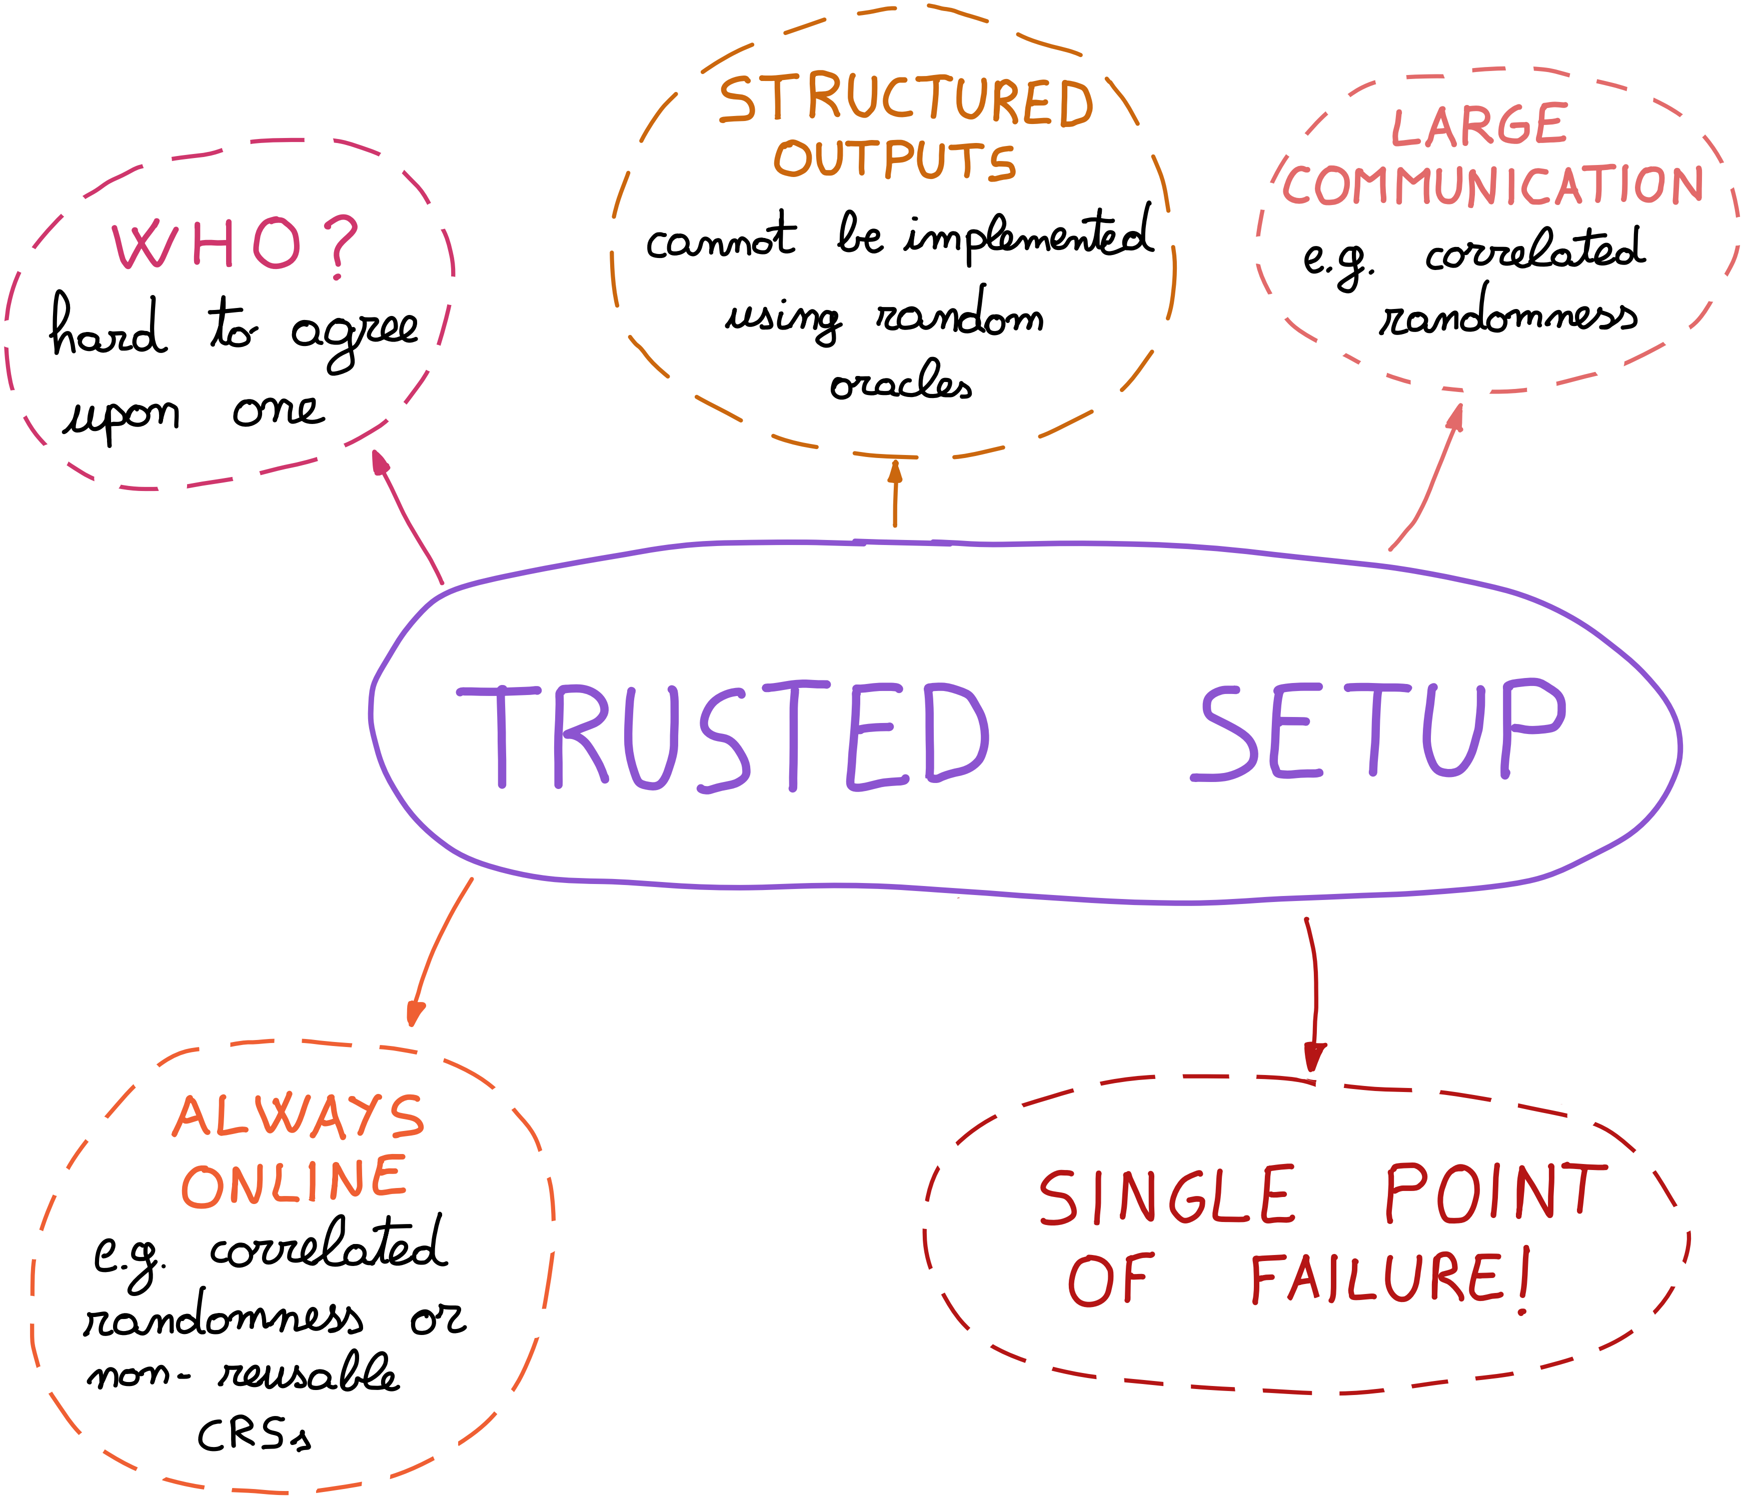
TRUSTED SETUP
WHO? hard to agree upon one
STRUCTURED OUTPUTS cannot be implemented using random oracles
LARGE COMMUNICATION e.g. correlated randomness
ALWAYS ONLINE e.g. correlated randomness or non-reusable CRSs
SINGLE POINT OF FAILURE!

DISTRIBUTED SAMPLERS

[EC: Abram, Scholl, Yakoubov 22]

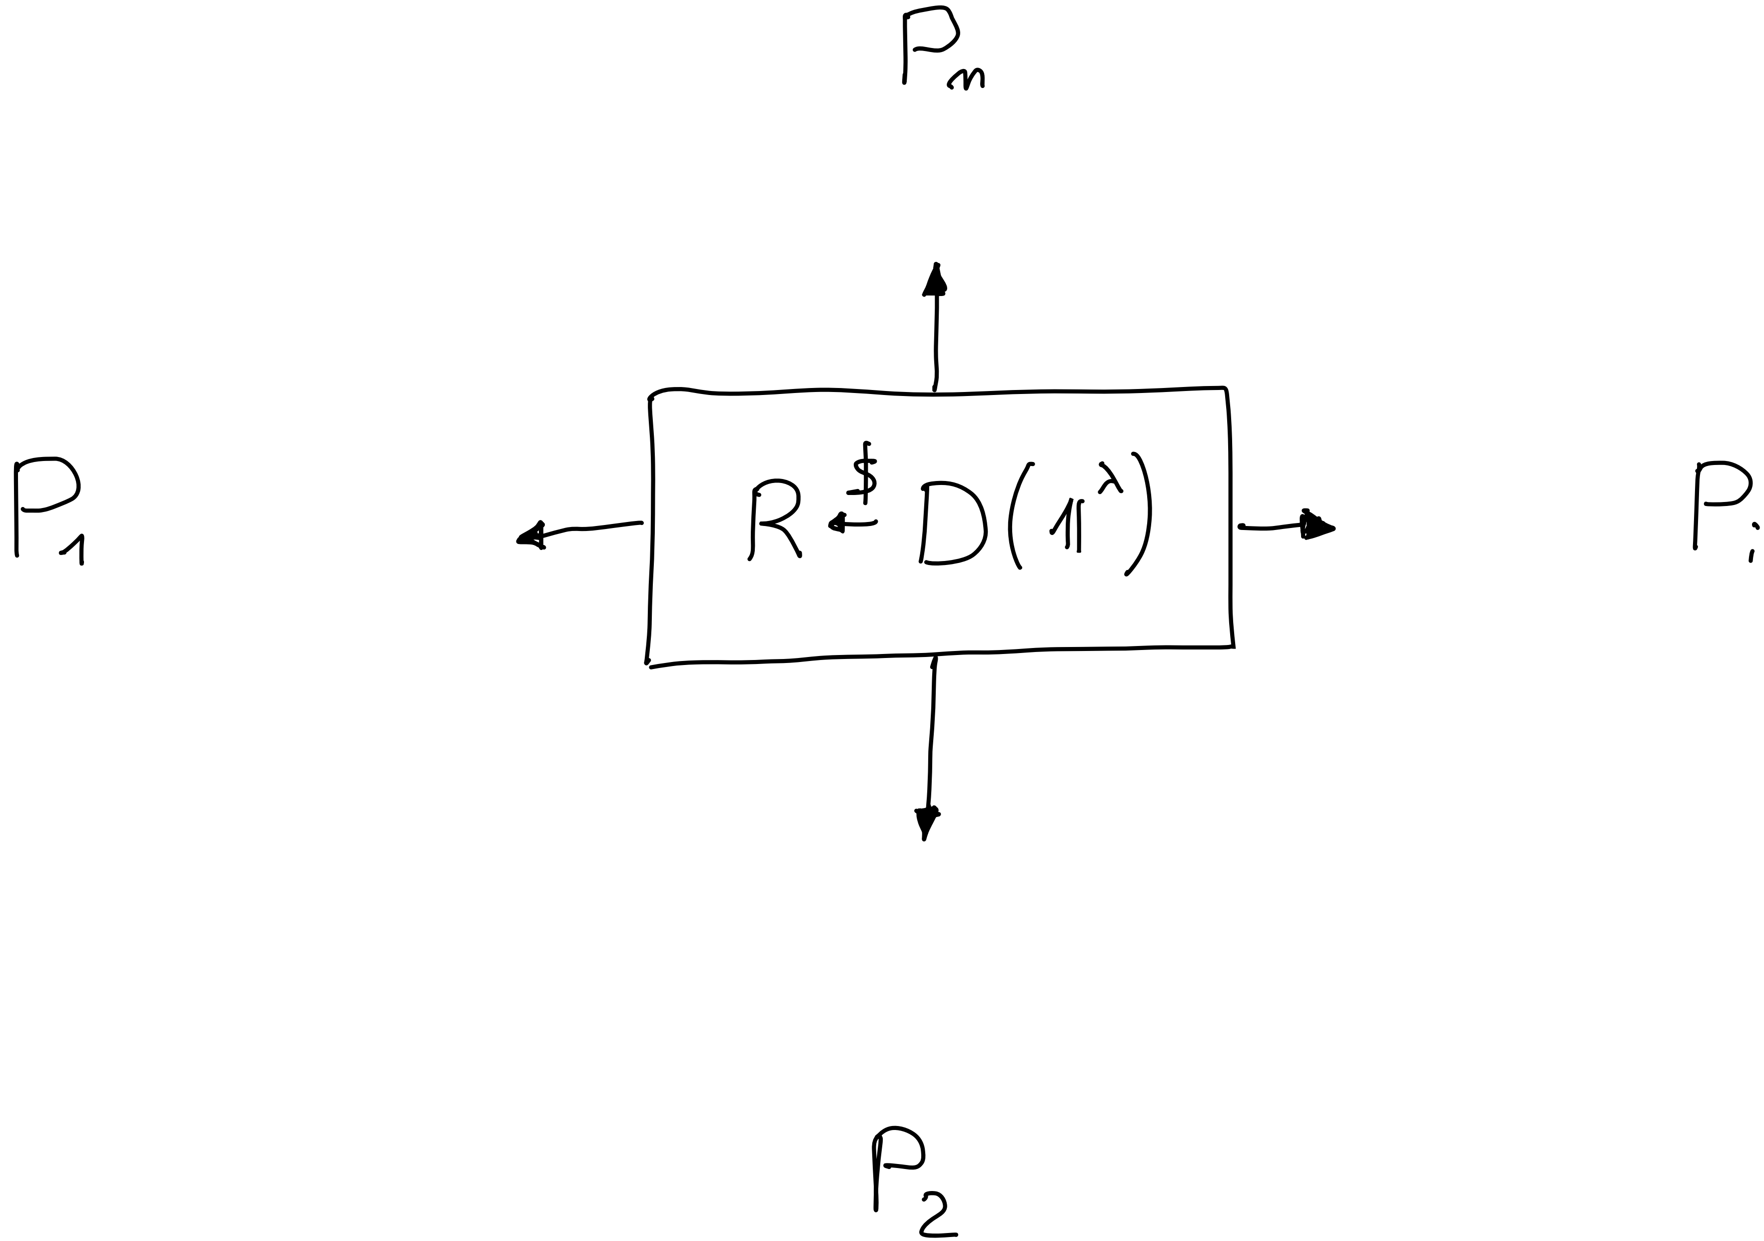

DISTRIBUTED SAMPLERS

[EC: ABRAM, SCHOLL, YAKOUBOV 22]

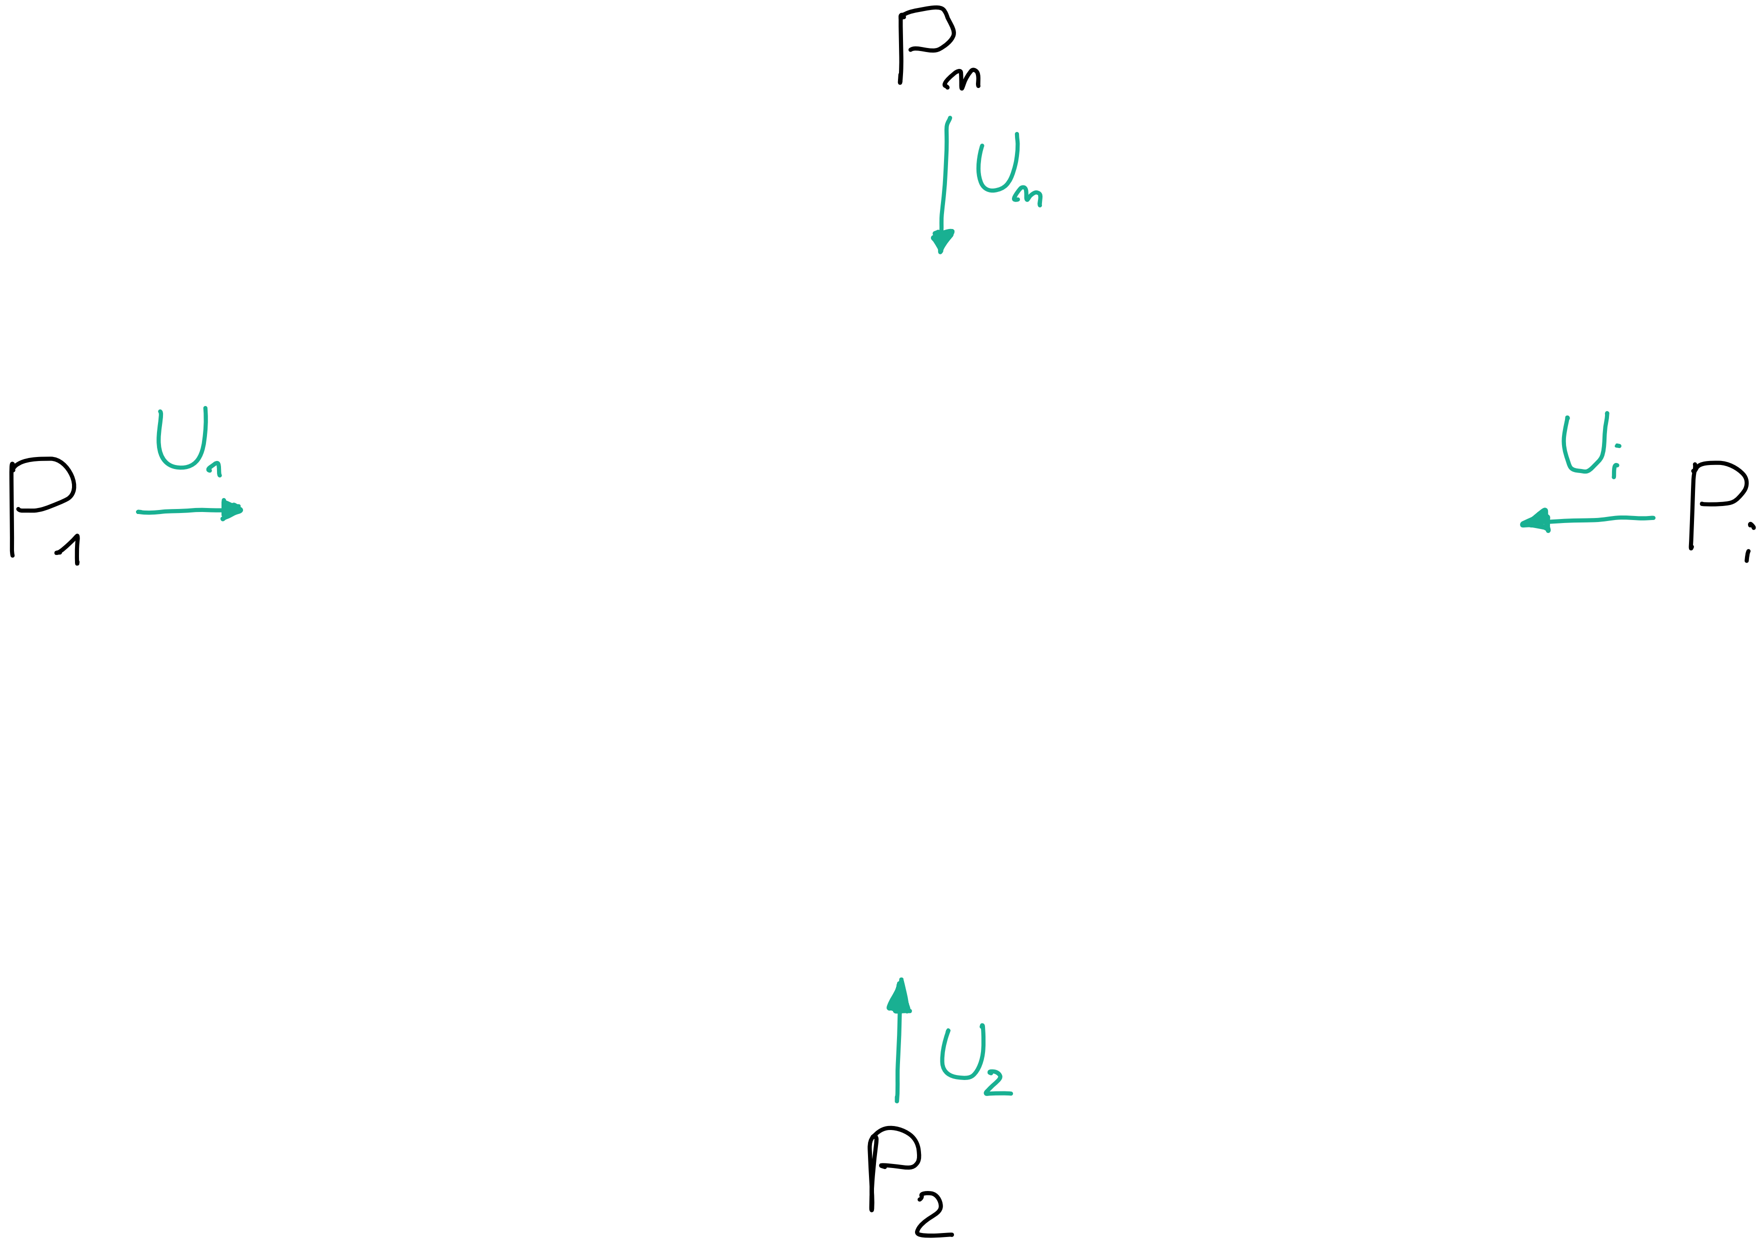

DISTRIBUTED SAMPLERS

[EC: Abram, Scholl, Yakoubov 22]

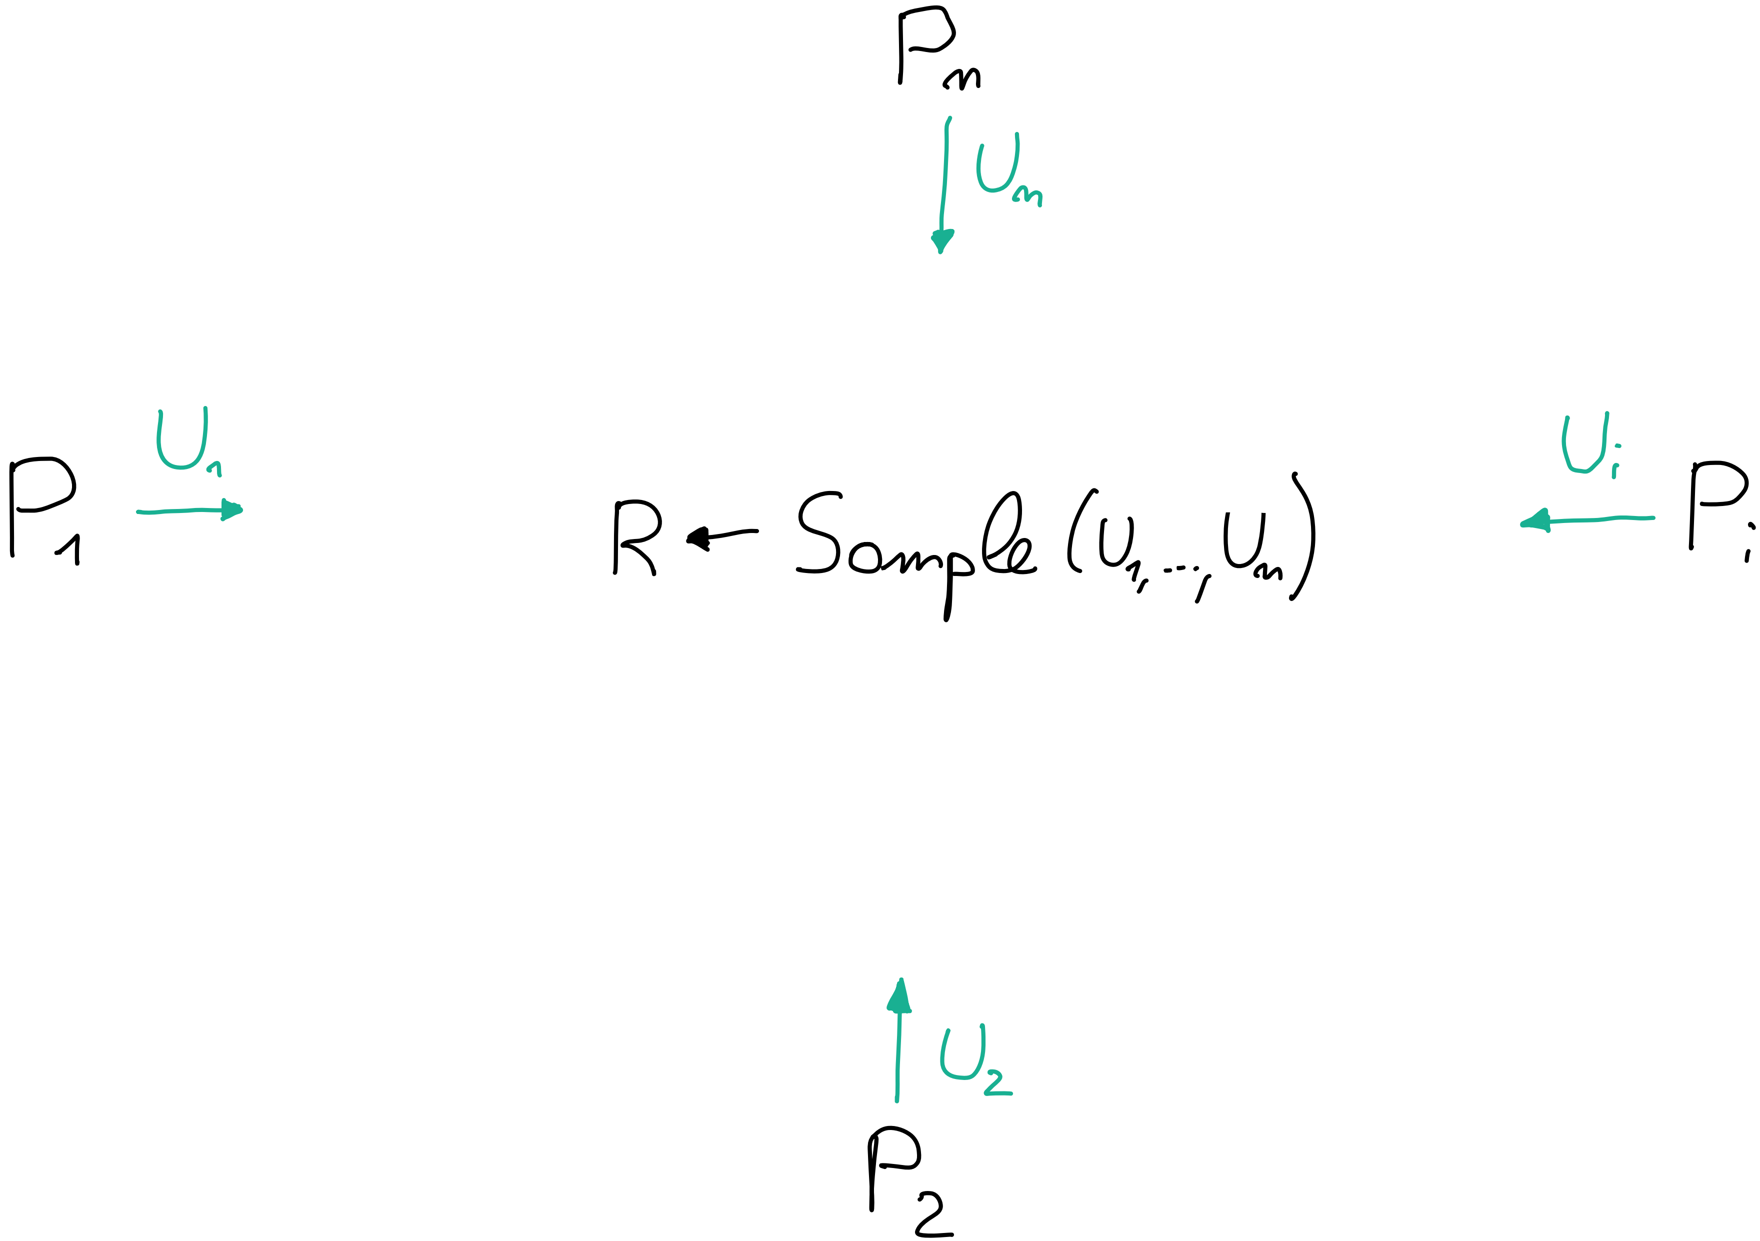

PREVIOUS WORK

POSITIVE RESULTS

[EC: ABRAM, SCHOLL, YAKOUBOV22]

SEMI-HONEST DISTRIBUTED SAMPLERS

- any efficient distribution $D(1^\lambda)$
- in the plain model
- dishonest majority
- iO + multi-Key FHE

PREVIOUS WORK

POSITIVE RESULTS

[EC: ABRAM, SCHOLL, YAKOUBOV22]

ACTIVE SECURITY?

PREVIOUS WORK

POSITIVE RESULTS

[EC: ABRAM, SCHOLL, YAKOUBOV22]

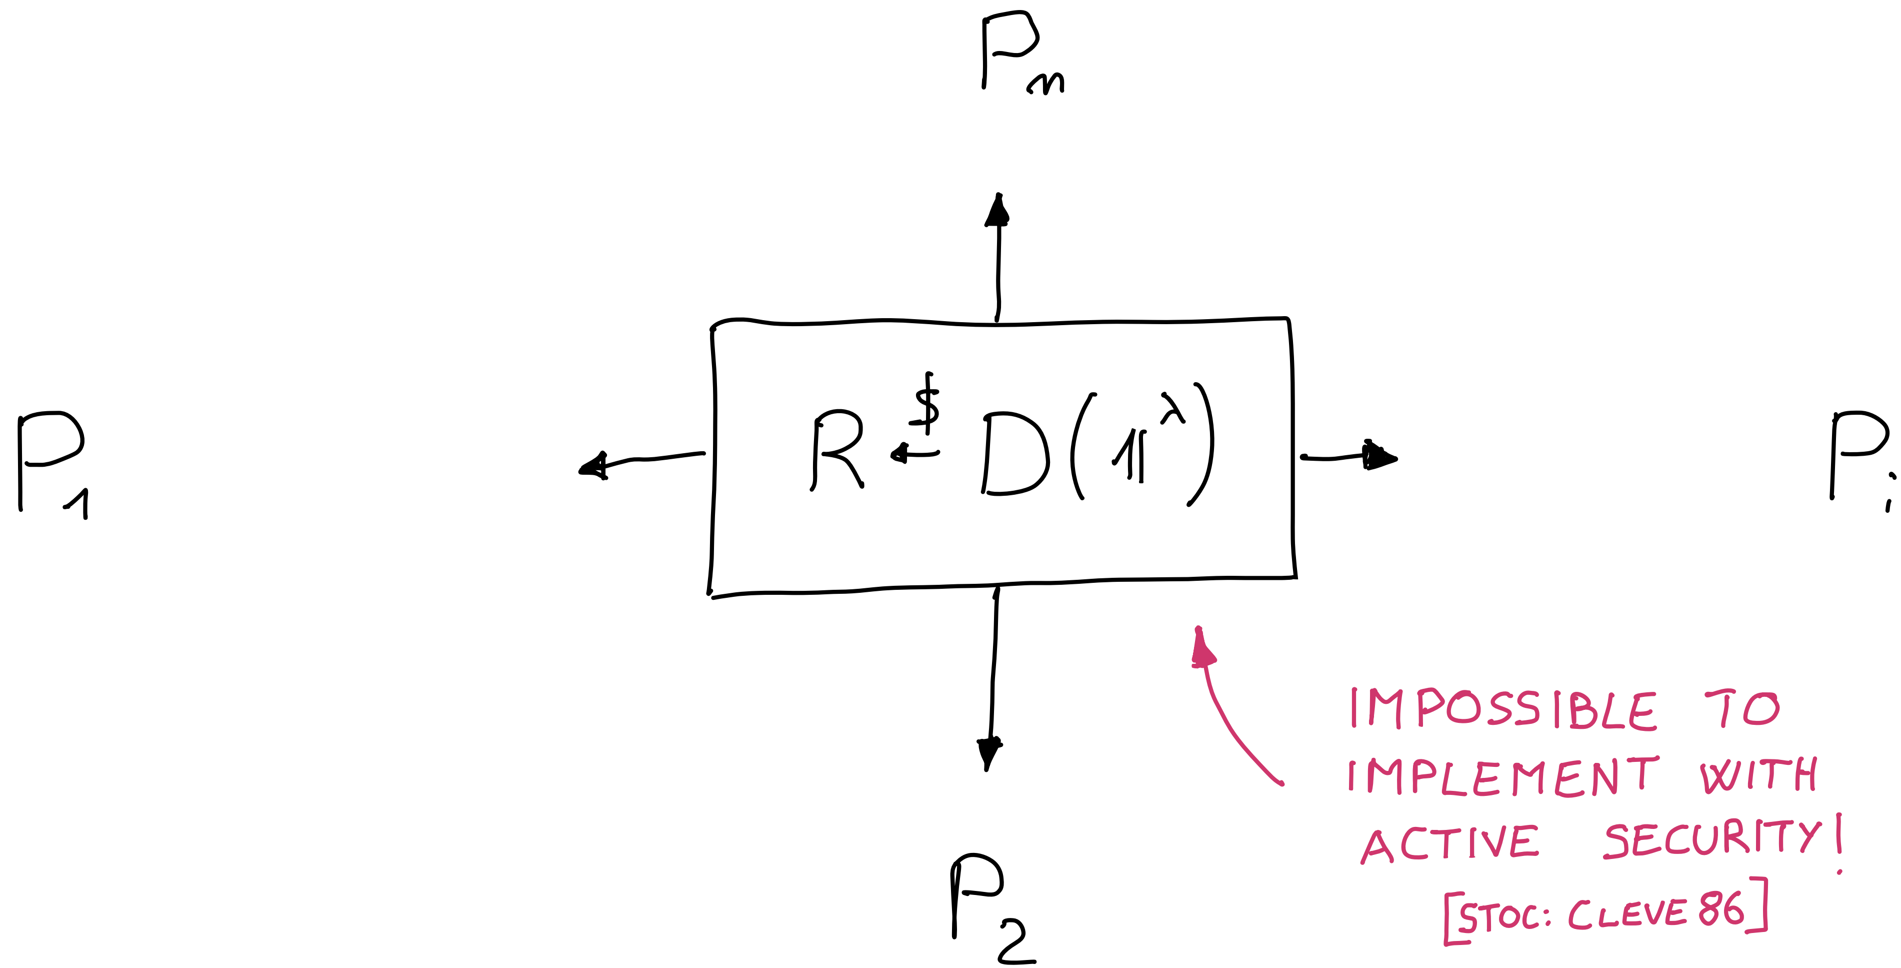

IMPOSSIBLE TO IMPLEMENT WITH ACTIVE SECURITY!

[STOC: CLEVE86]

PREVIOUS WORK

POSITIVE RESULTS

[EC: ABRAM, SCHOLL, YAKOUBOV22]

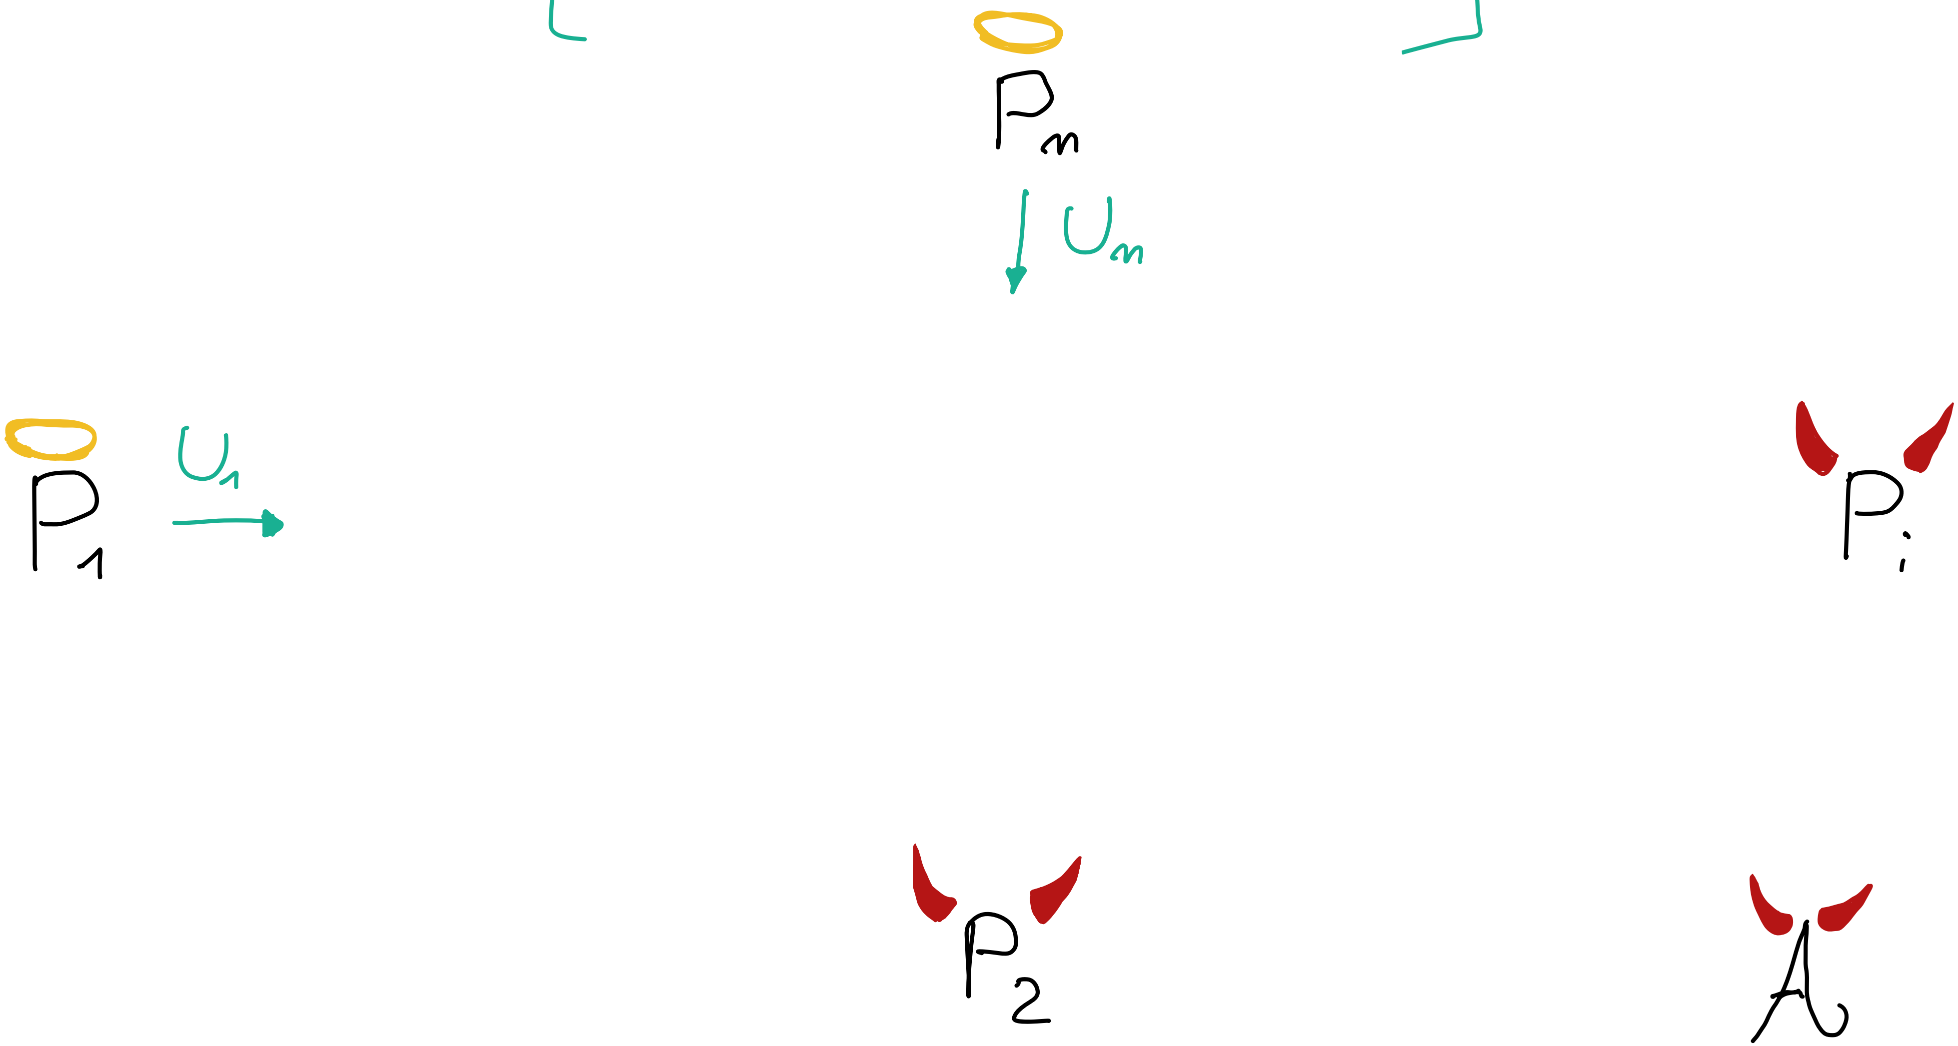

PREVIOUS WORK

POSITIVE RESULTS

[EC: ABRAM, SCHOLL, YAKOUBOV22]

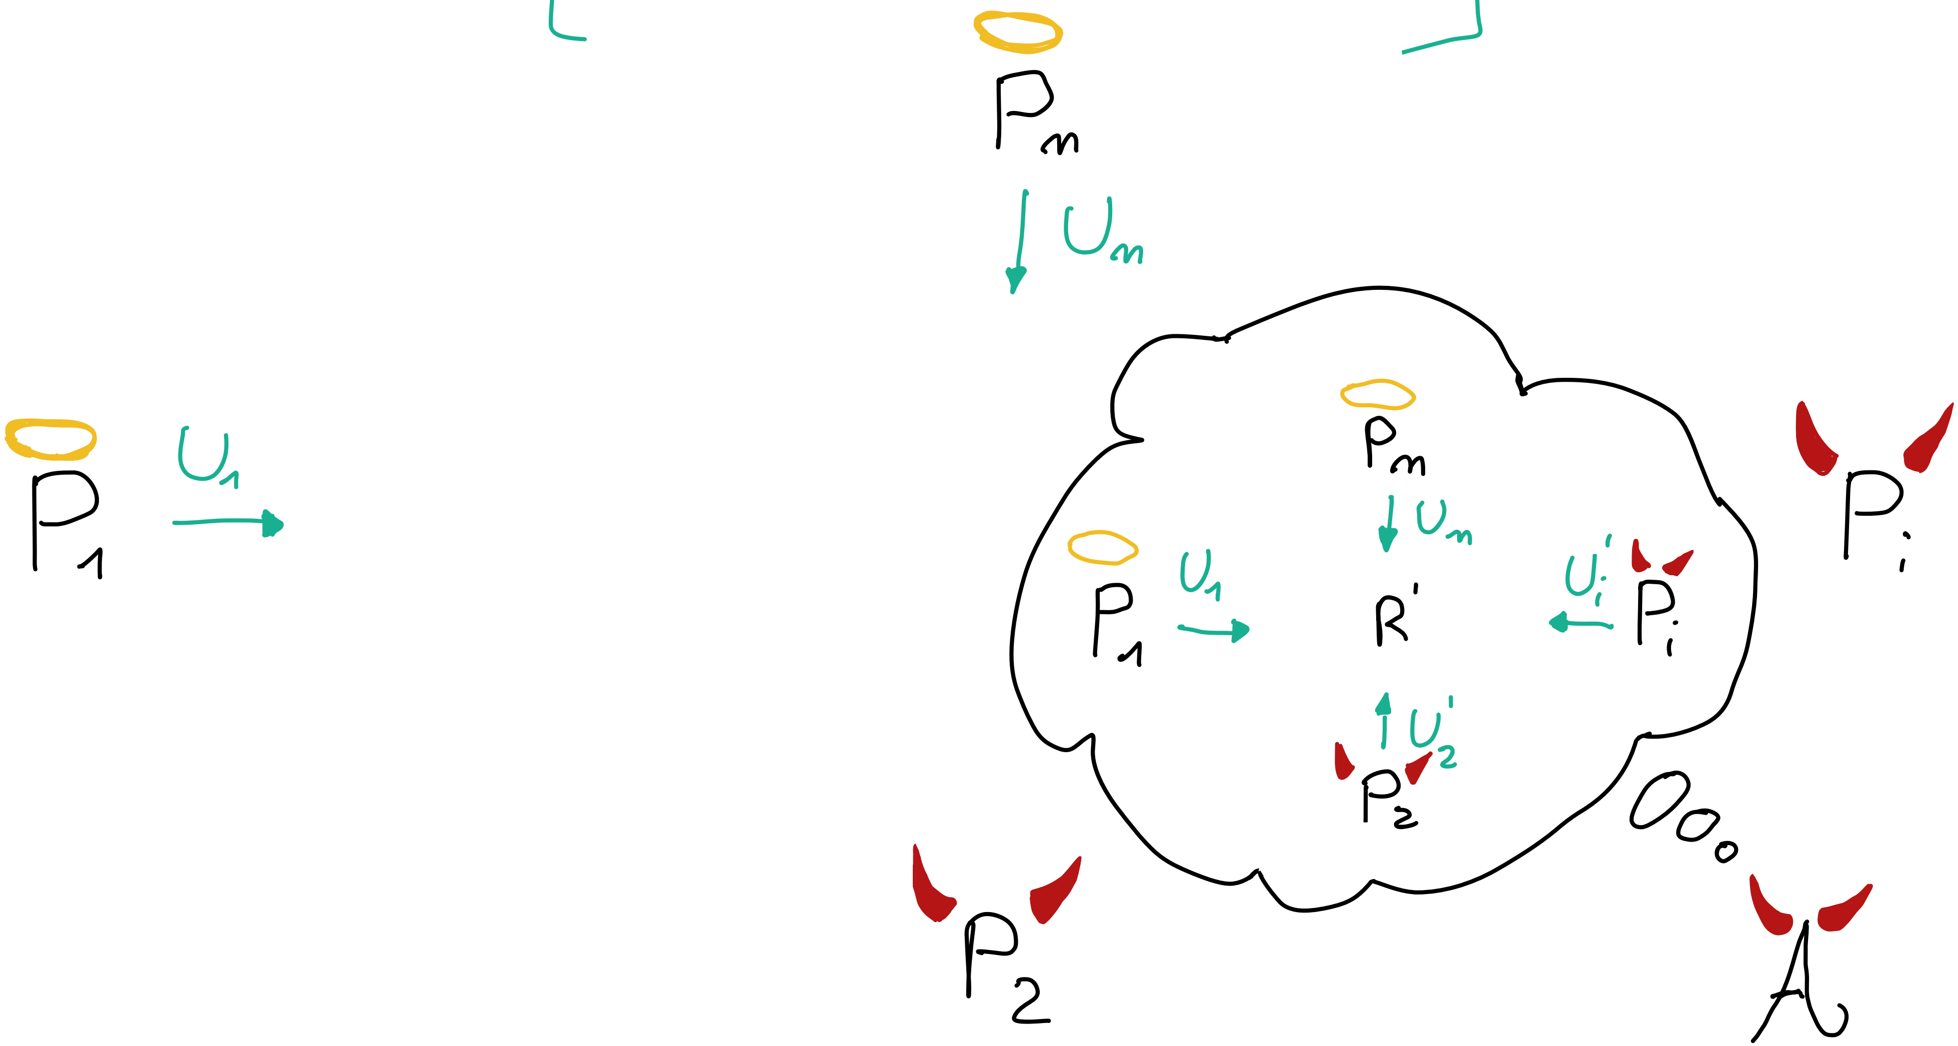

PREVIOUS WORK

POSITIVE RESULTS

[EC: Abram, Scholl, Yakoubov22]

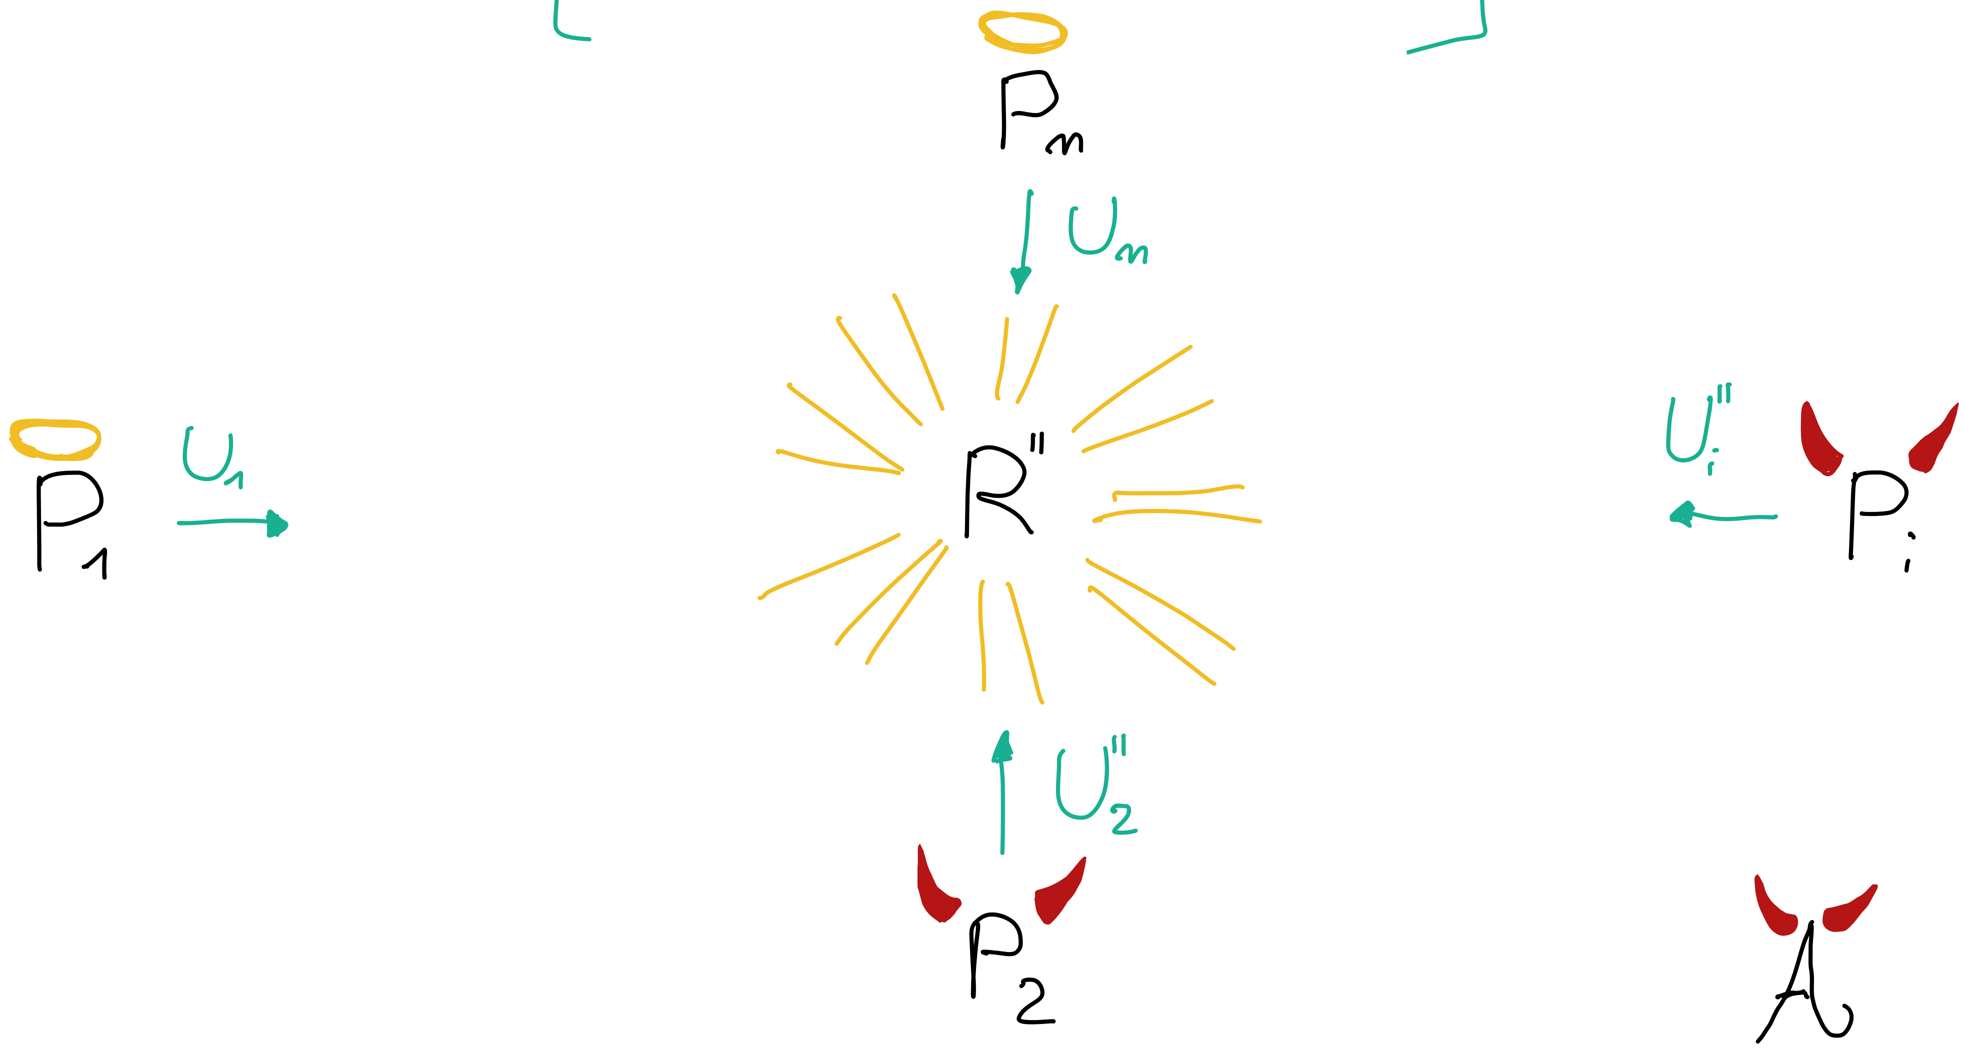

PREVIOUS WORK

POSITIVE RESULTS

[EC: Abram, Scholl, Yakoubov22]

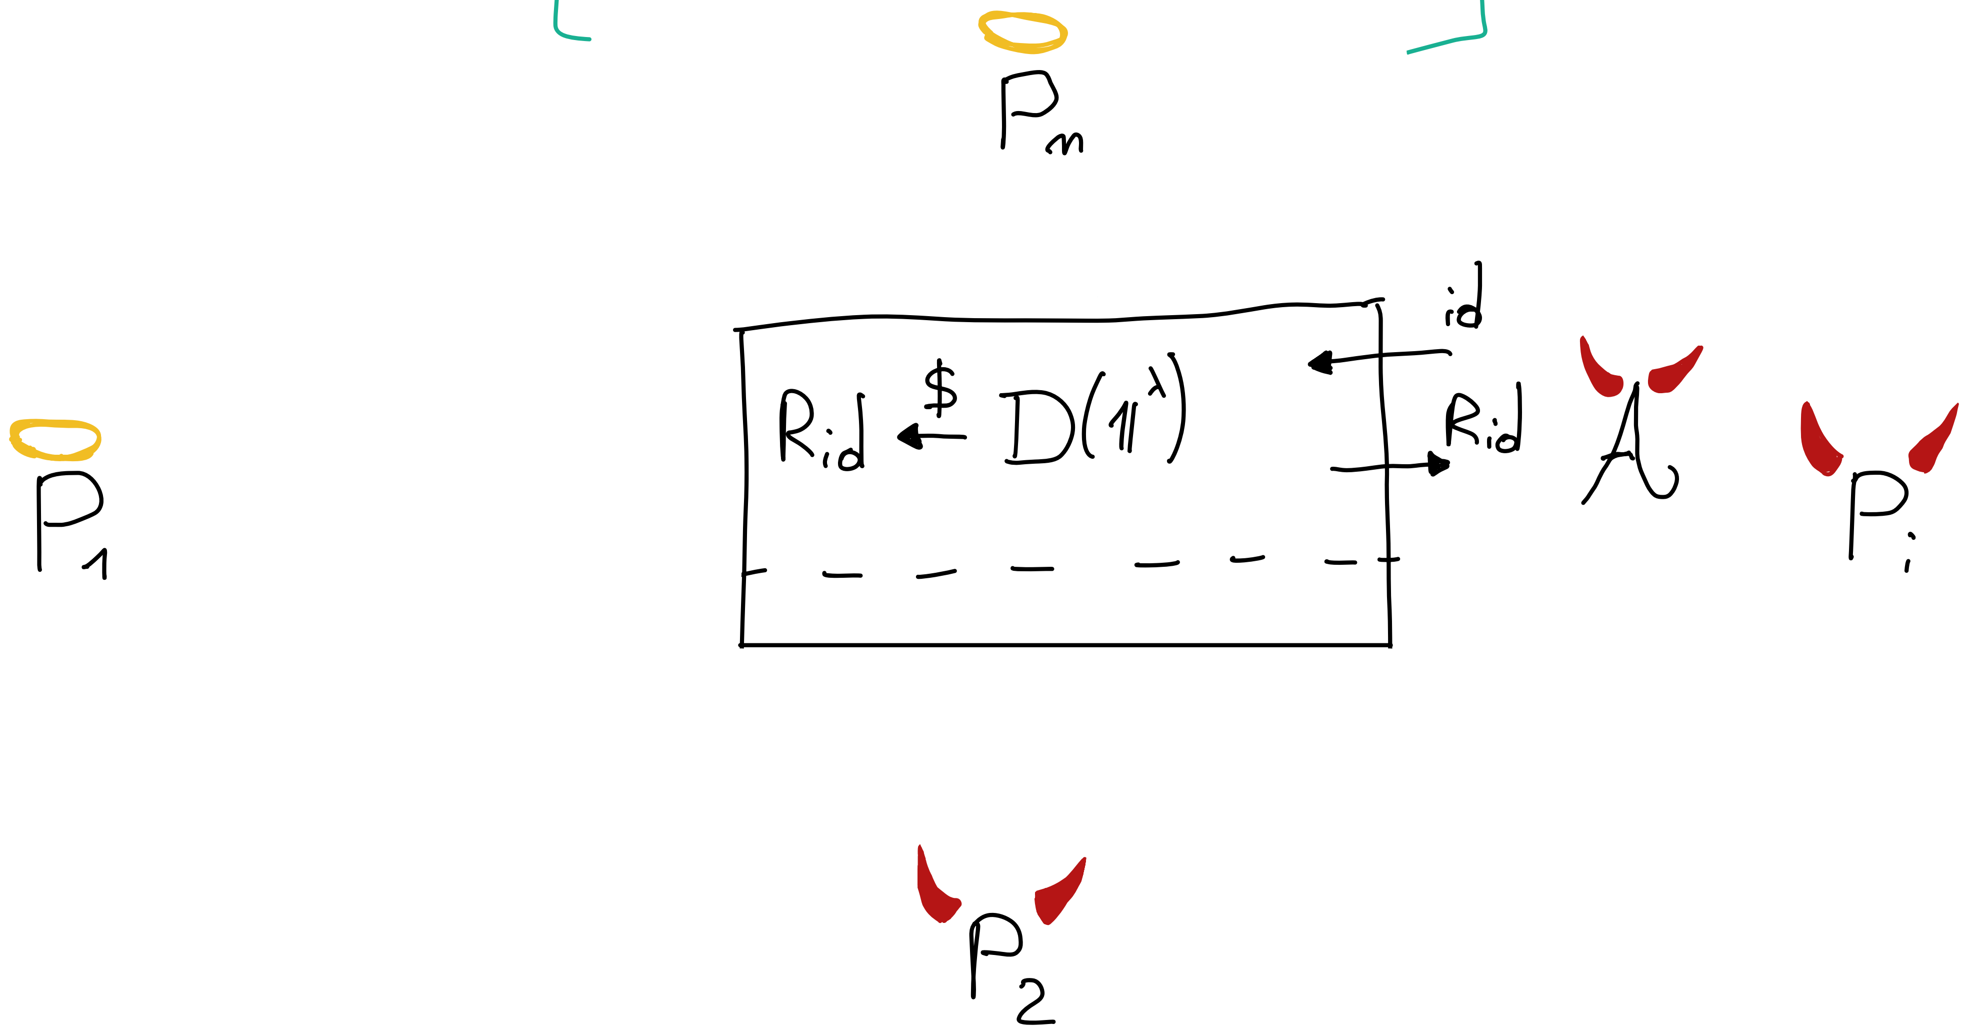

PREVIOUS WORK

POSITIVE RESULTS

[EC: ABRAM, SCHOLL, YAKOUBOV22]

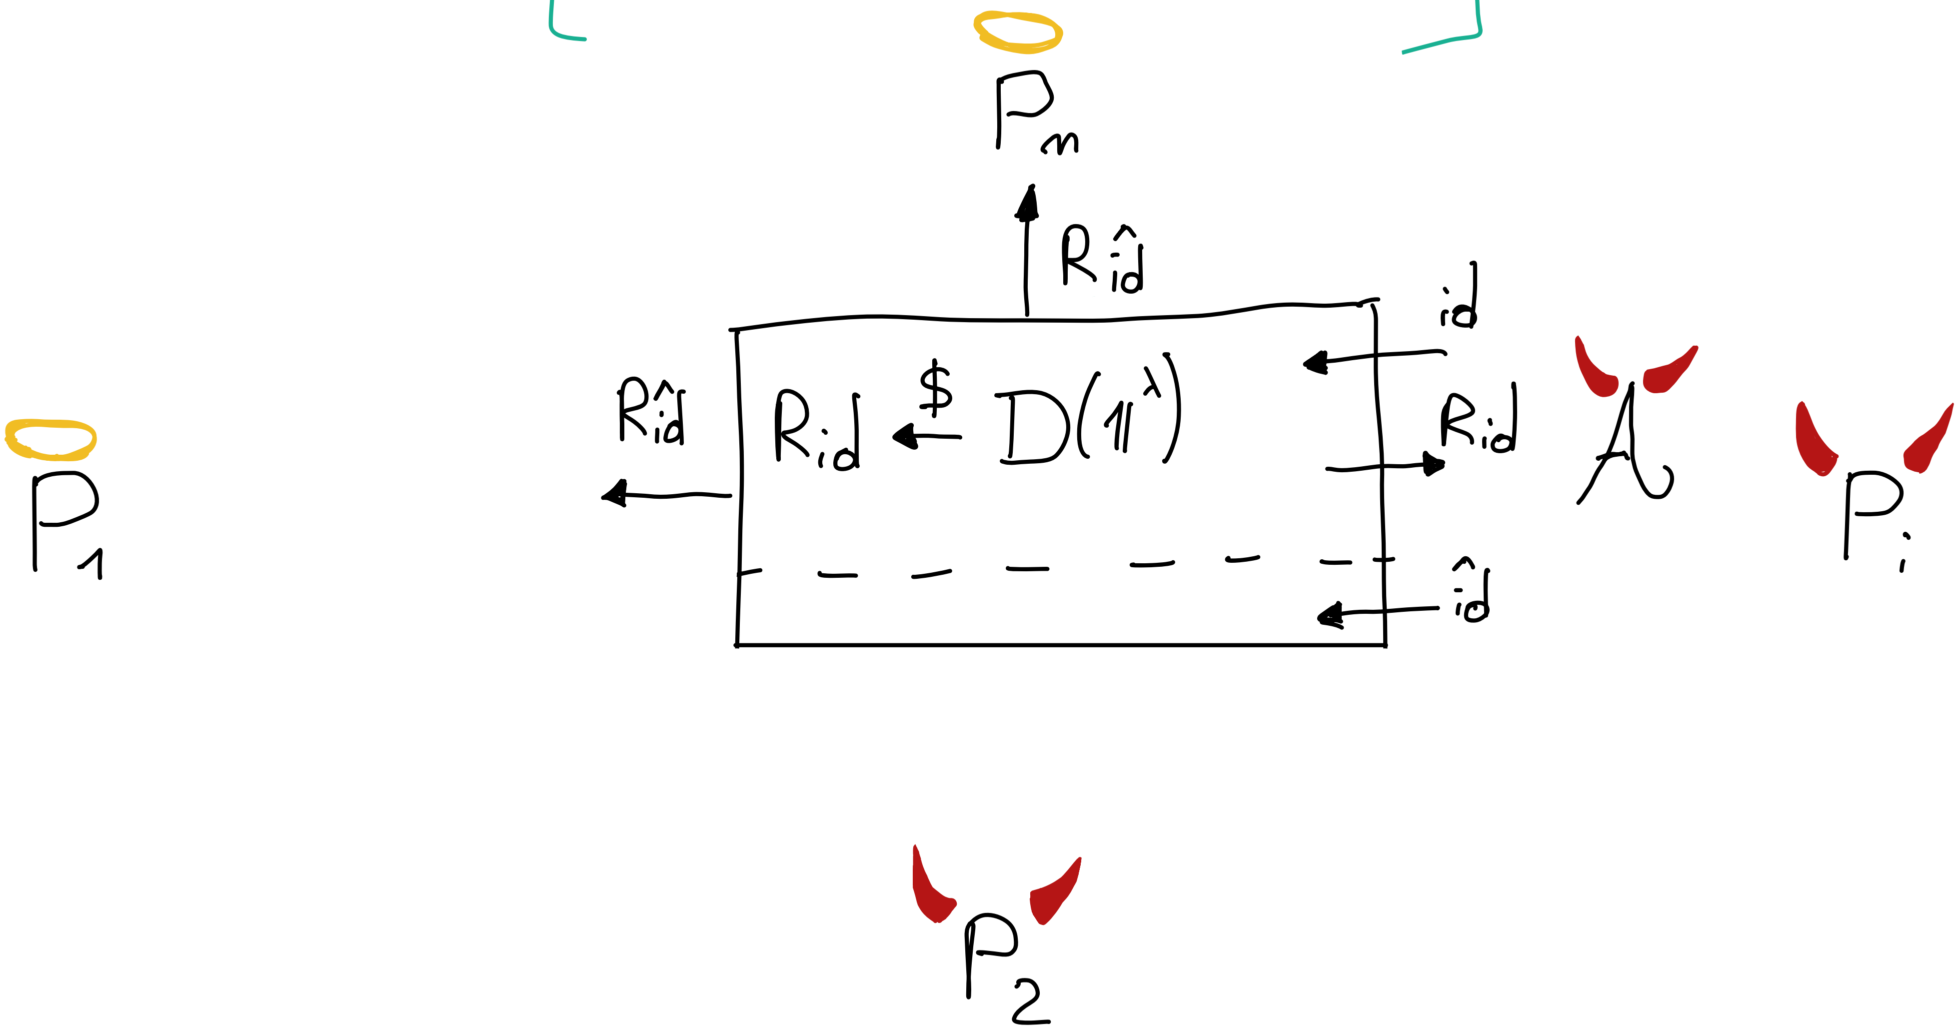

PREVIOUS WORK

POSITIVE RESULTS

[EC: ABRAM, SCHOLL, YAKOUBOV22]

ACTIVE DISTRIBUTED SAMPLERS

- any efficient distribution $D(1^\lambda)$
- in the programmable RO model
- dishonest majority, static corruption
- iO + multi-Key FHE + NIZKs

PREVIOUS WORK

NEGATIVE RESULTS

[EPRINT: ABRAM, OBREMSKI, SCHOLL 23]

PREVIOUS WORK

NEGATIVE RESULTS

[EPRINT: ABRAM, OBREMSKI, SCHOLL 23]

THEOREM

Suppose that $H_\infty(D) = \omega(\log \lambda)$.

Then, any actively secure distributed sampler for $D(1^\lambda)$ needs a CRS.

PREVIOUS WORK

NEGATIVE RESULTS

[EPRINT: ABRAM, OBREMSKI, SCHOLL 23]

THEOREM

Suppose that $H_\infty(D) = \omega(\log \lambda)$.

Then, any actively secure distributed sampler for $D(1^\lambda)$ needs a CRS. Furthermore, the CRS is:

- non-reusable

PREVIOUS WORK

NEGATIVE RESULTS

[EPRINT: ABRAM, OBREMSKI, SCHOLL 23]

THEOREM

Suppose that $H_\infty(D) = \omega(\log \lambda)$.

Then, any actively secure distributed sampler for $D(1^\lambda)$ needs a CRS. Furthermore, the CRS is:

- non-reusable
- at least $H_{YAO}(D) - O(\log \lambda)$ bits long

PREVIOUS WORK

NEGATIVE RESULTS

[EPRINT: ABRAM, OBREMSKI, SCHOLL 23]

THEOREM

Suppose that $H_\infty(D) = \omega(\log \lambda)$.

Then, any actively secure distributed sampler for $D(1^\lambda)$ needs a CRS. Furthermore, the CRS is:

- non-reusable
- at least $H_{YAO}(D) - O(\log \lambda)$ bits long
- unstructured only if D is obliviously samplable

PREVIOUS WORK

NEGATIVE RESULTS
[EPRINT: ABRAM, OBREMSKI, SCHOLL 23]

THEOREM

Suppose that $H_\infty(D) = \omega(\log \lambda)$.

Then, any actively secure distributed sampler for $D(1^\lambda)$ needs a CRS. Furthermore, the CRS is:

- non-reusable
- at least $H_{YAO}(D) - O(\log \lambda)$ bits long
- unstructured only if D is obliviously samplable

SUMMARY

WITHOUT RANDOM ORACLE, ACTIVELY SECURE DISTRIBUTED SAMPLERS CANNOT BE BETTER THAN THE TRUSTED SETUP!

OUR CONTRIBUTION

NEW DEFINITIONS OF ACTIVE DISTRIBUTED SAMPLERS THAT DON'T NEED RANDOM ORACLES

OUR CONTRIBUTION

NEW DEFINITIONS OF ACTIVE DISTRIBUTED SAMPLERS THAT DON'T NEED RANDOM ORACLES

HARDNESS-PRESERVING DISTRIBUTED SAMPLERS

INDISTINGUISHABILITY-PRESERVING DISTRIBUTED SAMPLERS

OUR CONTRIBUTION

NEW DEFINITIONS OF ACTIVE DISTRIBUTED SAMPLERS THAT DON'T NEED RANDOM ORACLES

HARDNESS-PRESERVING DISTRIBUTED SAMPLERS

preserving the hardness of search games with efficient challenger.

INDISTINGUISHABILITY-PRESERVING DISTRIBUTED SAMPLERS

OUR CONTRIBUTION

NEW DEFINITIONS OF ACTIVE DISTRIBUTED SAMPLERS THAT DON'T NEED RANDOM ORACLES

HARDNESS-PRESERVING DISTRIBUTED SAMPLERS

preserving the hardness of search games with efficient challenger.

INDISTINGUISHABILITY-PRESERVING DISTRIBUTED SAMPLERS

preserving the functionality of the compiled protocol if certain conditions are satisfied.

HARDNESS-PRESERVING DISTRIBUTED SAMPLERS

$\boxed{G}$

HARDNESS-PRESERVING DISTRIBUTED SAMPLERS

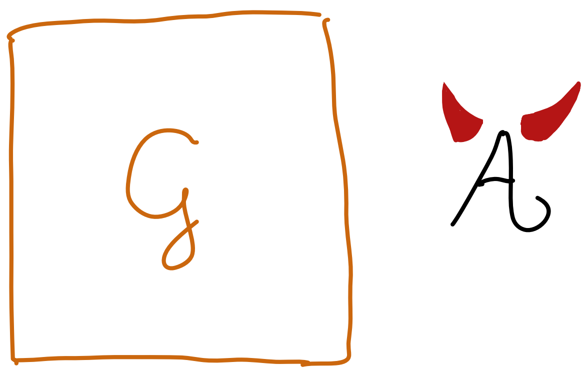

HARDNESS-PRESERVING DISTRIBUTED SAMPLERS

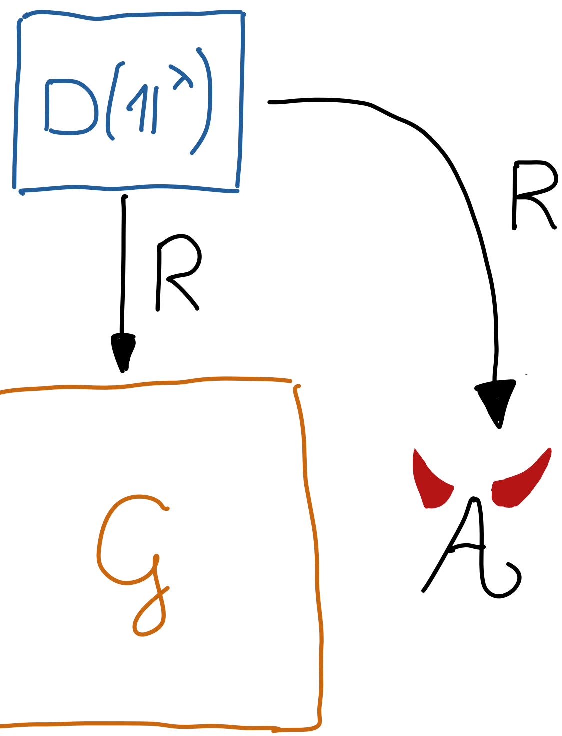

HARDNESS-PRESERVING DISTRIBUTED SAMPLERS

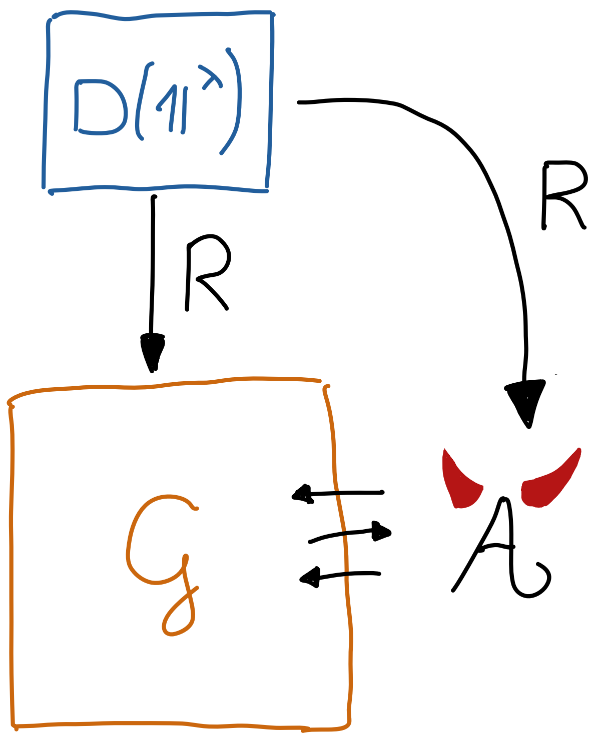

HARDNESS-PRESERVING DISTRIBUTED SAMPLERS

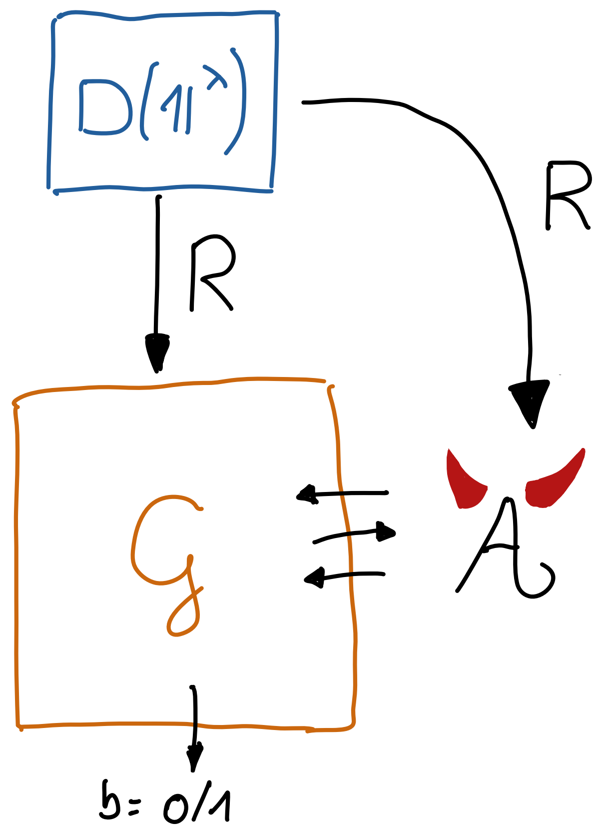

HARDNESS-PRESERVING DISTRIBUTED SAMPLERS

REAL WORLD	IDEAL WORLD

HARDNESS-PRESERVING DISTRIBUTED SAMPLERS

REAL WORLD

IDEAL WORLD

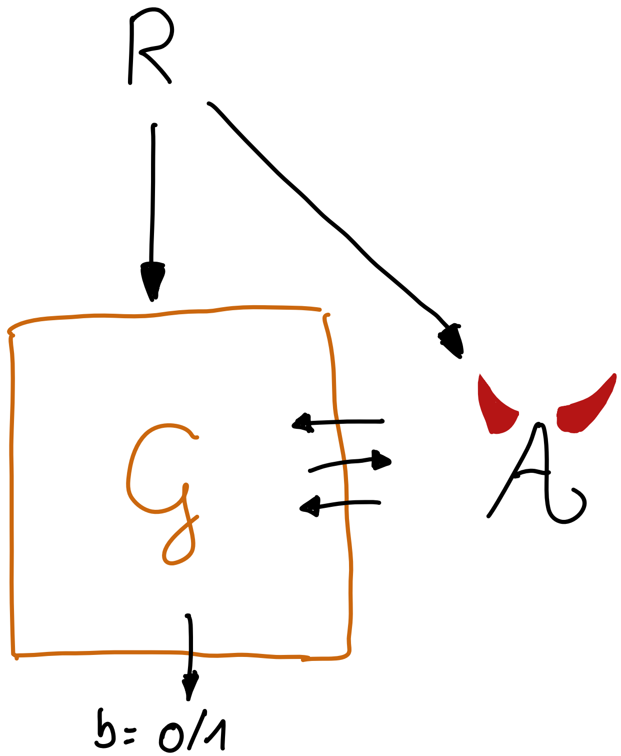

HARDNESS-PRESERVING DISTRIBUTED SAMPLERS

REAL WORLD

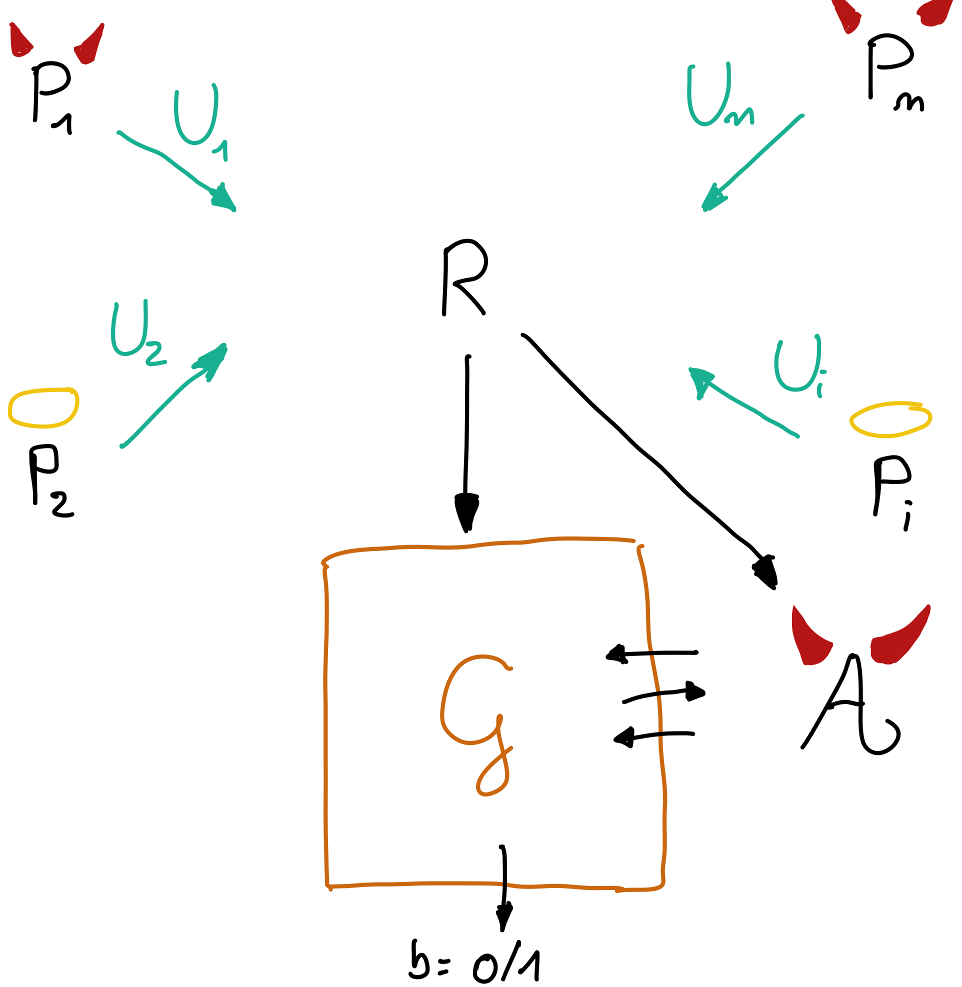

IDEAL WORLD

HARDNESS-PRESERVING DISTRIBUTED SAMPLERS

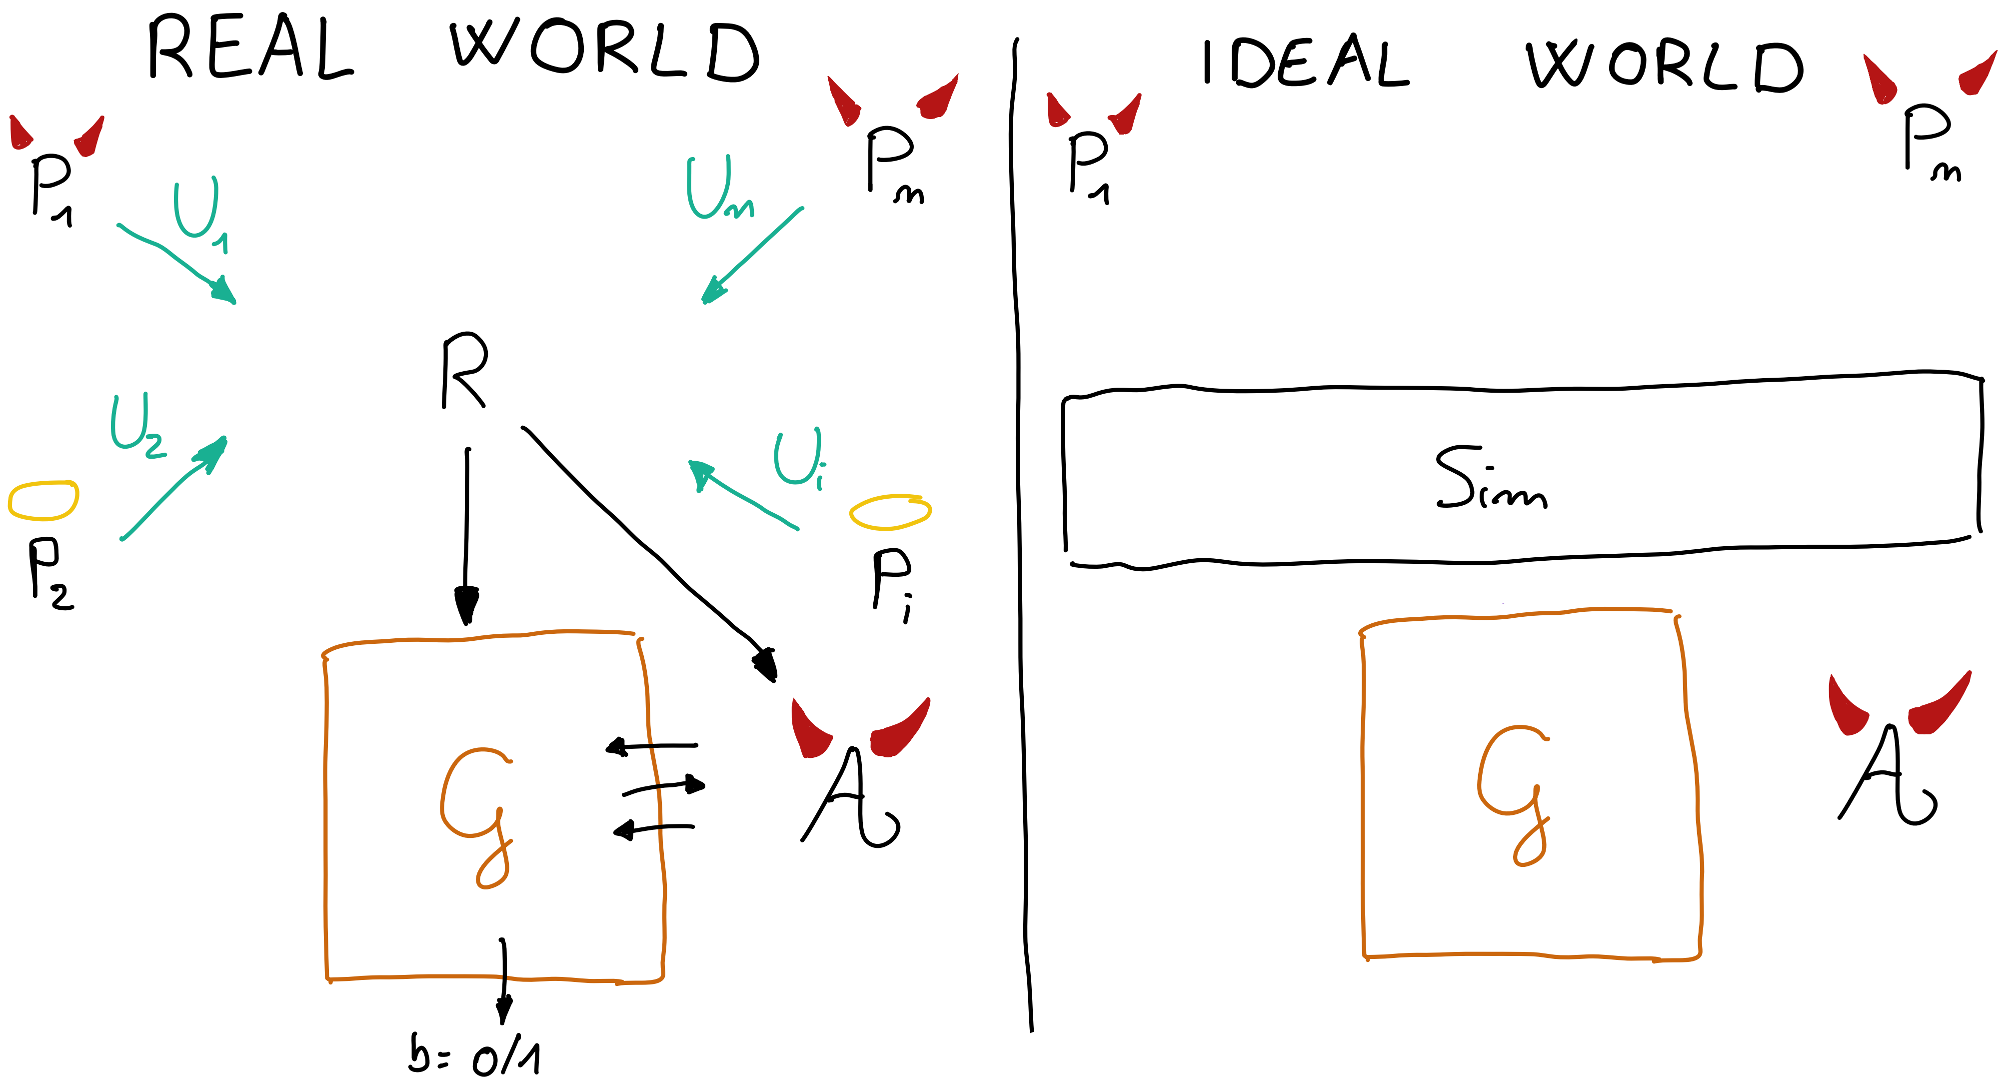

HARDNESS-PRESERVING DISTRIBUTED SAMPLERS

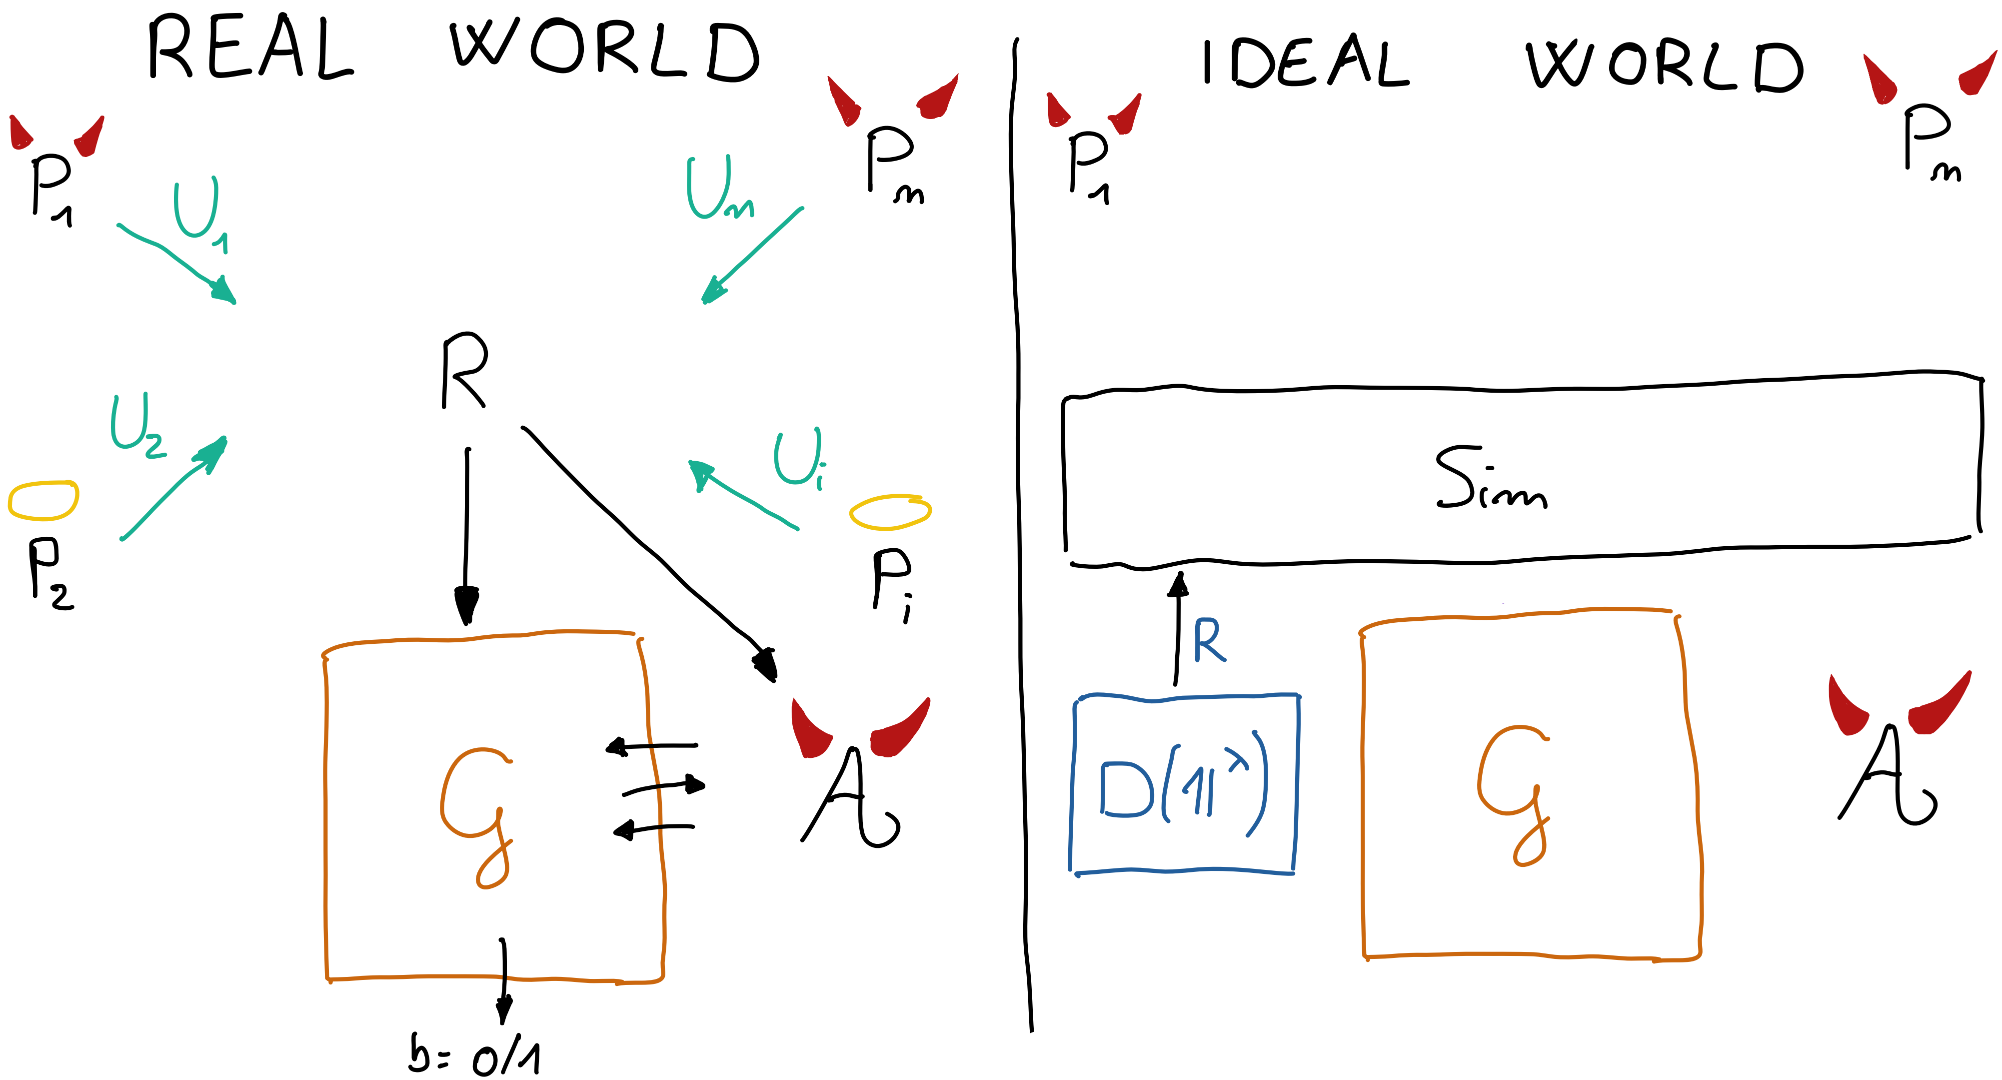

HARDNESS-PRESERVING DISTRIBUTED SAMPLERS

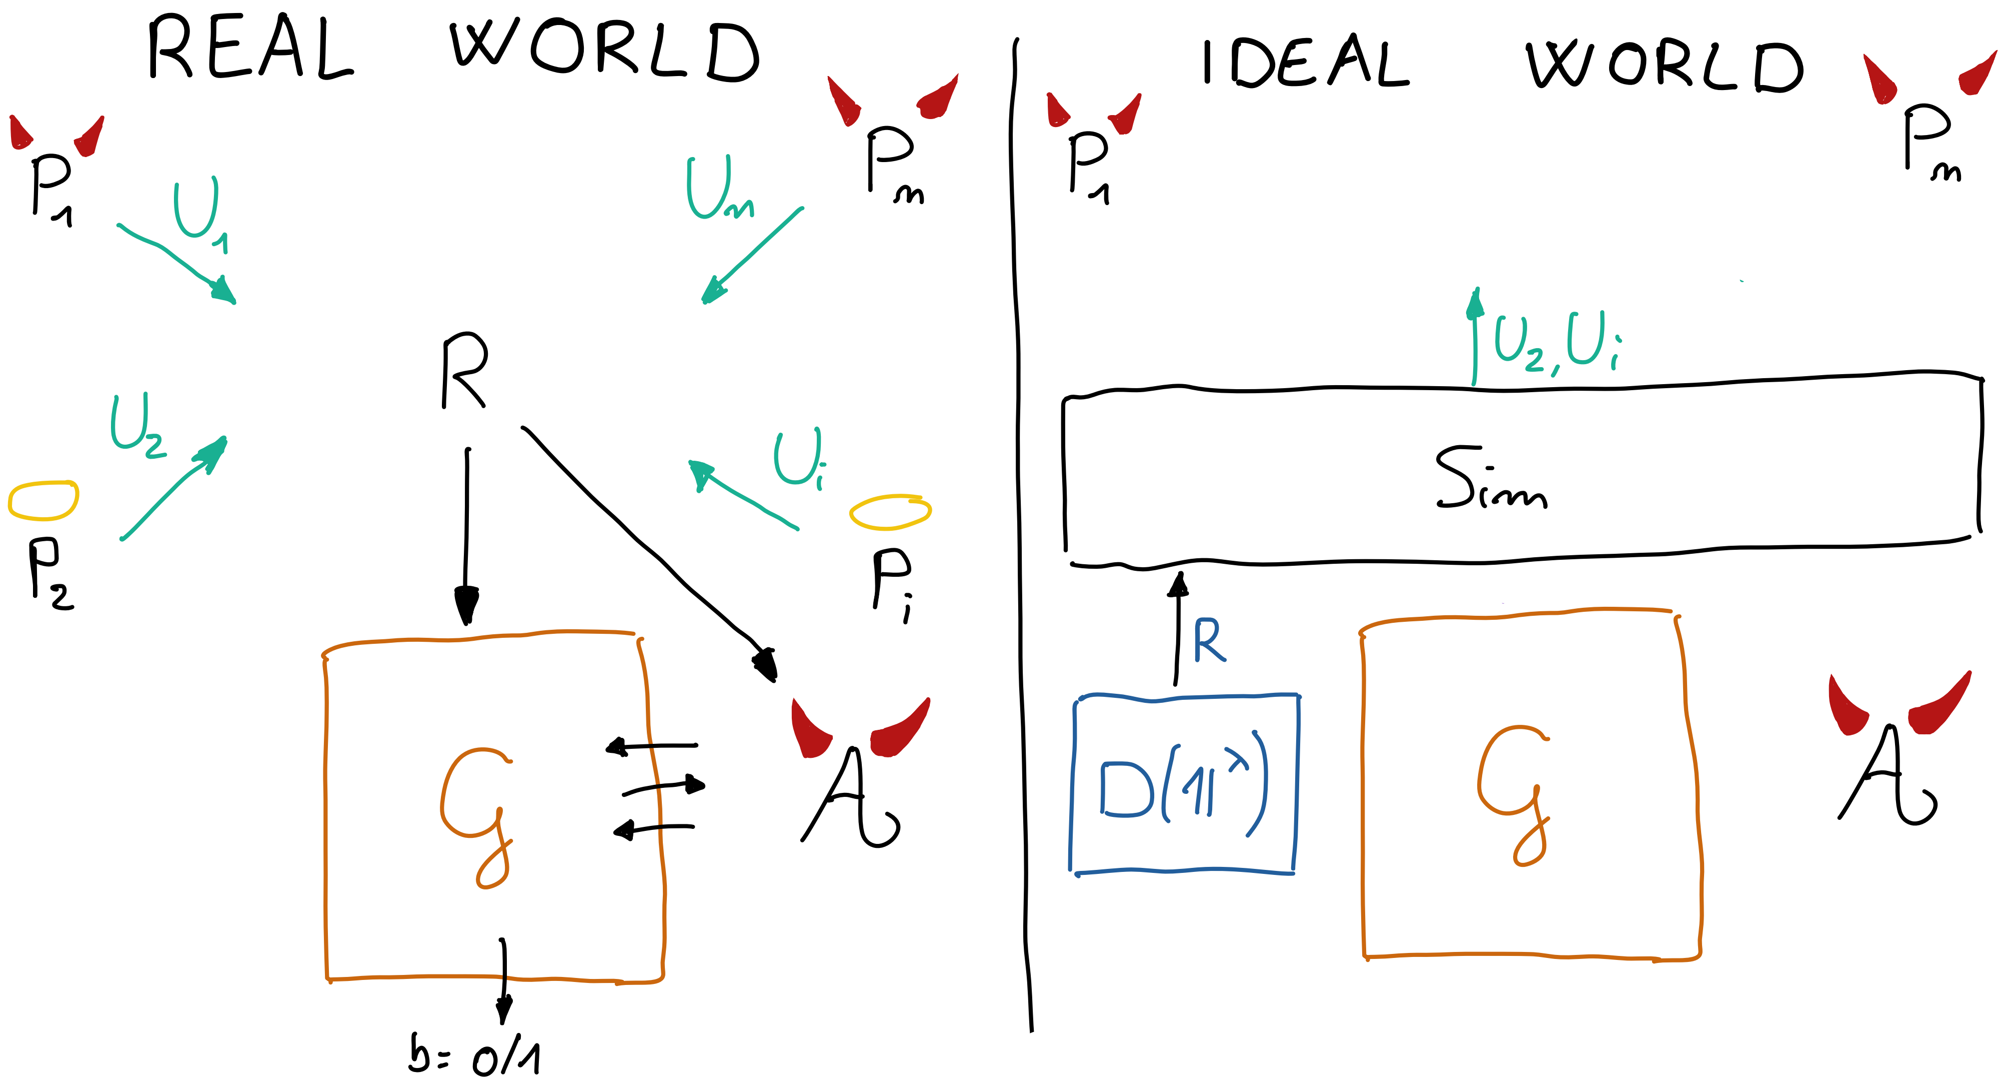

HARDNESS-PRESERVING DISTRIBUTED SAMPLERS

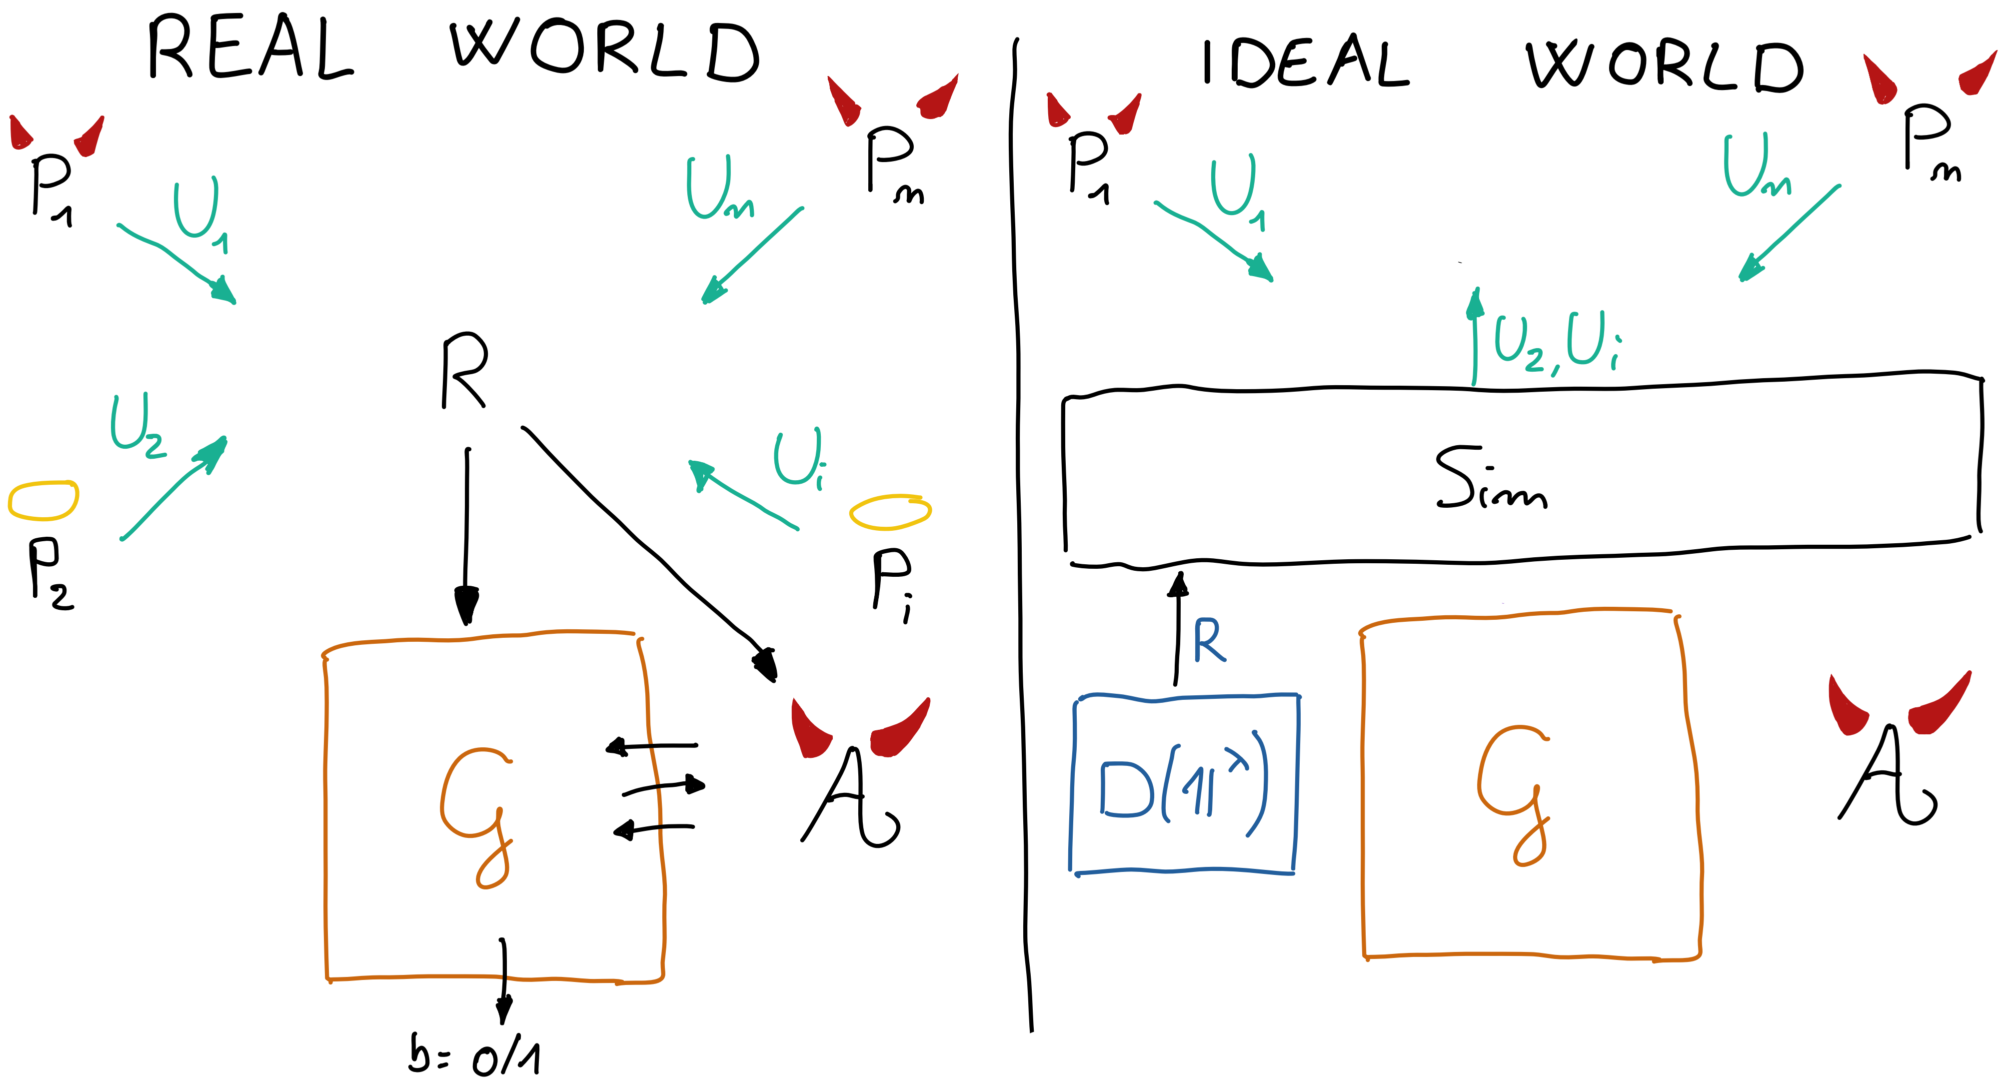

HARDNESS - PRESERVING DISTRIBUTED SAMPLERS

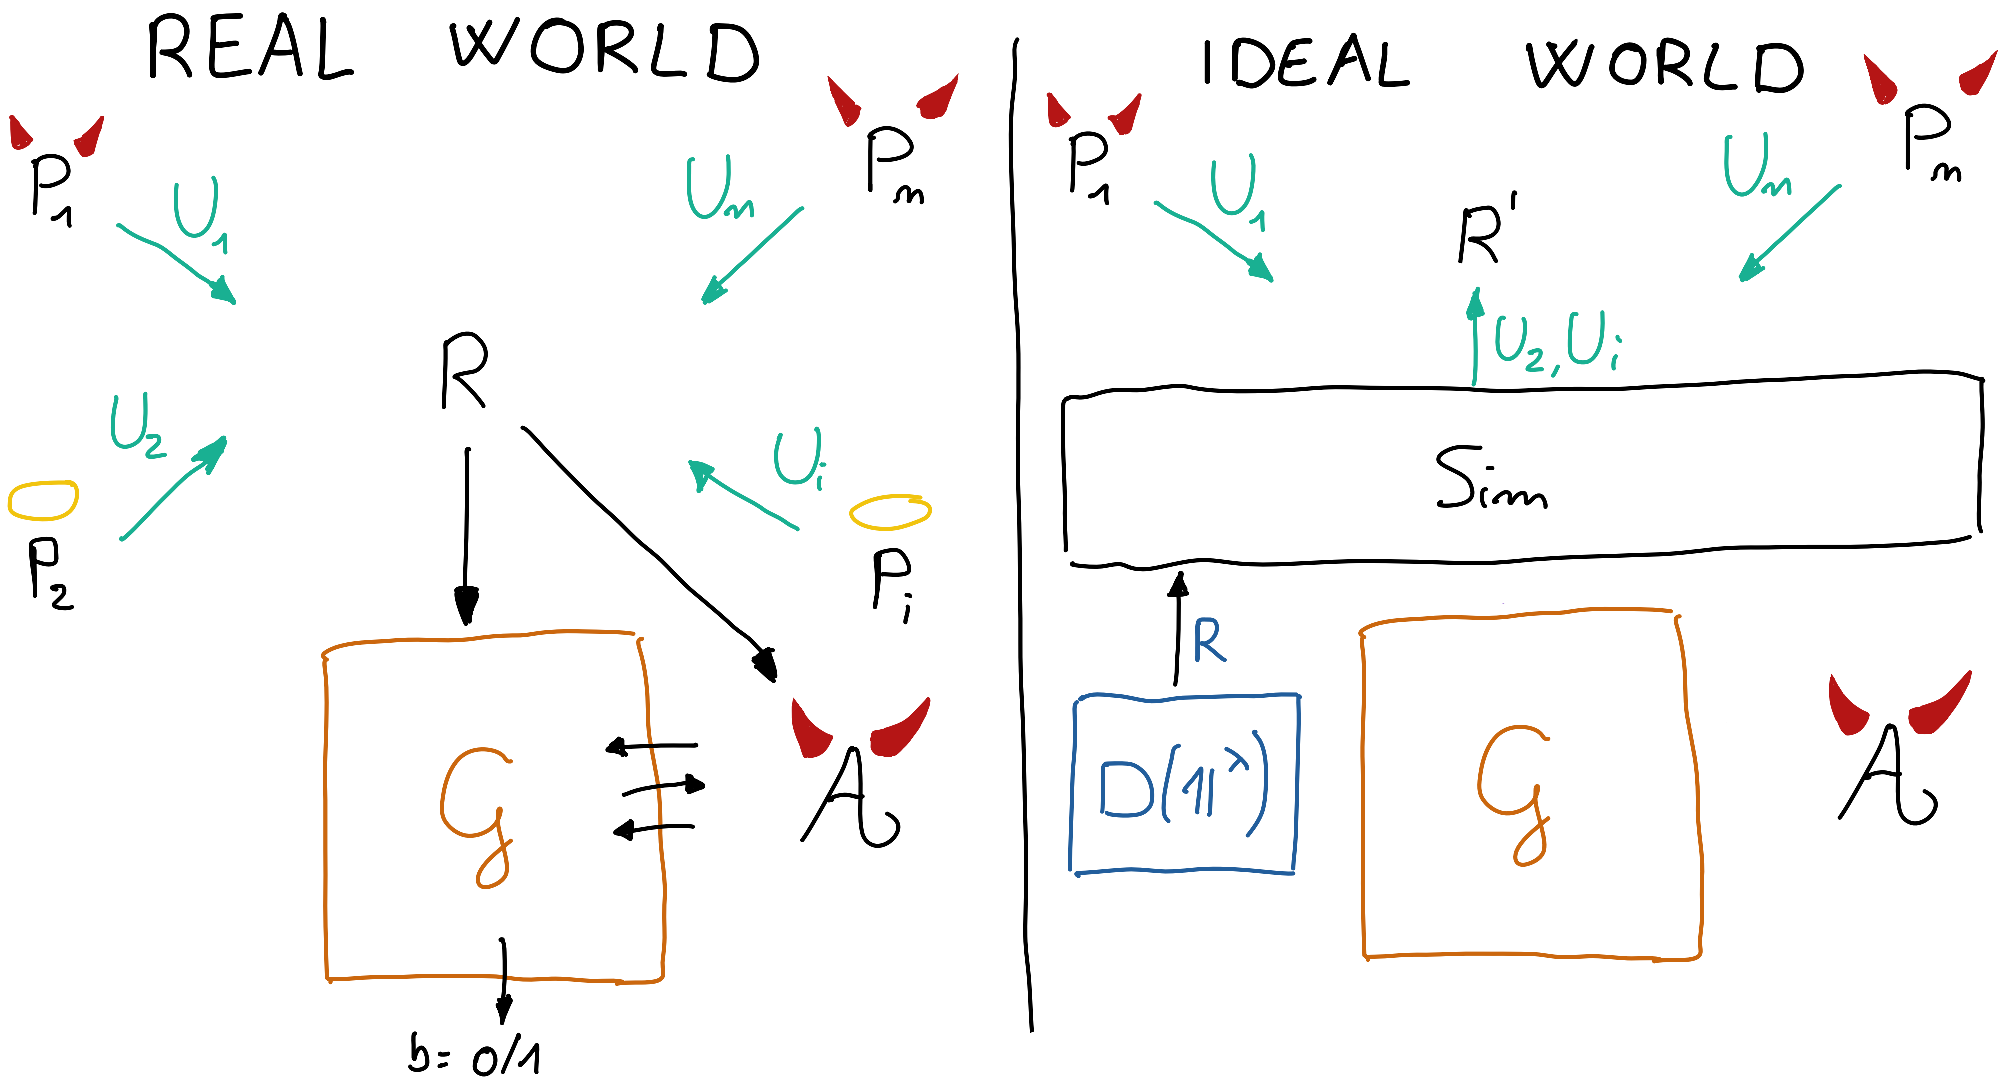

HARDNESS-PRESERVING DISTRIBUTED SAMPLERS

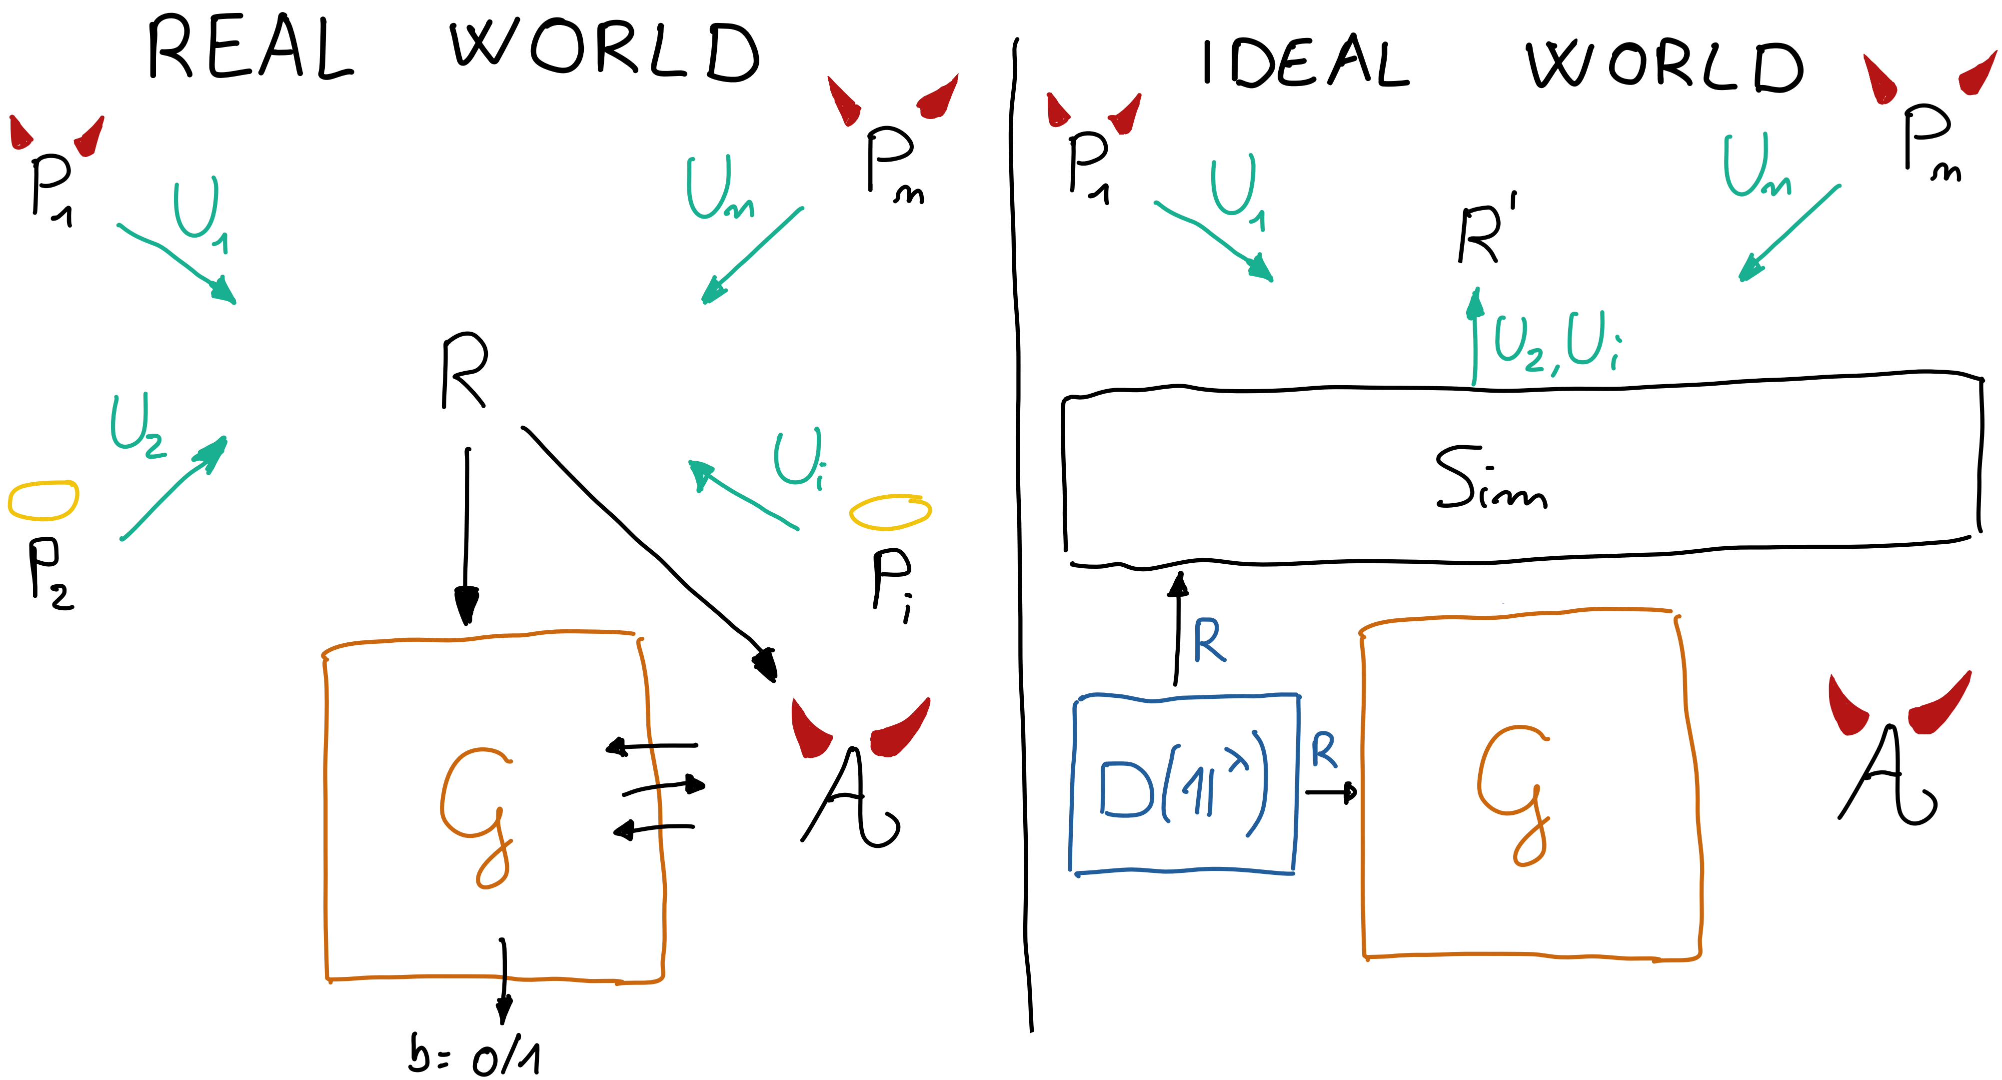

HARDNESS-PRESERVING DISTRIBUTED SAMPLERS

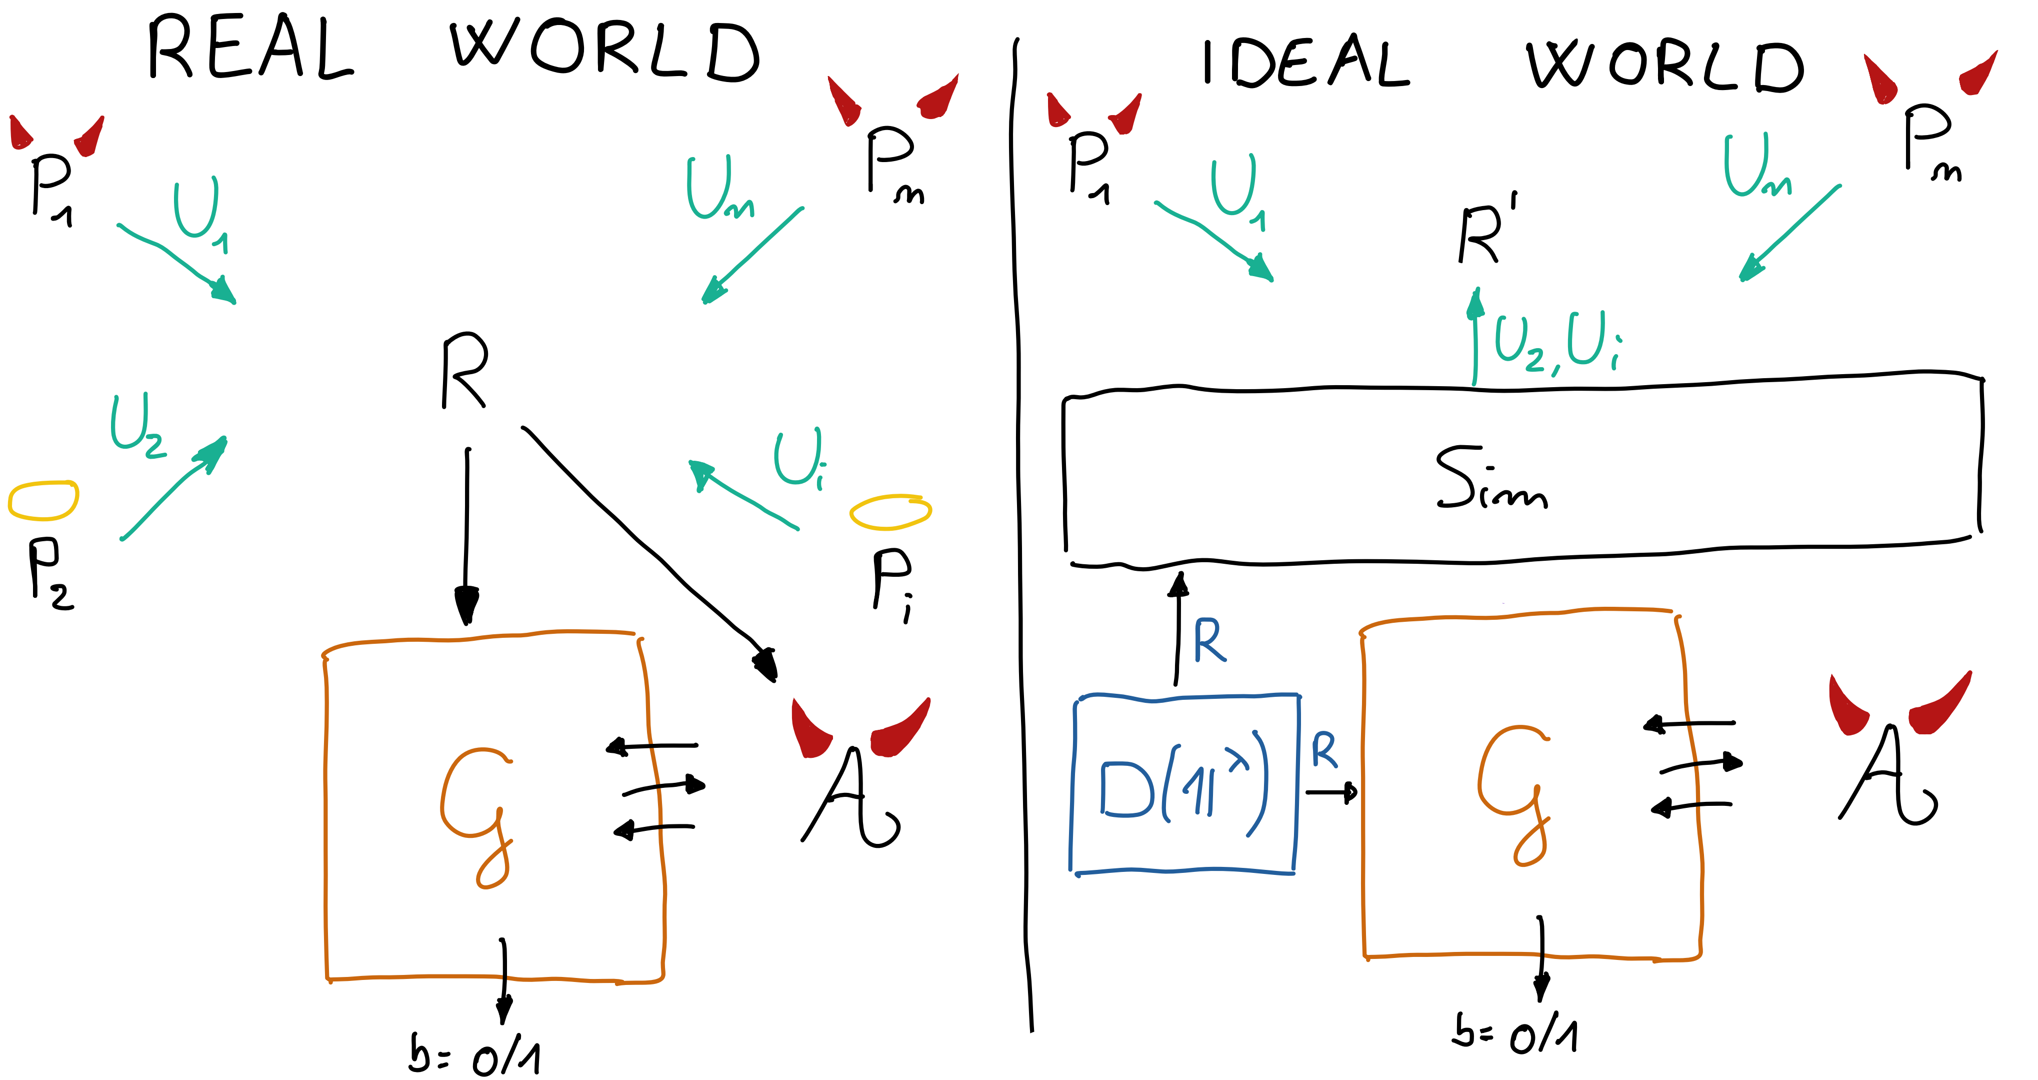

HARDNESS-PRESERVING DISTRIBUTED SAMPLERS

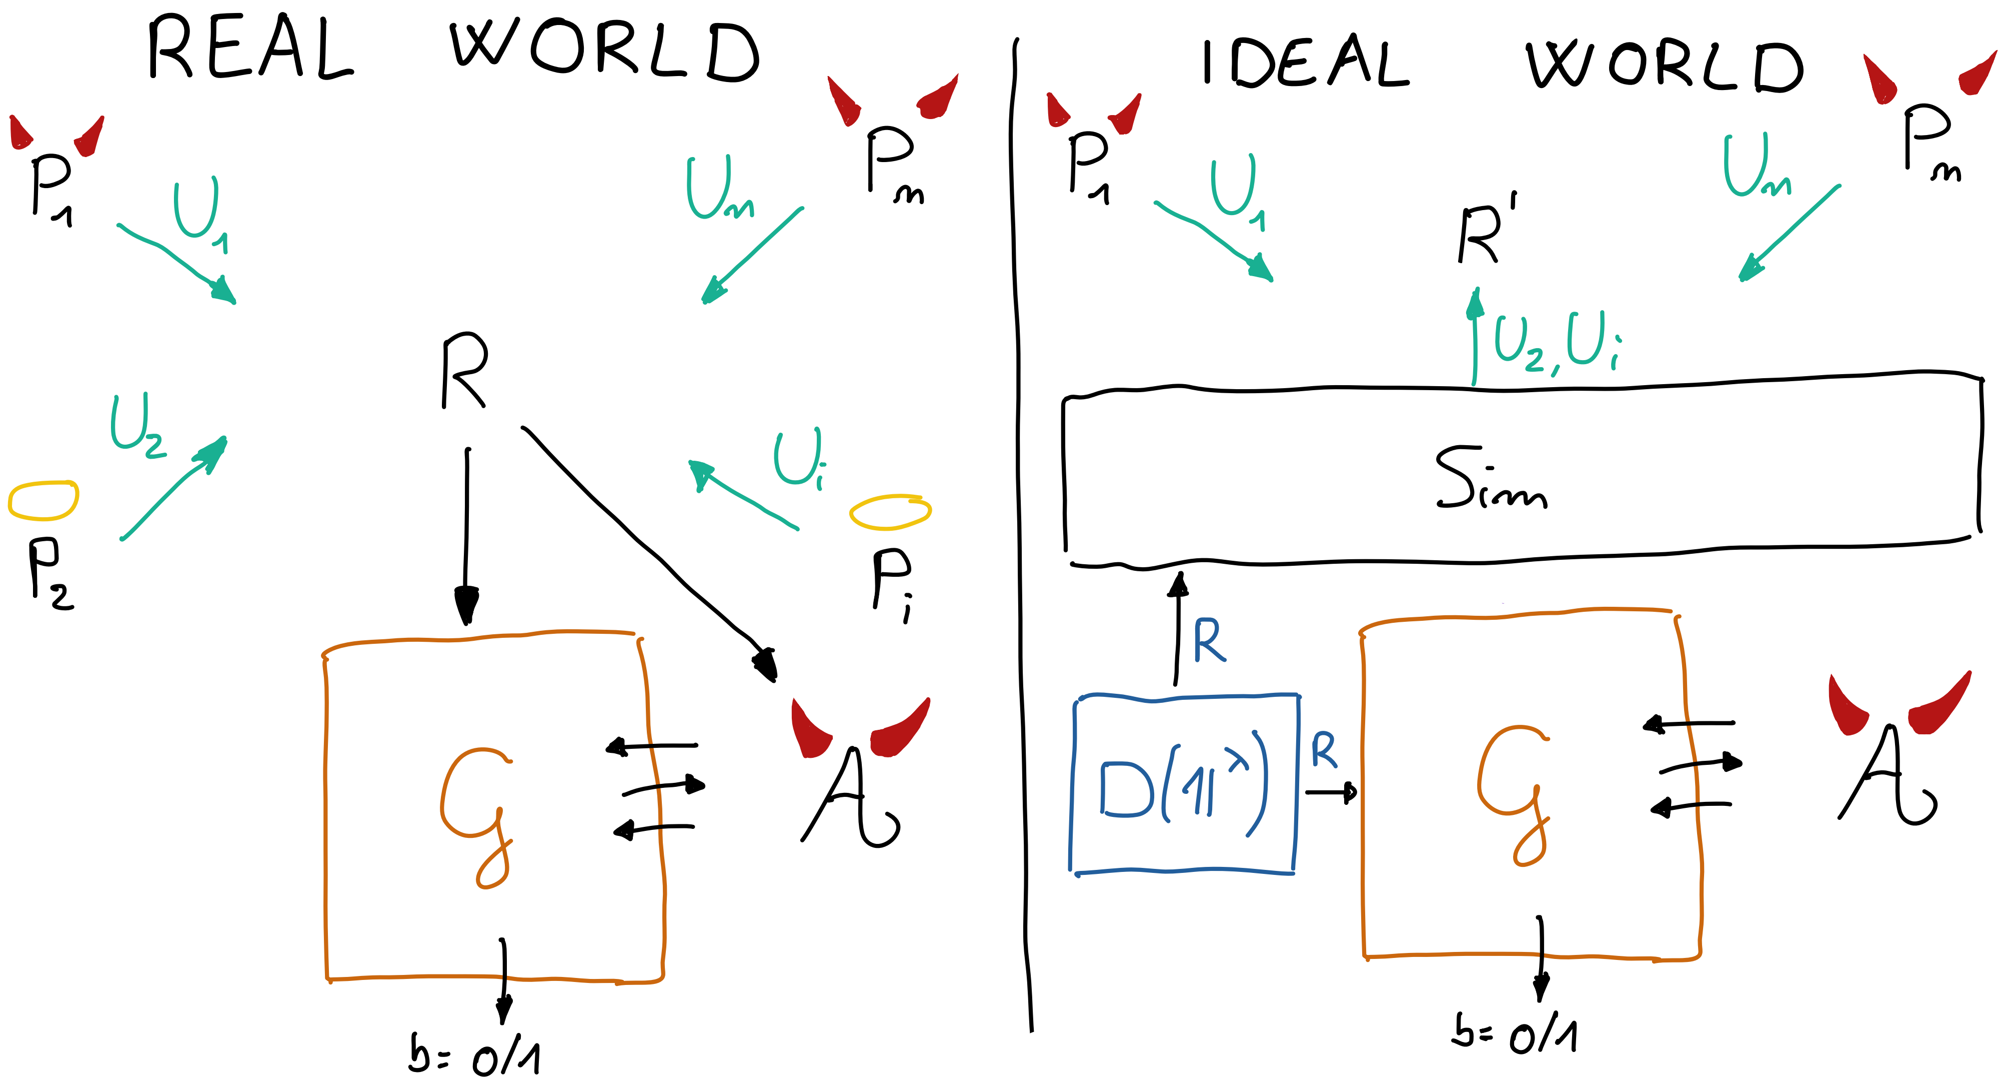

$$\mathbb{P}_{IDEAL}[b=1] \leq negl(\lambda)$$

HARDNESS-PRESERVING DISTRIBUTED SAMPLERS

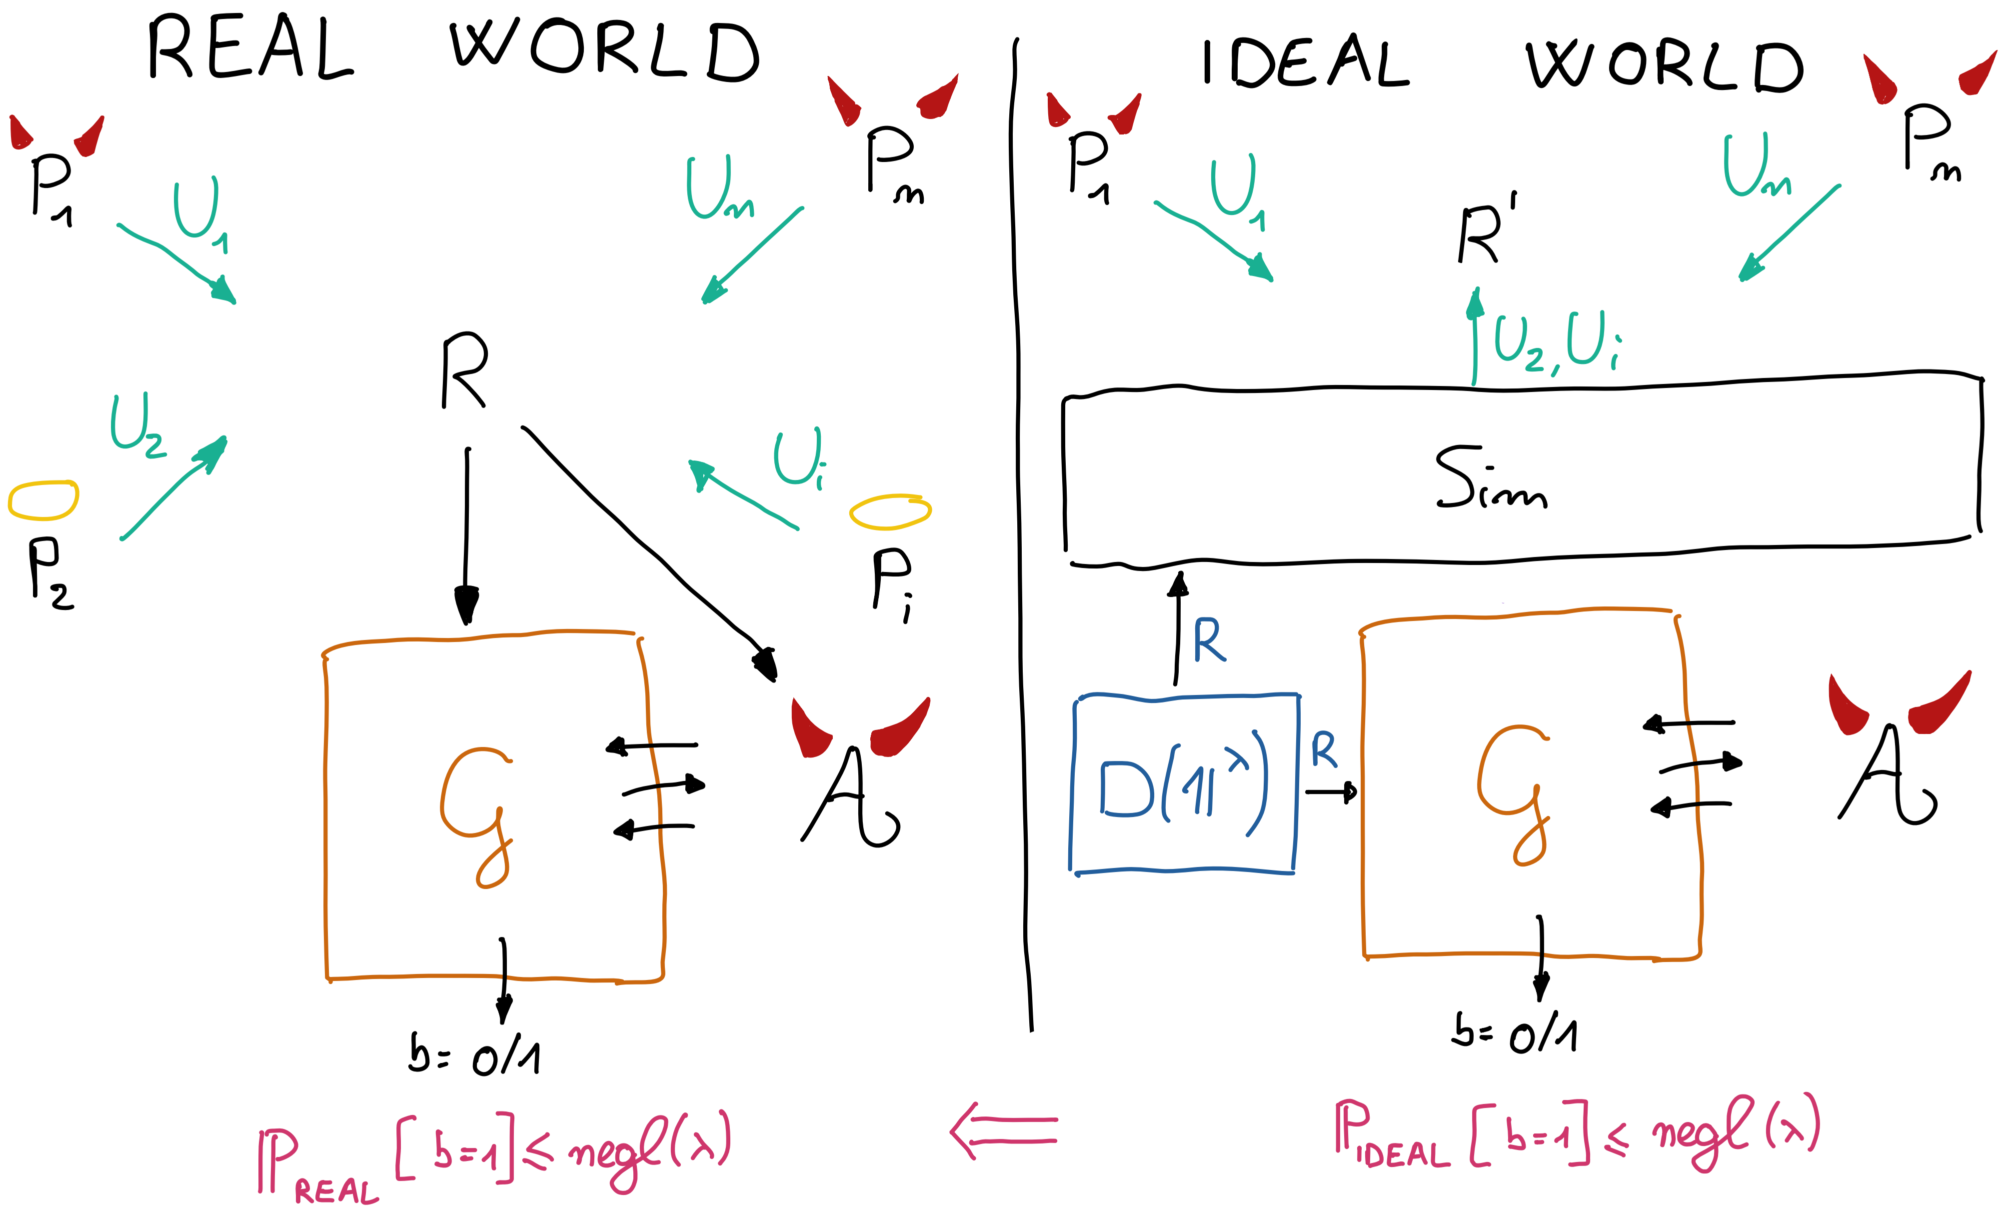

SECURITY GUARANTEES OF HARDNESS-PRESERVING DISTRIBUTED SAMPLERS

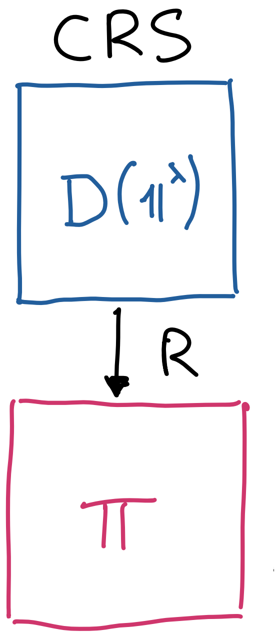

SECURITY GUARANTEES OF HARDNESS-PRESERVING DISTRIBUTED SAMPLERS

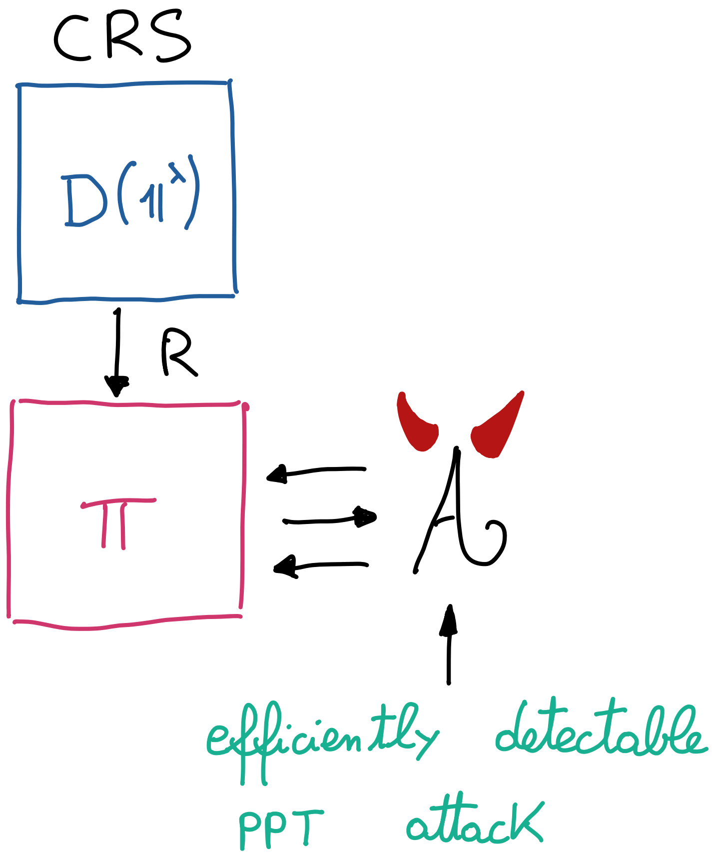

SECURITY GUARANTEES OF HARDNESS-PRESERVING DISTRIBUTED SAMPLERS

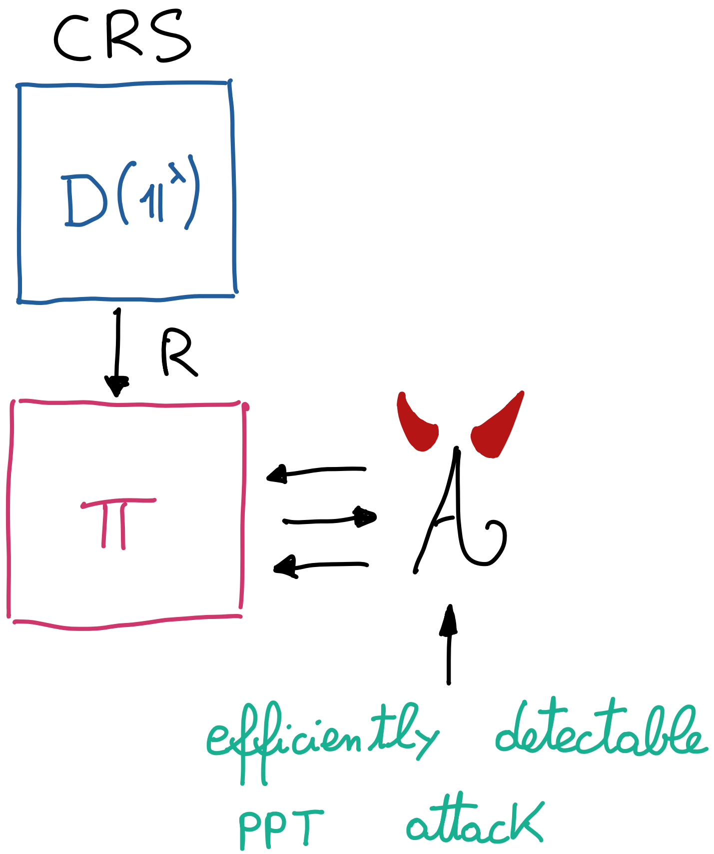

$$\mathbb{P}_{CRS}[\mathcal{A} \text{ succeeds}] < negl(\lambda)$$

SECURITY GUARANTEES OF HARDNESS-PRESERVING DISTRIBUTED SAMPLERS

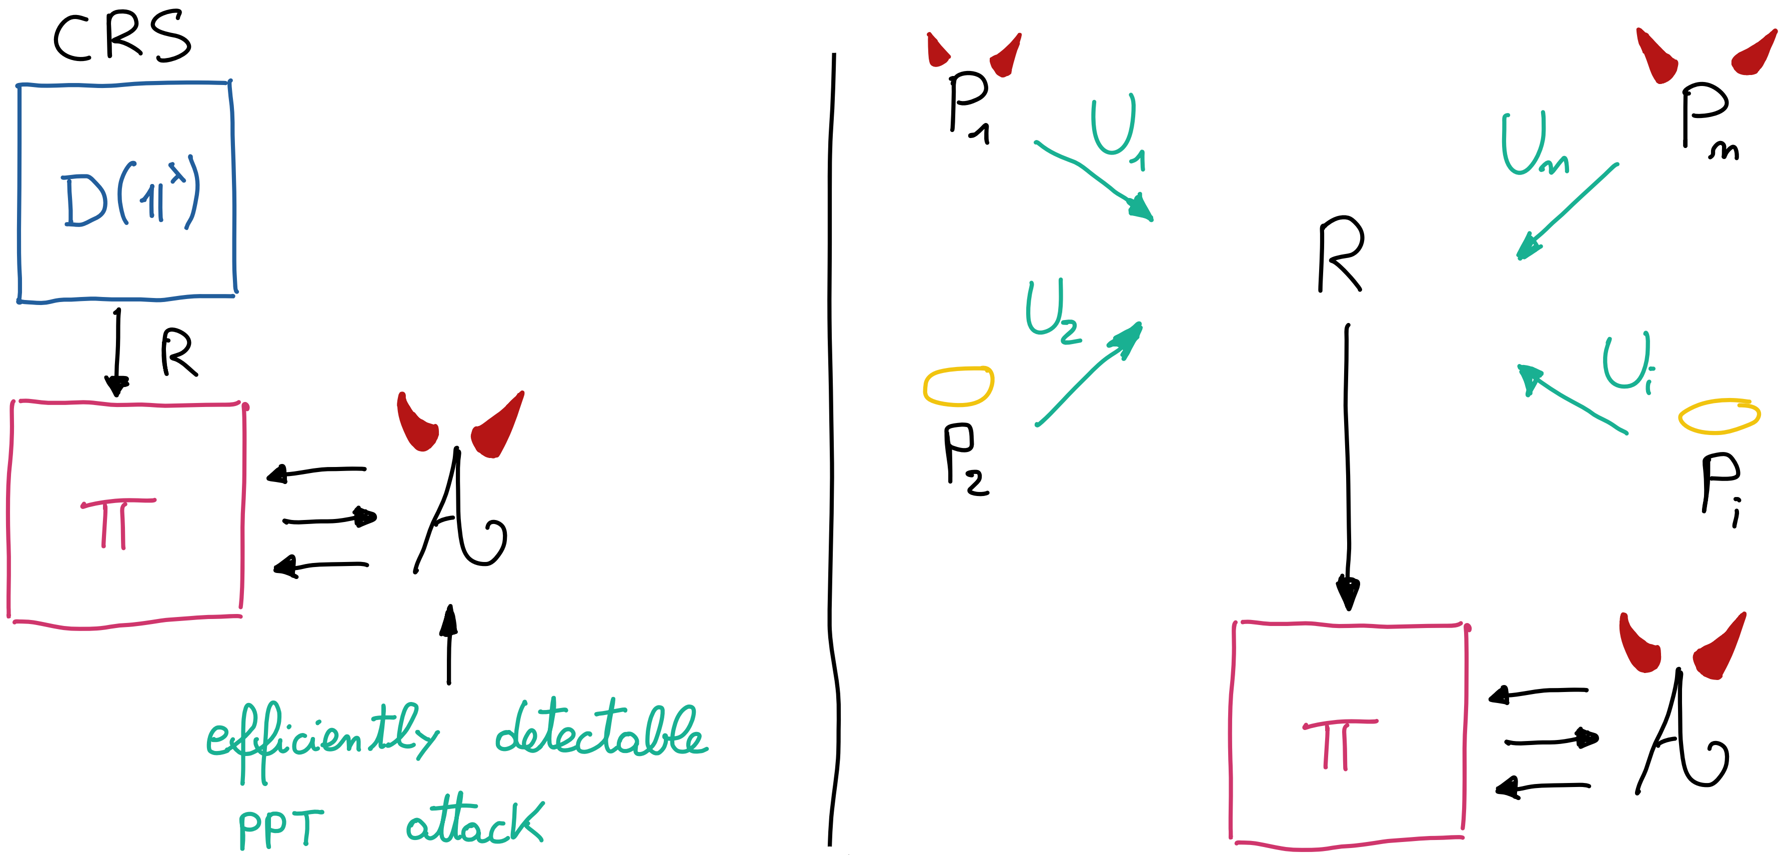

$$\mathbb{P}_{CRS}[\mathcal{A} \text{ succeeds}] < negl(\lambda)$$

SECURITY GUARANTEES OF HARDNESS-PRESERVING DISTRIBUTED SAMPLERS

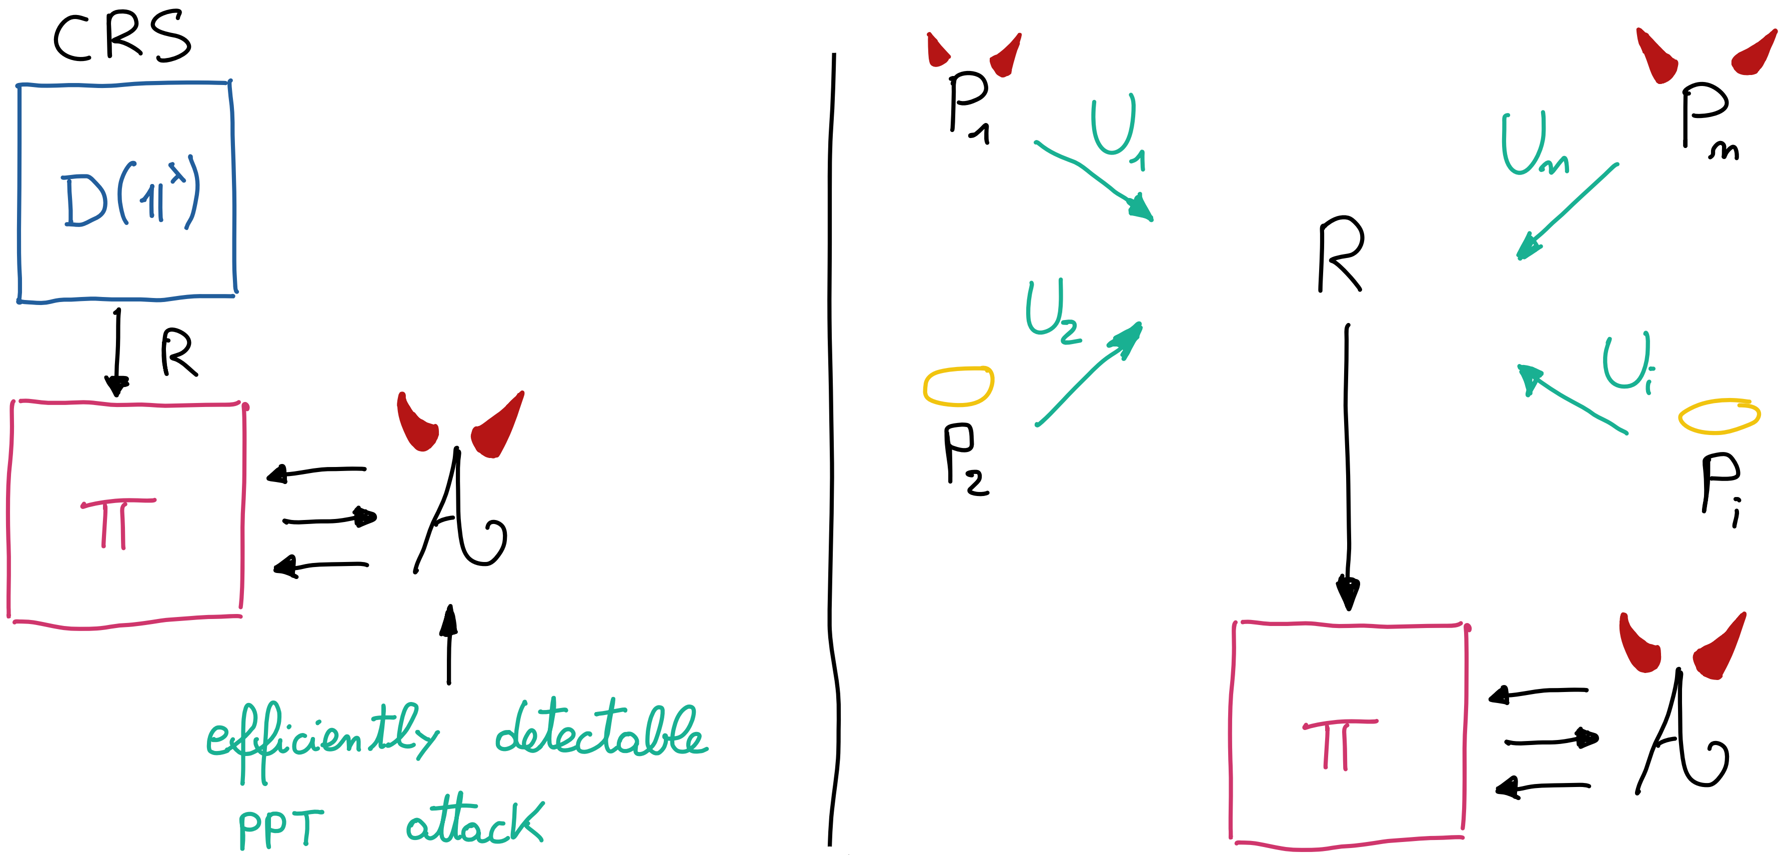

$$\mathbb{P}_{CRS}[\mathcal{A} \text{ succeeds}] < negl(\lambda) \Rightarrow \mathbb{P}_{DS}[\mathcal{A} \text{ succeeds}] < negl(\lambda)$$

LOSSY DISTRIBUTED SAMPLERS

distributed sampler message

$$\forall U: \qquad \Omega_U := \{ \mathrm{Sample}(U, U_2, \ldots, U_{n-1}) \mid U_2, \ldots, U_{n-1} \}$$

LOSSY DISTRIBUTED SAMPLERS

distributed sampler message

$\forall U$:

$$\Omega_U := \{ \mathrm{Sample}(U, U_2, \dots, U_{n-1}) \mid U_2, \dots, U_{n-1} \}$$

STANDARD MODE

LOSSY MODE

LOSSY DISTRIBUTED SAMPLERS

distributed sampler message

$$\forall U: \qquad \Omega_U := \{ \mathrm{Sample}(U, U_2, \dots, U_{n-1}) \mid U_2, \dots, U_{n-1} \}$$

polynomial

STANDARD MODE | LOSSY MODE (q)

LOSSY DISTRIBUTED SAMPLERS

distributed sampler message

$$\forall U: \qquad \Omega_U := \{ \mathrm{Sample}(U, U_2, \ldots, U_{n-1}) \mid U_2, \ldots, U_{n-1} \}$$

polynomial

STANDARD MODE | LOSSY MODE (q)

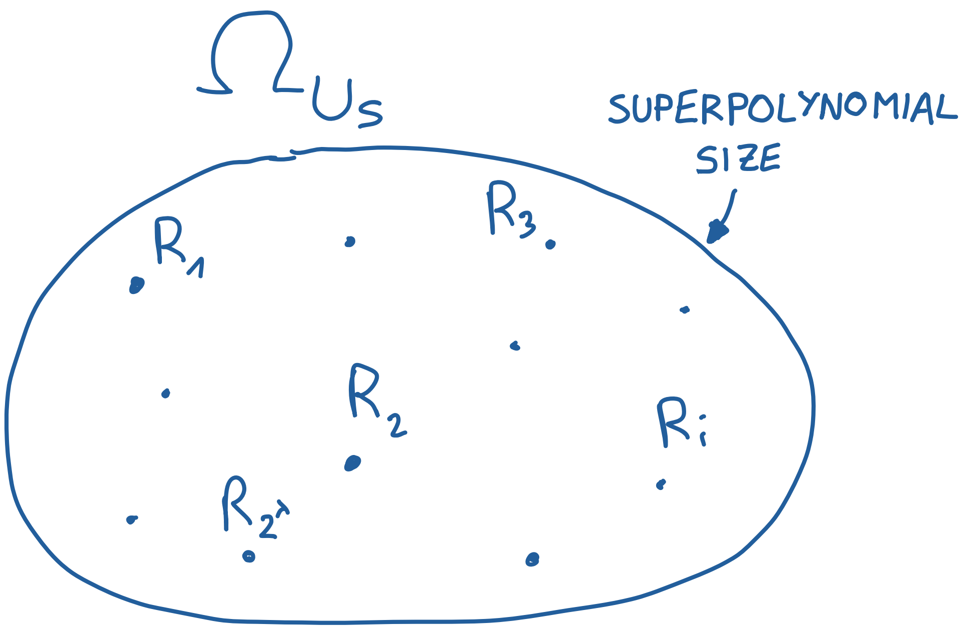

LOSSY DISTRIBUTED SAMPLERS

distributed sampler message

$$\forall U: \qquad \Omega_U := \left\{ \mathsf{Sample}(U, U_2, \dots, U_{n-1}) \;\middle|\; U_2, \dots, U_{n-1} \right\}$$

polynomial

STANDARD MODE | LOSSY MODE (q)

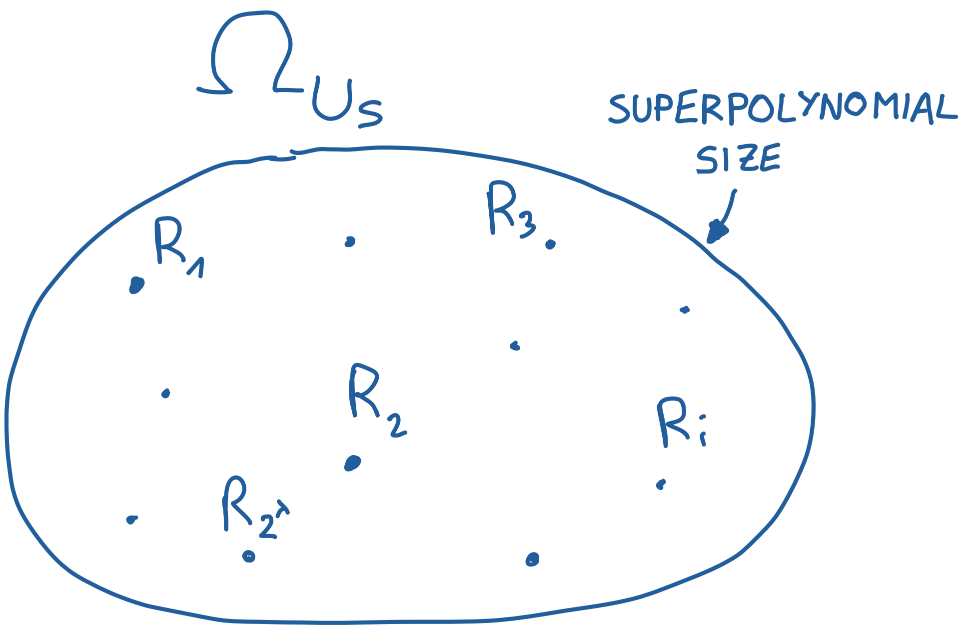

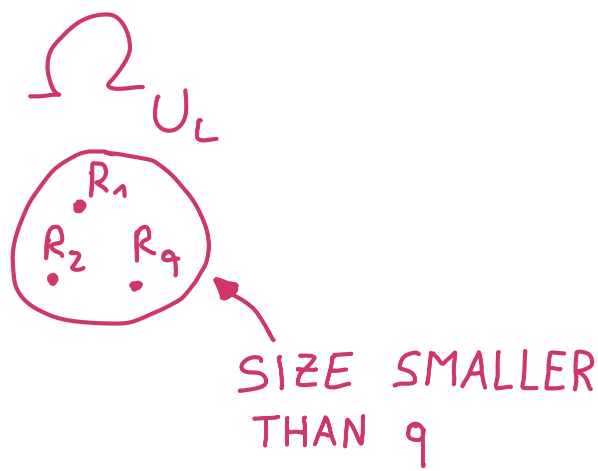

LOSSY DISTRIBUTED SAMPLERS

distributed sampler message

$\forall U$: $\Omega_U := \{ \text{Sample}(U, U_2, \dots, U_{n-1}) \mid U_2, \dots, U_{n-1} \}$

polynomial

STANDARD MODE | LOSSY MODE (q)

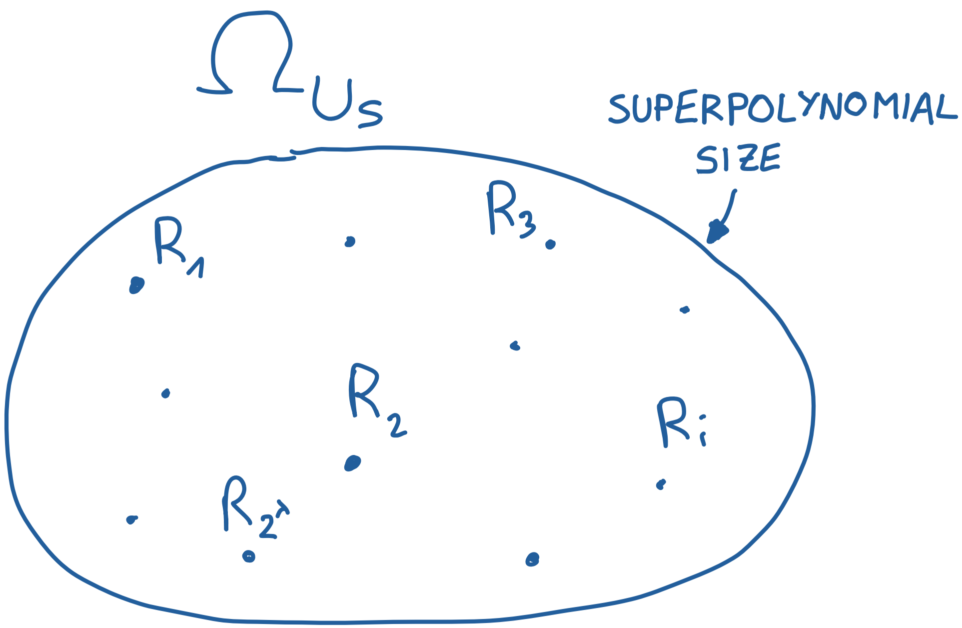

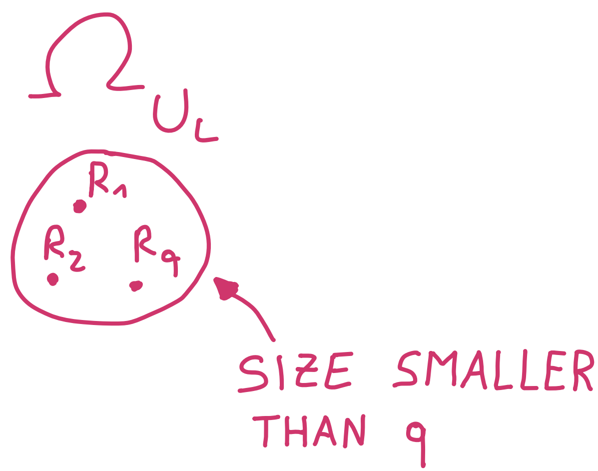

DISTINGUISHABLE WITH ARBITRARILY SMALL INVERSE-POLYNOMIAL ADVANTAGE!

PROGRAMMABILITY OF LOSSY DISTRIBUTED SAMPLERS

PROGRAMMABILITY OF LOSSY DISTRIBUTED SAMPLERS

LOSSY MODE (q)

U_L

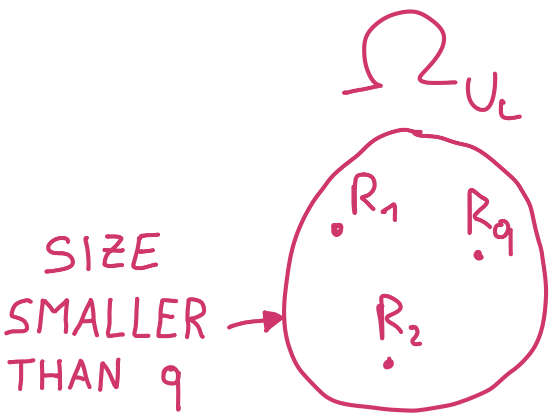

PROGRAMMED MODE (q)

$\hat{U}_L$

$R \xleftarrow{\$} D(1^\lambda)$

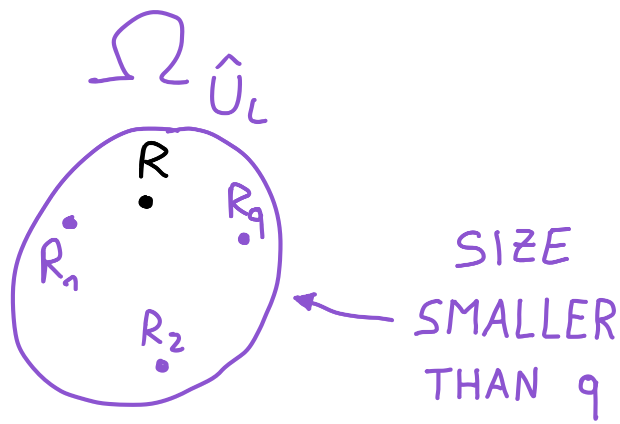

PROGRAMMABILITY OF LOSSY DISTRIBUTED SAMPLERS

LOSSY MODE (q) | **PROGRAMMED MODE (q)**

U_L

$\hat{U}_L$ $\quad R \xleftarrow{\$} D(1^\lambda)$

SIZE SMALLER THAN q

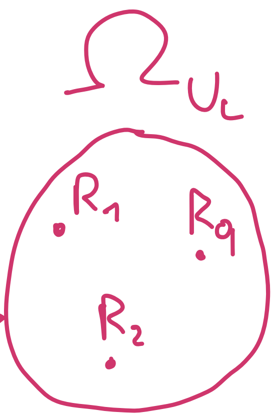

SIZE SMALLER THAN q

INDISTINGUISHABLE

BUILDING LOSSY DISTRIBUTED SAMPLERS

THEOREM

Assume the existence of

- subexp iO
- subexp multi-Key FHE
- extremely lossy functions (ELFs)
- subexp collision resistant hash functions

BUILDING LOSSY DISTRIBUTED SAMPLERS

THEOREM

Assume the existence of

- subexp iO
- subexp multi-Key FHE
- extremely lossy functions (ELFs)
- subexp collision resistant hash functions
- almost-everywhere-extractable NIZKs

subexp injective OWFs
perfectly correct IBE
perfectly sound NIWI

BUILDING LOSSY DISTRIBUTED SAMPLERS

THEOREM

Assume the existence of

- subexp iO
- subexp multi-Key FHE
- extremely lossy functions (ELFs)
- subexp collision resistant hash functions
- almost-everywhere-extractable NIZKs

subexp injective OWFs
perfectly correct IBE
perfectly sound NIWI

Then, there exists a programmable lossy distributed sampler with a short (poly λ), reusable CRS.

can be made unstructured

FROM LOSSY TO HARDNESS-PRESERVING DISTRIBUTED SAMPLERS

REAL WORLD

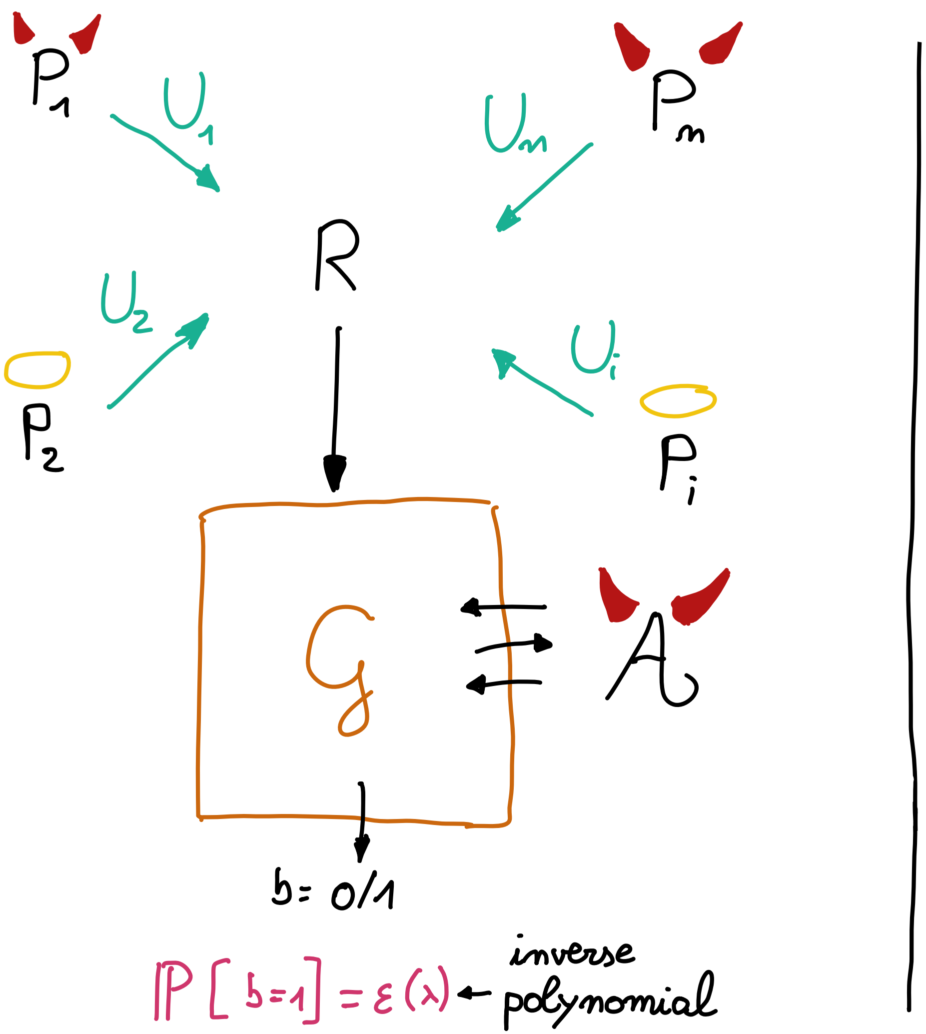

FROM LOSSY TO HARDNESS-PRESERVING DISTRIBUTED SAMPLERS

REAL WORLD

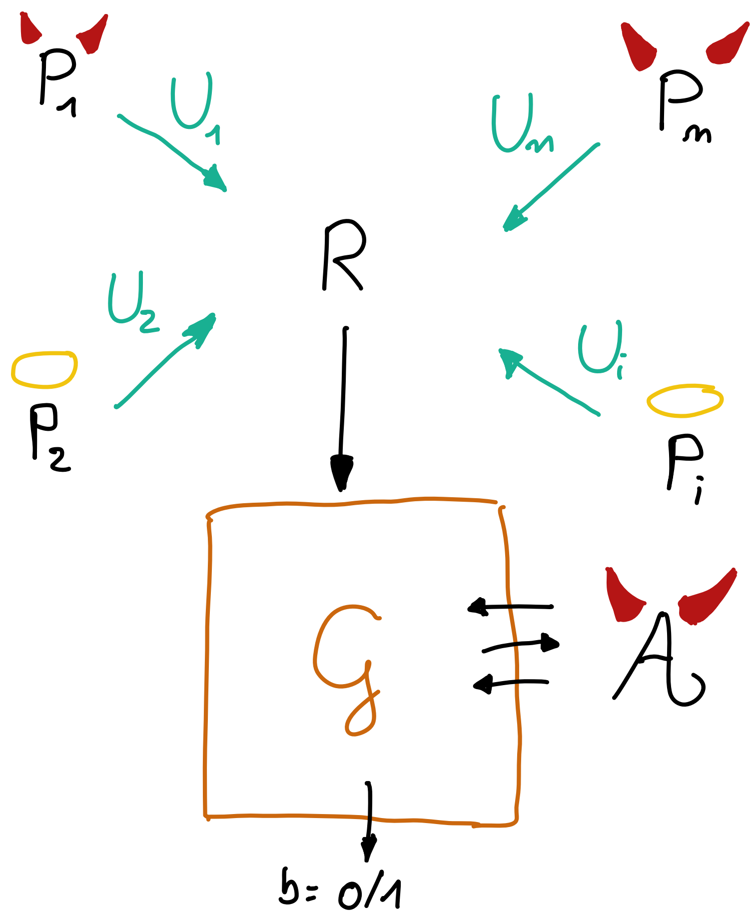

$\mathbb{P}[b=1] = \varepsilon(\lambda)$ ← inverse polynomial

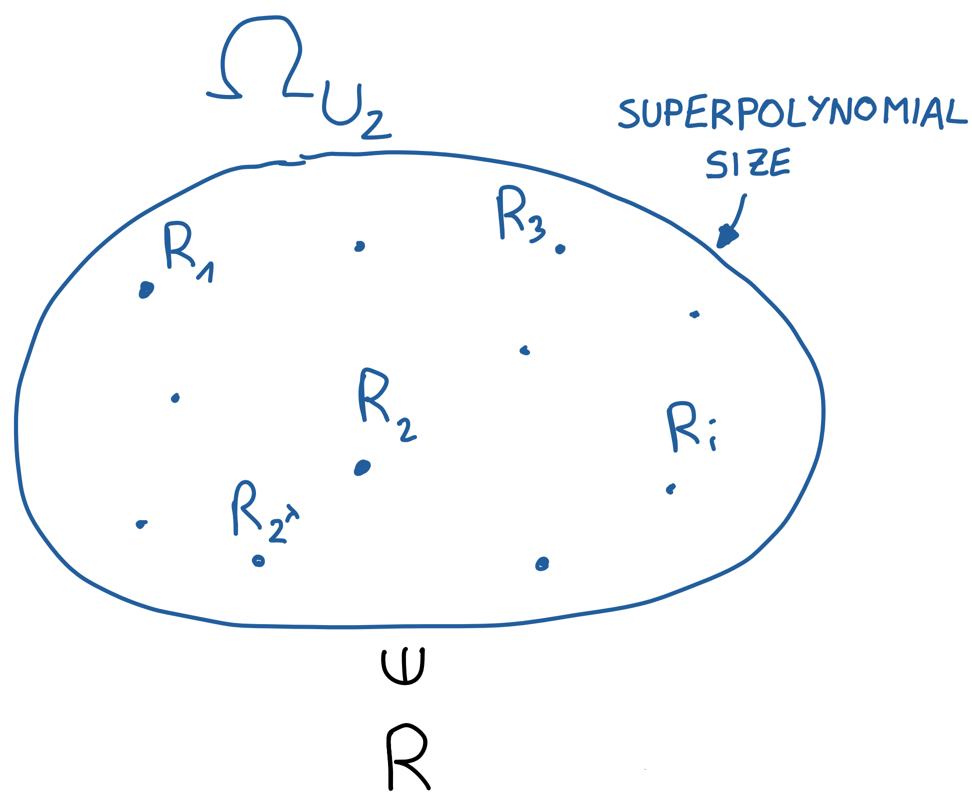

FROM LOSSY TO HARDNESS-PRESERVING DISTRIBUTED SAMPLERS

HYBRID WORLD 1

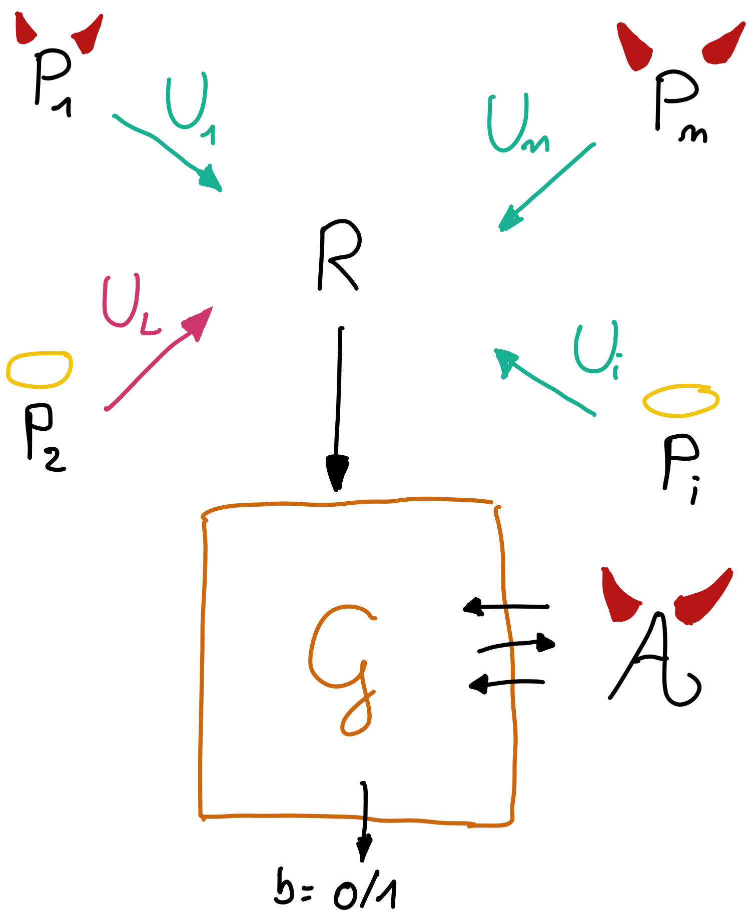

FROM LOSSY TO HARDNESS-PRESERVING DISTRIBUTED SAMPLERS

HYBRID WORLD 1

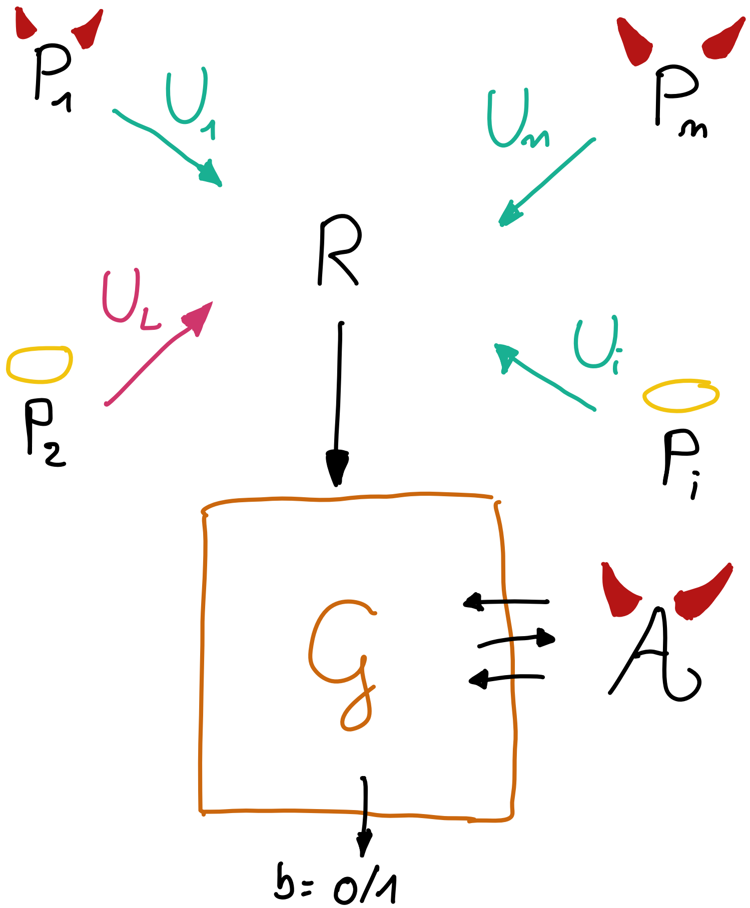

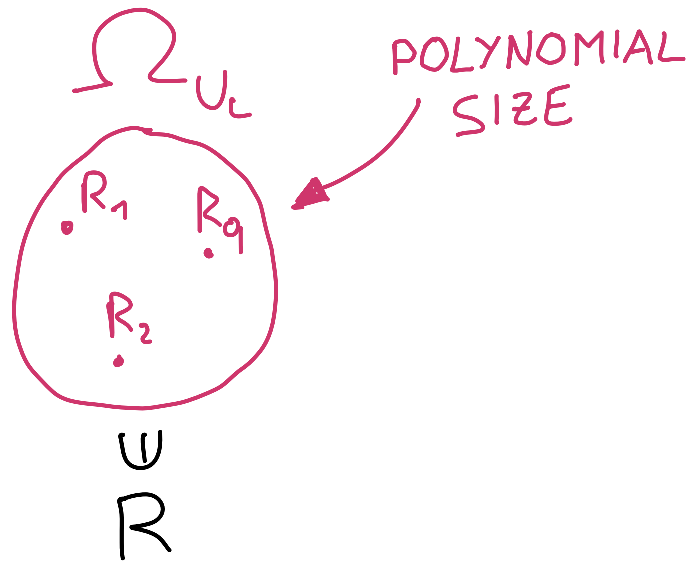

FROM LOSSY TO HARDNESS-PRESERVING DISTRIBUTED SAMPLERS

HYBRID WORLD 1

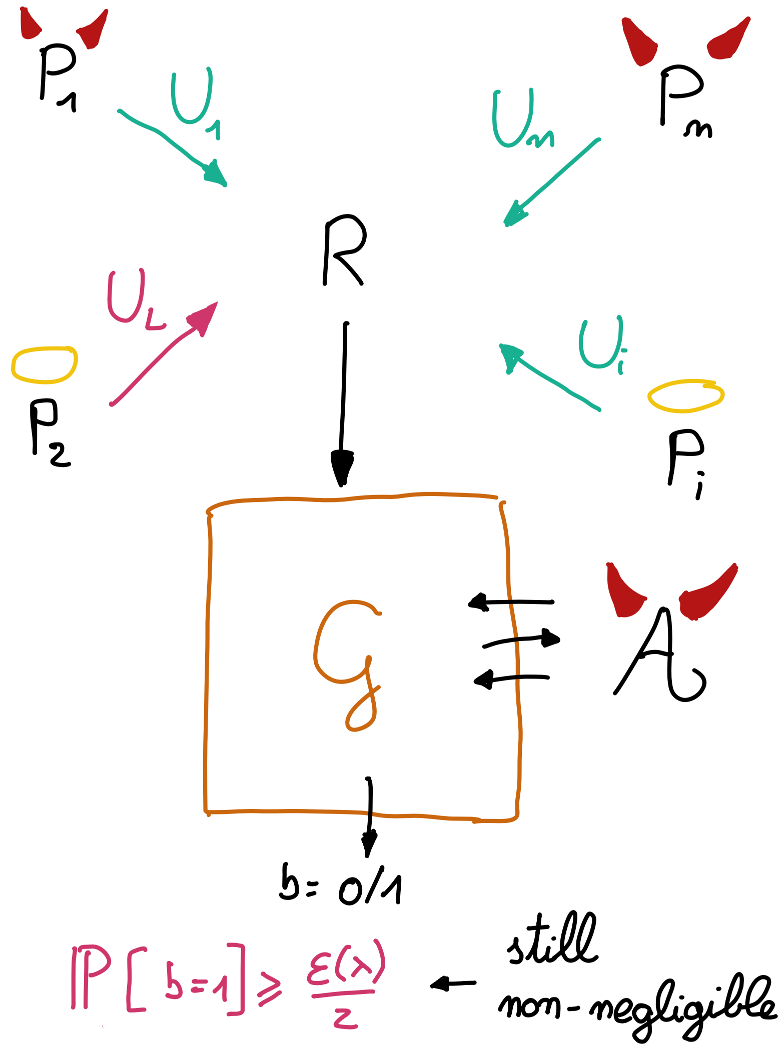

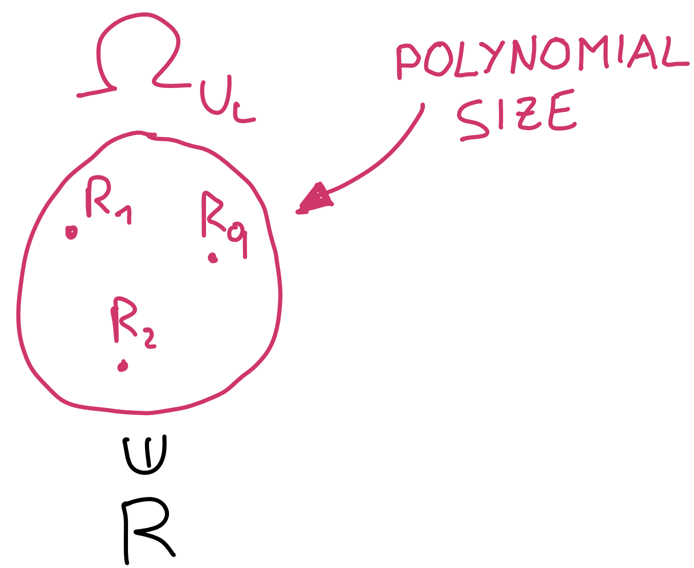

FROM LOSSY TO HARDNESS-PRESERVING DISTRIBUTED SAMPLERS

HYBRID WORLD 2

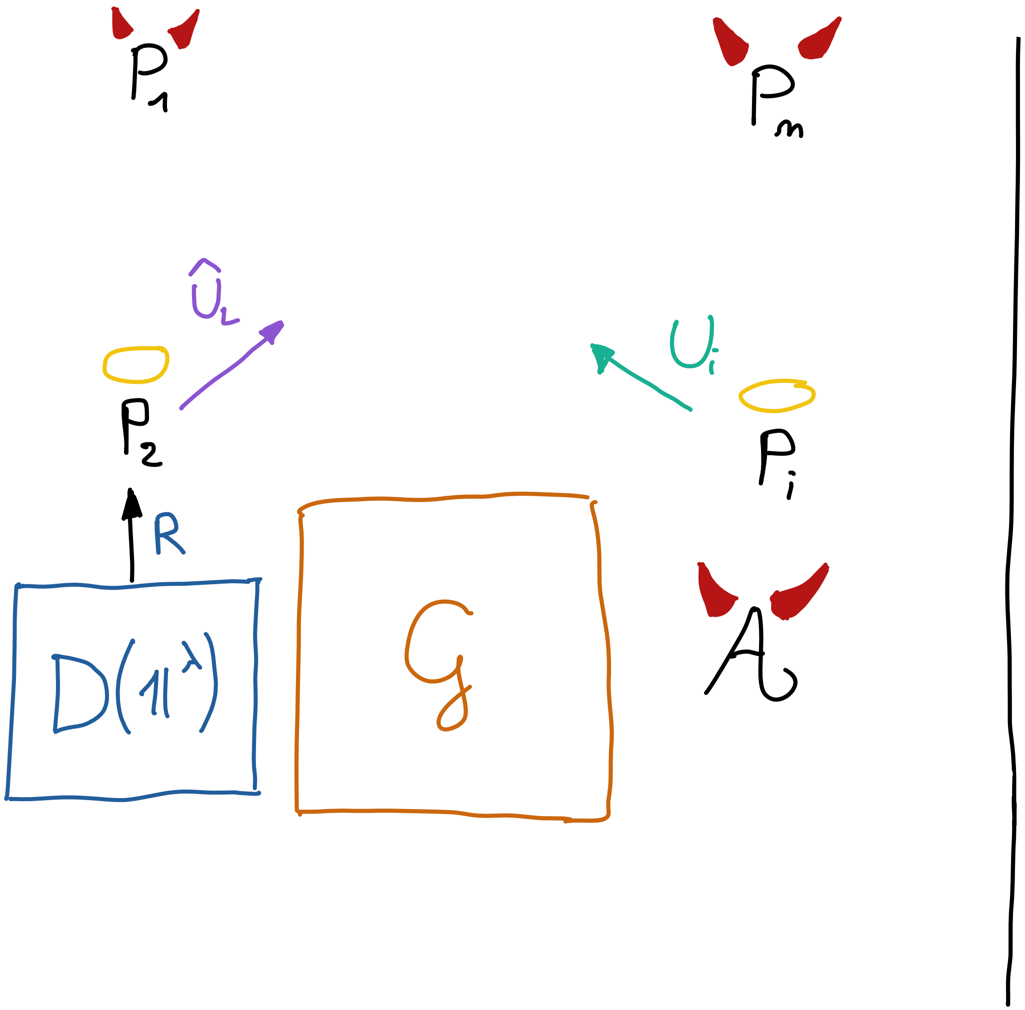

FROM LOSSY TO HARDNESS-PRESERVING DISTRIBUTED SAMPLERS

HYBRID WORLD 2

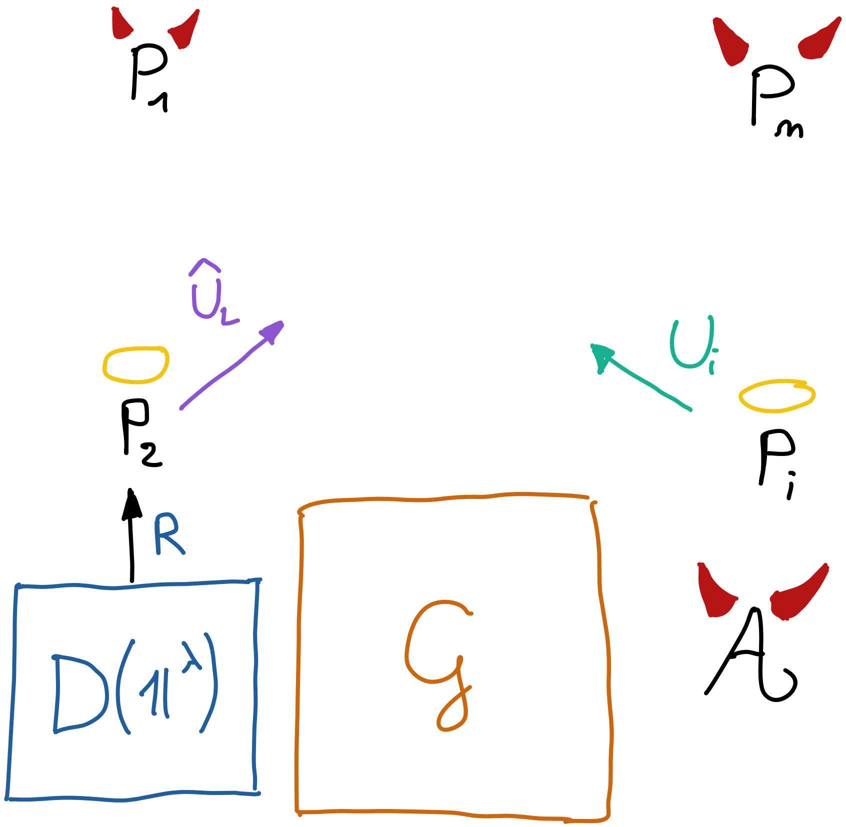

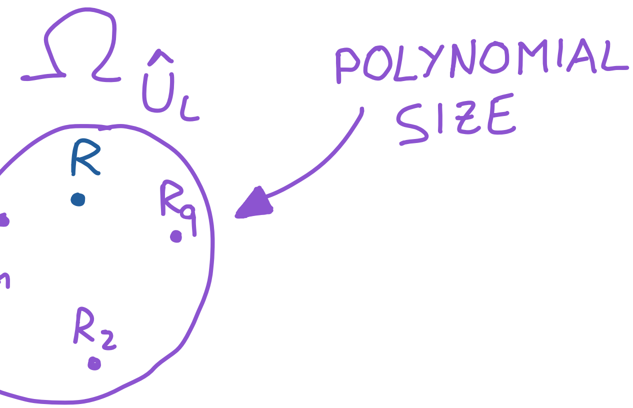

FROM LOSSY TO HARDNESS-PRESERVING DISTRIBUTED SAMPLERS

HYBRID WORLD 2

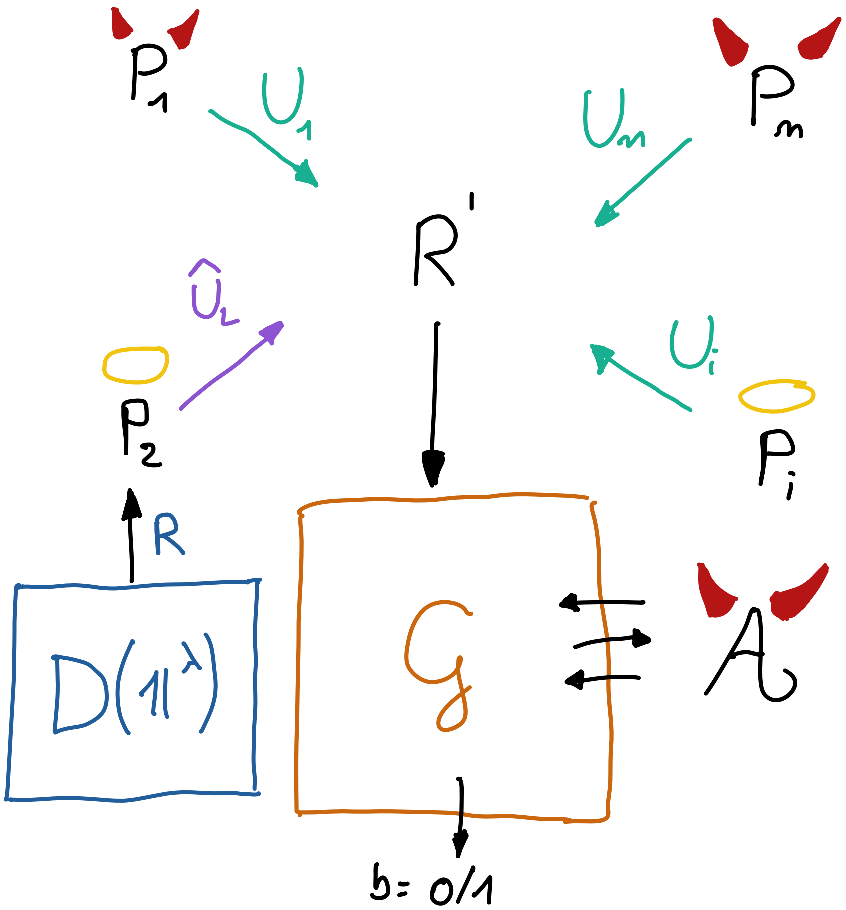

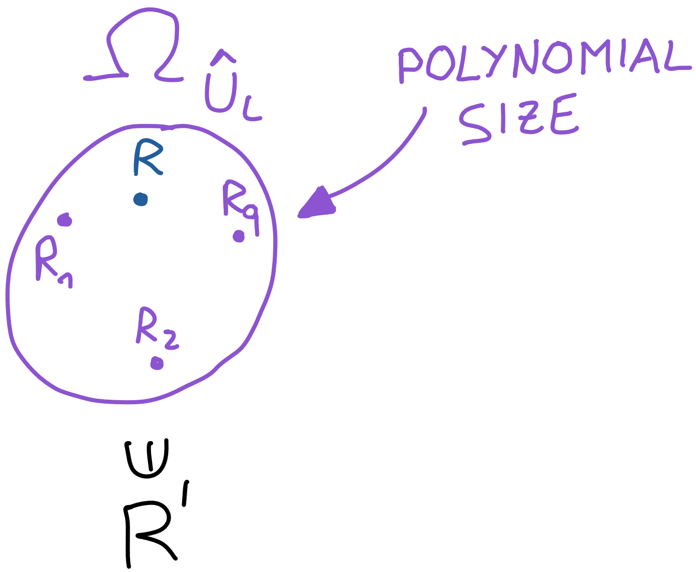

FROM LOSSY TO HARDNESS-PRESERVING DISTRIBUTED SAMPLERS

HYBRID WORLD 2

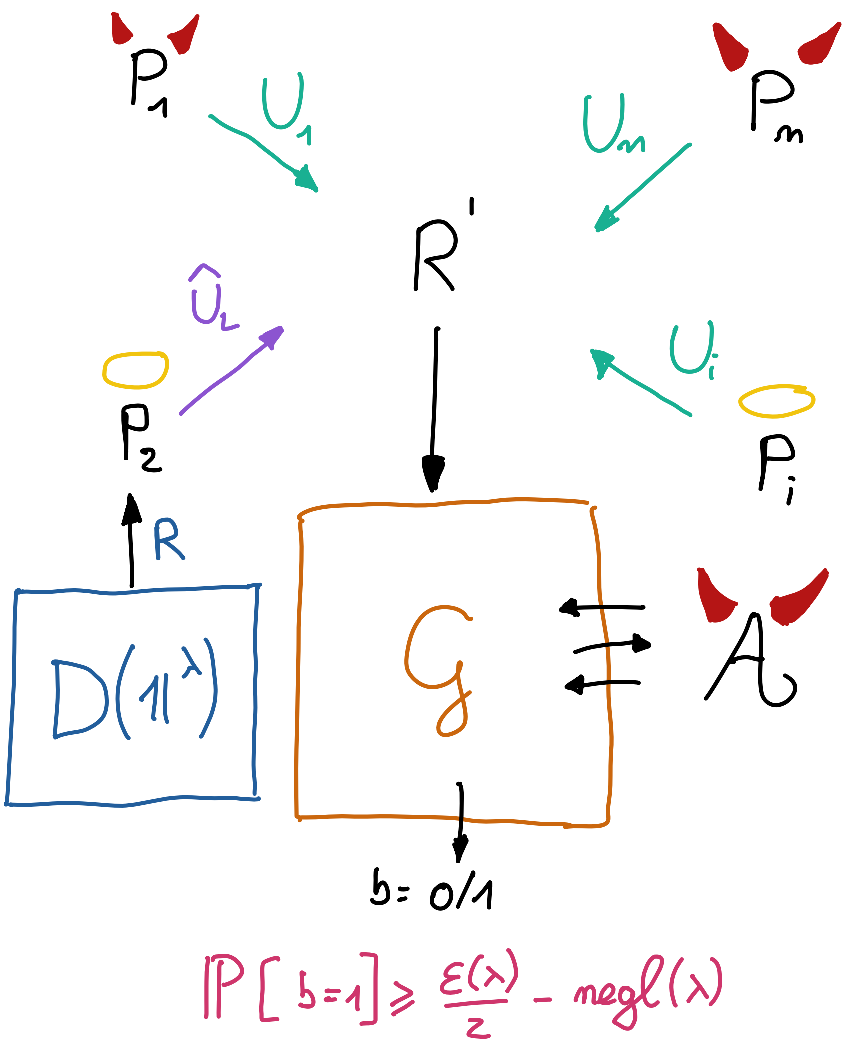

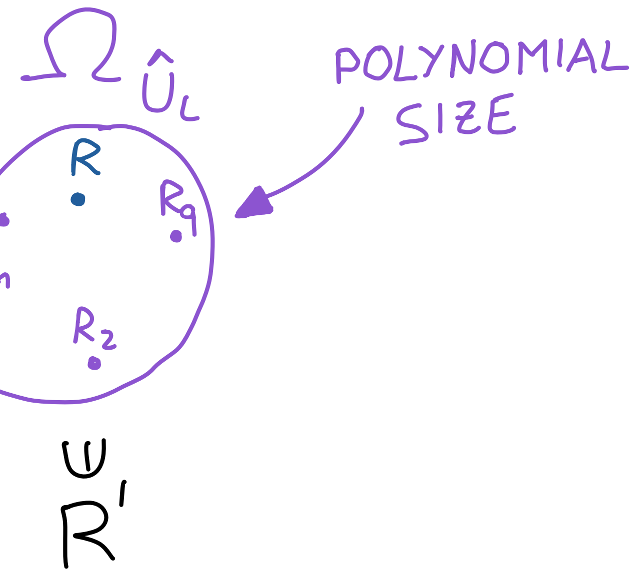

FROM LOSSY TO HARDNESS-PRESERVING DISTRIBUTED SAMPLERS

HYBRID WORLD 2

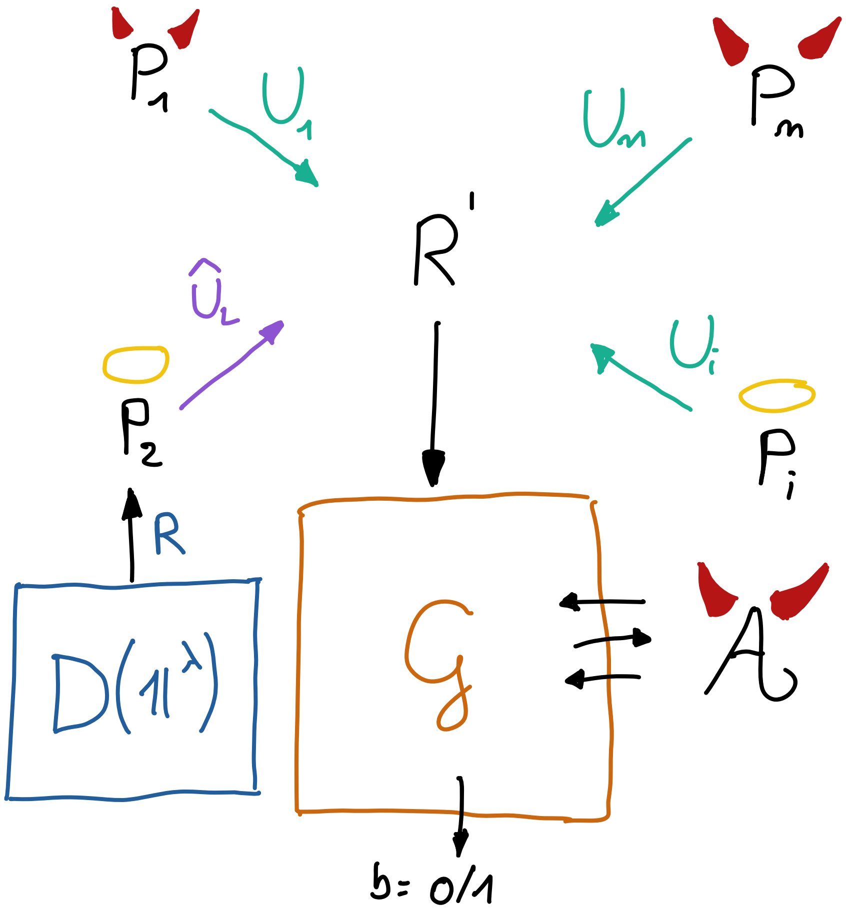

$$\mathbb{P}[b=1] \geq \frac{\varepsilon(\lambda)}{2} - negl(\lambda)$$

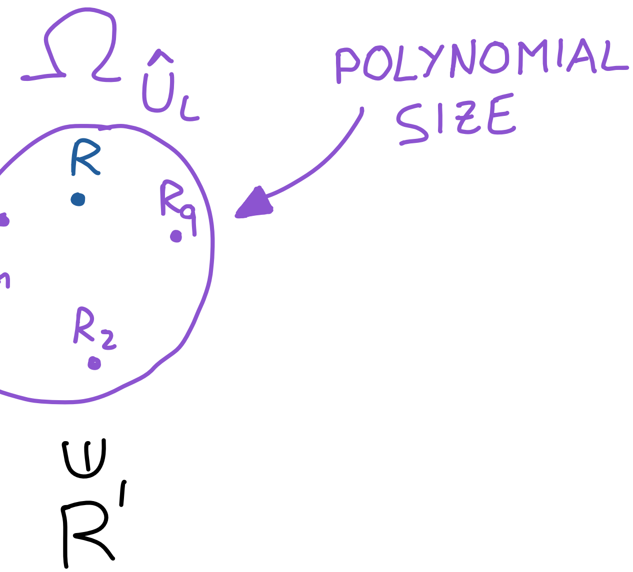

$$\mathbb{P}[b=1,\ R=R'] \geq \left(\frac{\varepsilon(\lambda)}{2} - negl(\lambda)\right) \cdot \frac{1}{poly(\lambda)}$$

FROM LOSSY TO HARDNESS-PRESERVING DISTRIBUTED SAMPLERS

HYBRID WORLD 2

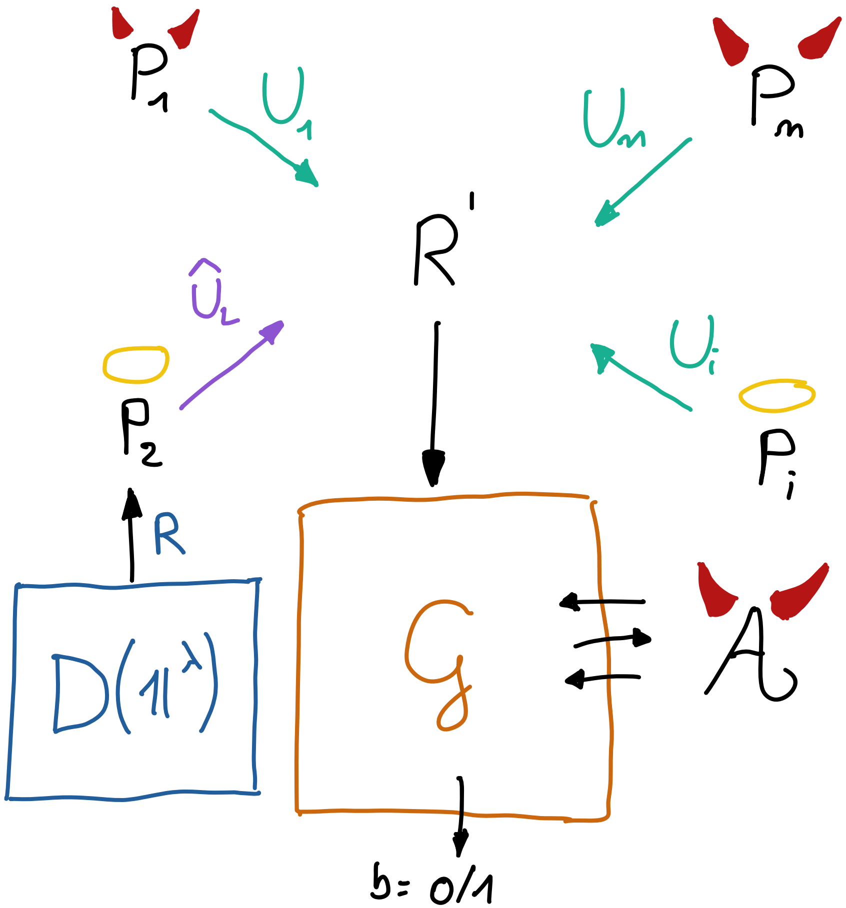

$$\mathbb{P}[b=1] \geq \frac{\varepsilon(\lambda)}{2} - negl(\lambda)$$

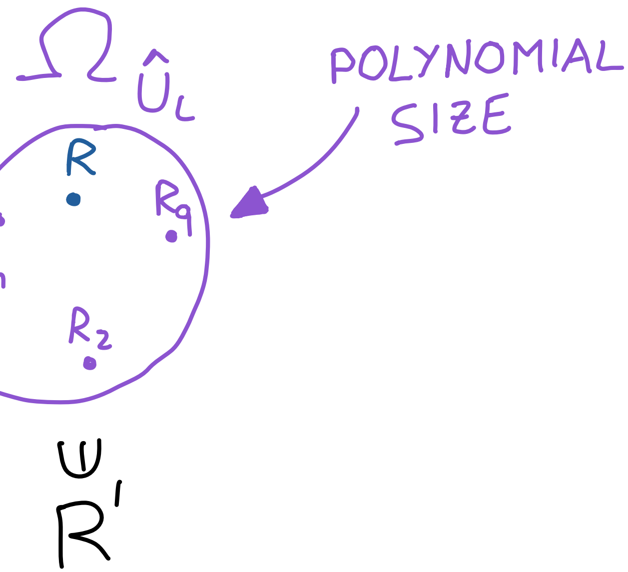

$$\mathbb{P}[b=1,\ R=R'] \geq \overbrace{\left(\frac{\varepsilon(\lambda)}{2} - negl(\lambda)\right) \cdot \frac{1}{poly(\lambda)}}^{\text{non-negligible}}$$

INDISTINGUISHABILITY-PRESERVING DISTRIBUTED SAMPLERS

REAL WORLD

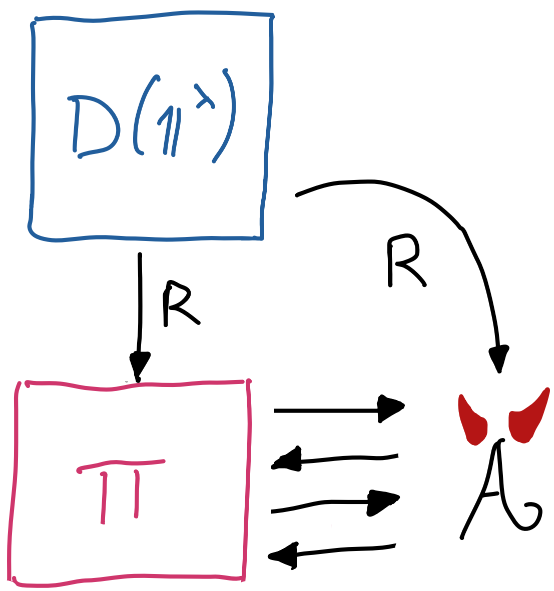

INDISTINGUISHABILITY-PRESERVING DISTRIBUTED SAMPLERS

REAL WORLD

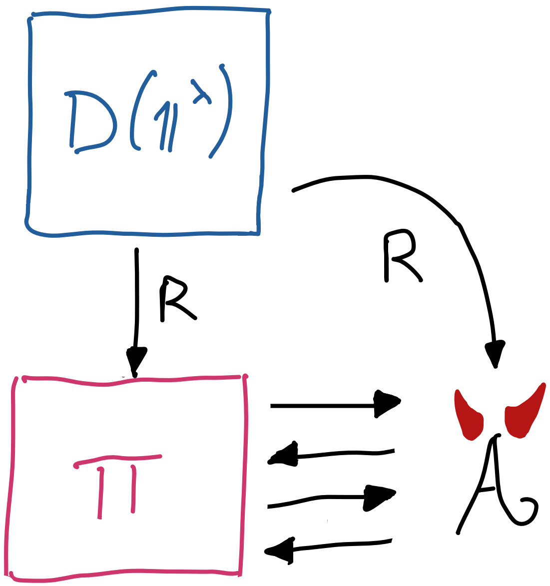

IDEAL WORLD

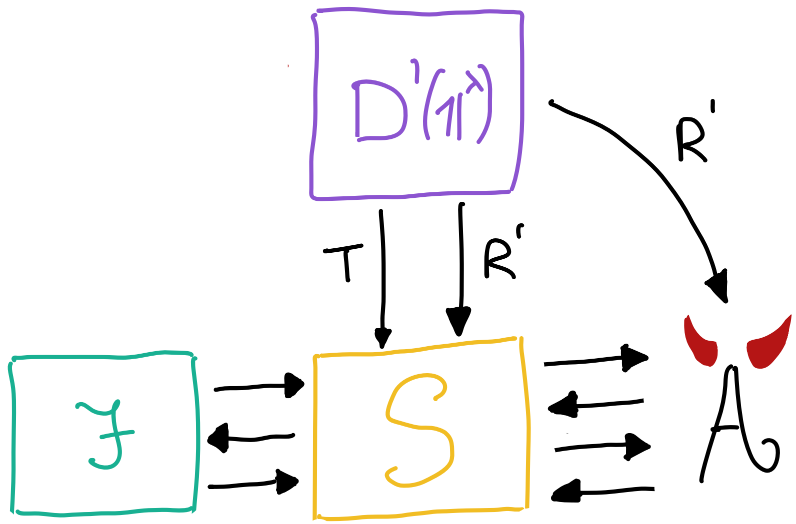

INDISTINGUISHABILITY-PRESERVING DISTRIBUTED SAMPLERS

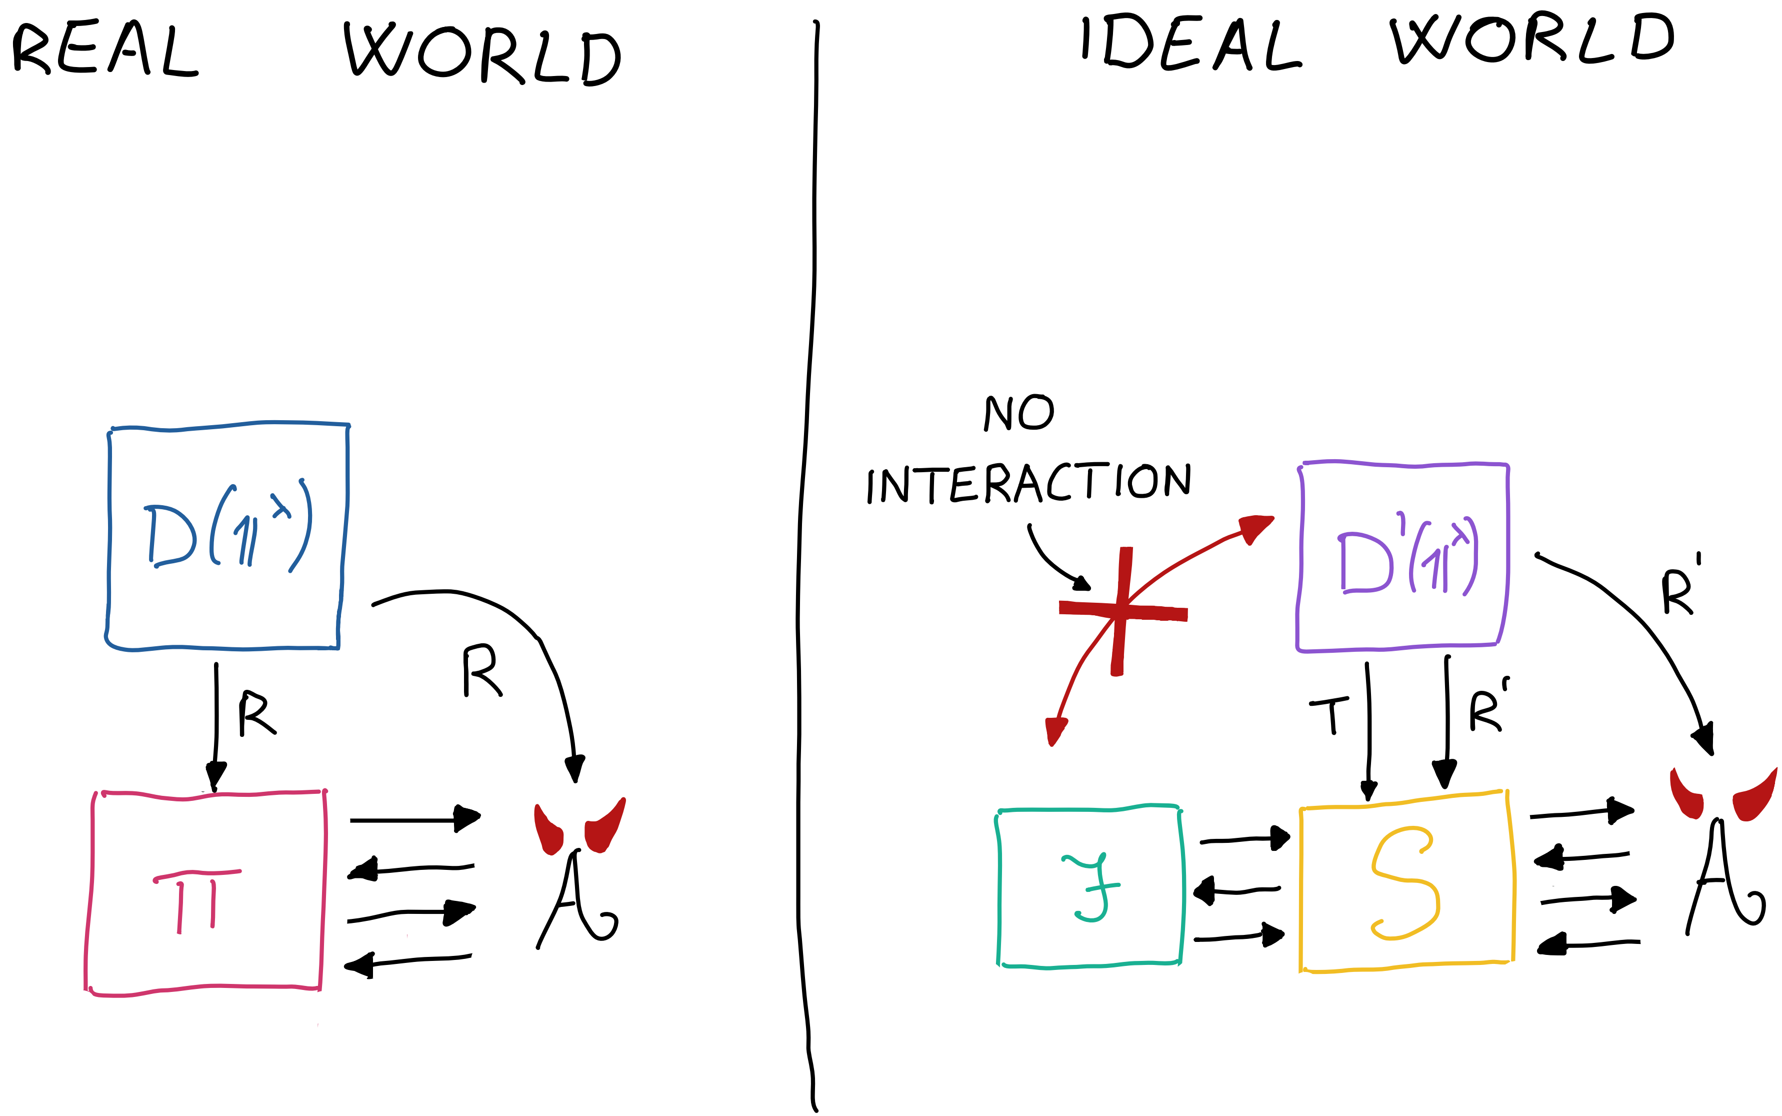

BUILDING INDISTINGUISHABILITY-PRESERVING DISTRIBUTED SAMPLERS

THEOREM

Our lossy distributed sampler is indistinguishability-preserving.

ON THE NEED FOR CRS'S

Our distributed samplers have CRS's that are:

ON THE NEED FOR CRS'S

Our distributed samplers have CRS's that are:

- reusable

ON THE NEED FOR CRS'S

Our distributed samplers have CRS's that are:

- reusable
- short

ON THE NEED FOR CRS'S

Our distributed samplers have CRS's that are:

- reusable
- short
- unstructured

ON THE NEED FOR CRS'S

Our distributed samplers have CRS's that are:

- reusable
- short
- unstructured

Can we get rid of CRSs?

ON THE NEED FOR CRS'S

Our distributed samplers have CRS's that are:

- reusable
- short
- unstructured

Can we get rid of CRSs?

- for indistinguishability-preserving distributed samplers NO!

ON THE NEED FOR CRS'S

Our distributed samplers have CRS's that are:

- reusable
- short
- unstructured

Can we get rid of CRSs?

- for indistinguishability-preserving distributed samplers NO!

  BY COMPILING [PVW08], WE WOULD GET 3-ROUND ACTIVE OT IN THE PLAIN MODEL ↯

ON THE NEED FOR CRS'S

Our distributed samplers have CRS's that are:

- reusable
- short
- unstructured

Can we get rid of CRSs?

- for indistinguishability-preserving distributed samplers NO!

  BY COMPILING [PVW08], WE WOULD GET 3-ROUND ACTIVE OT IN THE PLAIN MODEL ↯

- for hardness-preserving distributed samplers OPEN!

ON THE NEED FOR CRS'S

We can build security-preserving distributed samplers without CRS if:

ON THE NEED FOR CRS'S

We can build security-preserving distributed samplers without CRS if:

- we restrict to uniform adversaries

ON THE NEED FOR CRS'S

We can build security-preserving distributed samplers without CRS if:

- we restrict to uniform adversaries
- we allow non-uniform simulators

ON THE NEED FOR CRS'S

We can build security-preserving distributed samplers without CRS if:

- we restrict to uniform adversaries
- we allow non-uniform simulators

We built CRS-less simulation-extractable NIZKs!

SUMMARY

NEW DEFINITIONS OF ACTIVE DISTRIBUTED SAMPLERS THAT DON'T NEED RANDOM ORACLES:

HARDNESS-PRESERVING DISTRIBUTED SAMPLERS

preserving the hardness of search games

INDISTINGUISHABILITY-PRESERVING DISTRIBUTED SAMPLERS

preserving the functionality for a large class of protocols.

BUILT FROM SUBEXP iO, SUBEXP MK-FHE, ELFs, ...

WITH REUSABLE, SHORT AND UNSTRUCTURED CRS.

NEW NIZK NOTIONS:

- CRS-LESS NIZKs
- ALMOST-EVERYWHERE-EXTRACTABILITY