Cosmic Security

Benjamin Chan

Cornell Tech

@ Reading Group, NYU. April 12 2022

Joint work with Cody Freitag & Rafael Pass

Observation

We may be theorists, but at the end of the day, we ideally want to build cryptosystems secure for real-world attackers.

Recall:
the *(Physical) Extended Church Turing Hypothesis*

Recall:
the *(Physical) Extended Church Turing Hypothesis*

All Real-Life Attackers

Recall:
the *(Physical) Extended Church Turing Hypothesis*

All Real-Life Attackers can be simulated by

$\approx$

Recall:
the *(Physical) Extended Church Turing Hypothesis*

All Real-Life Attackers | can be simulated by | (nu)PPT Turing Machines

$\approx$

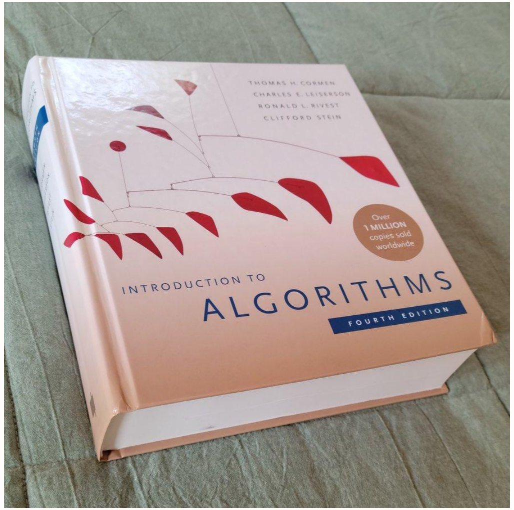

computational

How Cryptography Models Attacks Today

Suppose ∃ (nu)PPT **A** that breaks an encryption scheme **C**:

$$\begin{tikzcd} \mathbf{C} \arrow[r, shift left=2, "\mathrm{Enc}(m_b)"] & \boxed{\mathbf{A}} \arrow[l, shift left=6, "m_0, m_1"'] \arrow[l, shift left=-2, "b"] \end{tikzcd}$$

(nu)PPT (interactive) Turing Machine

Then ∃ (nu)PPT **A'** that inverts a one way function C':

$$\begin{tikzcd} \mathbf{C'} \arrow[r, shift left=2, "f(x)"] & \mathbf{A'} \arrow[l, shift left=2, "f^{-1}(f(x))"] \end{tikzcd}$$

Contradiction! (assuming security of the OWF for (nu)PPT attackers.)

How computational Cryptography Models Attacks Today

Suppose ∃ (nu)PPT **A** that breaks an encryption scheme **C**:

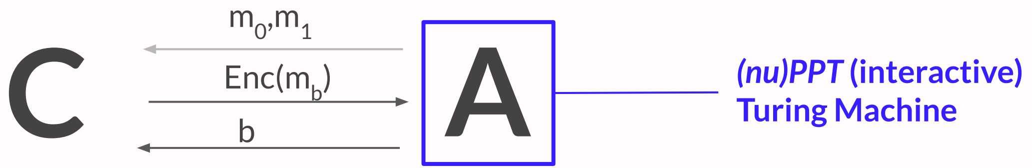

Then (black-box) (nu)PPT $\mathbf{R}^A$ inverts a one way function C’:

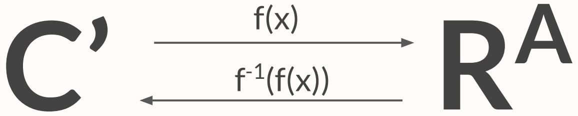

Contradiction! (assuming security of the OWF for (nu)PPT attackers.)

How *computational* Cryptography Models Attacks Today

Suppose ∃ (nu)PPT **A** that breaks an encryption scheme **C**:

$$\mathbf{C} \xleftarrow{m_0, m_1} \quad \xrightarrow{Enc(m_b)} \quad \xleftarrow{b} \boxed{\mathbf{A}}$$

(nu)PPT (interactive) Turing Machine

Then (black-box) (nu)PPT $\mathbf{R}^{A}$ inverts a one way function C':

$$\mathbf{C'} \xrightarrow{f(x)} \quad \xleftarrow{f^{-1}(f(x))} \mathbf{R}^{\boxed{A}}$$

Classically, R can rewind or restart the attacker, because it is just a piece of code!

Contradiction! (assuming security of the OWF for (nu)PPT attackers.)

A very productive study

From OWFs, we can build,

- Hardcore Bits [GL89]
- Pseudorandom Generators [HILL99]
- Private-Key Encryption [GGM85+86]
- Commitment Schemes [Nao91]
- Digital Signatures [Lam79, Rom90, Mer90]
- Zero-Knowledge Proofs for all of NP [GMW86]

A very productive study

From OWFs, we can build,

- Hardcore Bits [GL89]
- Pseudorandom Generators [HILL99]
- Private-Key Encryption [GGM85+86]
- Commitment Schemes [Nao91]
- Digital Signatures [Lam79, Rom90, Mer90]
- Zero-Knowledge Proofs for all of NP [GMW86]

But throughout, an explicit assumption:
The attacker is a rewindable algorithm that can be reused!

So, what if the physical extended Church Turing Hypothesis is wrong?

Enter Quantum Computing.

No-Cloning Theorem:
can’t copy qubits without collapsing superposition.

How computational Cryptography Models Attacks Today

Suppose ∃ (nu)PPT **A** that breaks an encryption scheme **C**:

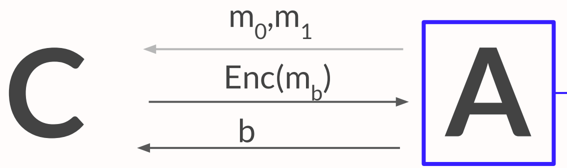

(nu)PPT (interactive)
Turing Machine

Then (black-box) $\mathbf{R}^A$ inverts a one way function C':

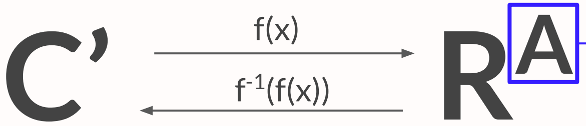

Classically, R can rewind or restart the attacker, because it is just a piece of code!

Contradiction! (assuming security of the OWF for (nu)PPT attackers.)

How computational Cryptography Models Attacks Today

Suppose ∃ (nu)QPT **A** that breaks an encryption scheme **C**:

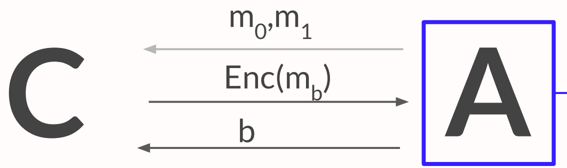

(nu)QPT (interactive) Turing Machine

Then (black-box) (nu)QPT $\mathbf{R}^A$ inverts a one way function C':

Quantum rewinding requires extra care

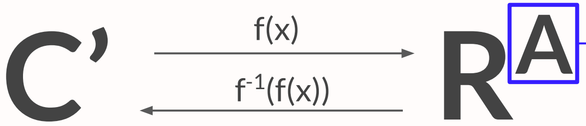

~~Classically, R can rewind or restart the attacker, because it is just a piece of code!~~

Contradiction! (assuming security of the OWF for (nu)QPT attackers.)

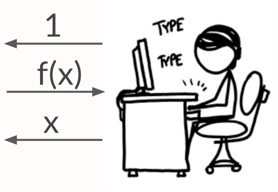

Quantum rewinding requires extra care

~~*Classically, R can rewind or restart the attacker, because it is just a piece of code!*~~

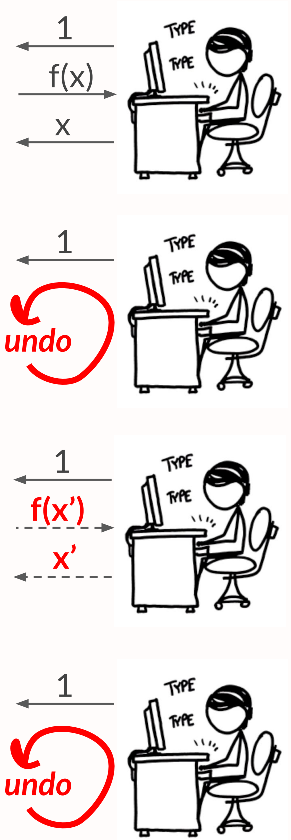

Hard to “reset” quantum resources that have been measured!

Quantum rewinding requires extra care

~~*Classically, R can rewind or restart the attacker, because it is just a piece of code!*~~

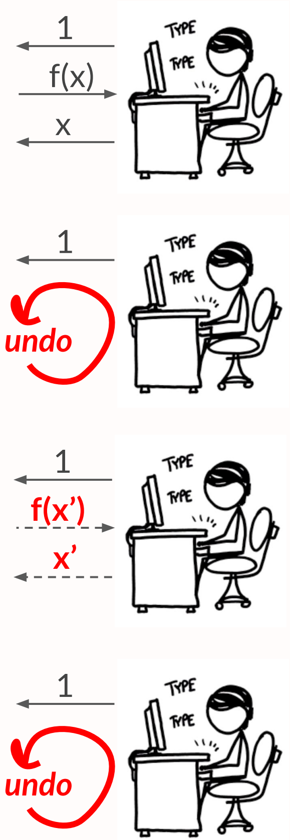

Need to revisit:

- Hardcore Bits from OWFs [AC01*]
- PRGs from OWFs [Aar09]
- Digital Signatures [Son14]
- Zero Knowledge [Wat09*]
- etc

Quantum rewinding requires extra care

~~Classically, R can rewind or restart the attacker, because it is just a piece of code!~~

Hard to "reset" quantum resources that have been measured!

Another assumption: if reduction uses qubits, then it also assumes that QPT is "feasible"! (is it?)

Need to revisit:

- Hardcore Bits from OWFs [AC01*]
- PRGs from OWFs [Aar09]
- Digital Signatures [Son14]
- Zero Knowledge [Wat09*]
- etc

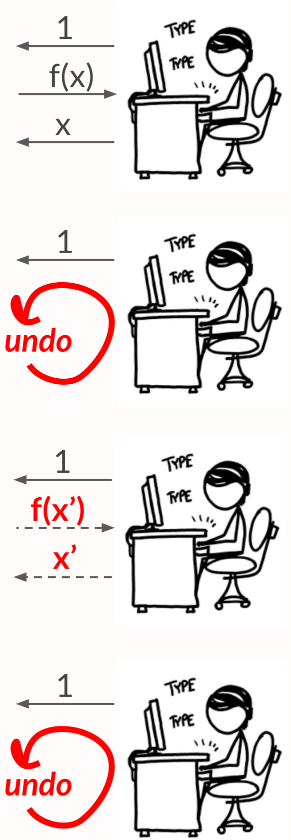

Hard to "reset" quantum resources that have been measured!

Quantum rewinding requires extra care

~~*Classically, R can rewind or restart the attacker, because it is just a piece of code!*~~

Claim: currently, cryptography depends heavily on *extended physical Church Turing* assumptions

"It must be possible for a scientific theory to be refuted by experience." - Karl Popper

- Not clear how to decisively prove
  that any version of the Church Turing Hypothesis is correct.

A

Could easily imagine "fully stateful" attackers lurking somewhere in the "Cosmos" that we do not know how to "restart" or "rewind"

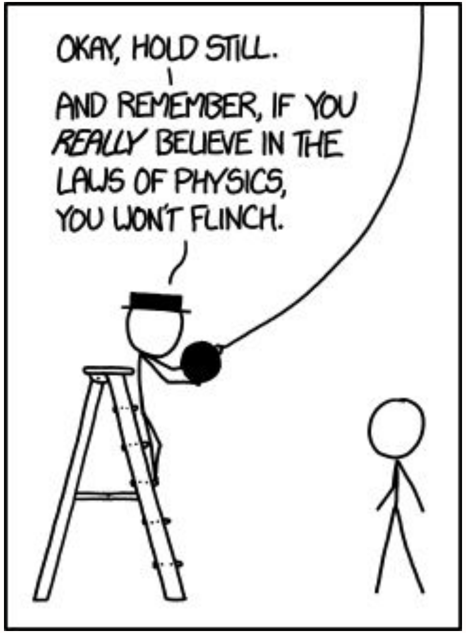

Claim: currently, cryptography depends heavily on *extended physical Church Turing* assumptions

"It must be possible for a scientific theory to be refuted by experience." - Karl Popper

- Not clear how to decisively prove that any version of the Church Turing Hypothesis is correct. Seems non-ideal!

$\boxed{A}$

Could easily imagine "fully stateful" attackers lurking somewhere in the "Cosmos" that we do not know how to "restart" or "rewind"

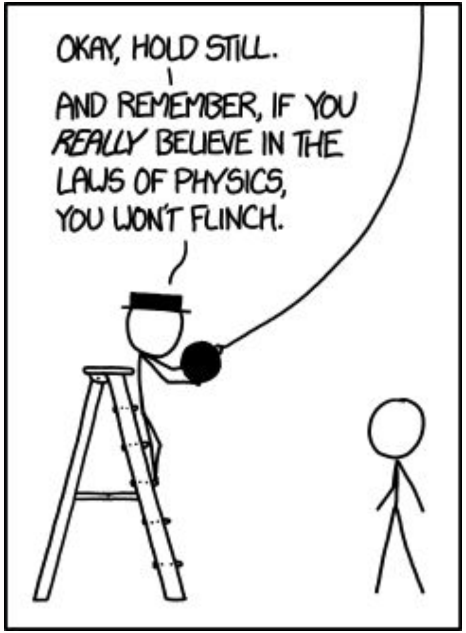

This talk:

Can we build a reduction-based
theory of computational cryptography
with minimal assumptions on
the Nature of real-world attackers?

This talk:

Can we build a reduction-based
theory of computational cryptography
with minimal assumptions on
the Nature of real-world attackers?

No rewinding

This talk:

Can we build a reduction-based
theory of computational cryptography
with minimal assumptions on
the Nature of real-world attackers?

No rewinding

Letting the attacker "keep state"
between uses.
(e.g. maybe your genius neighbor can
decrypt things)

Our work: Cosmic Security (Informal)

Suppose PPT **A** uses **Nat** to 'break' **C** (e.g. breaking an encryption scheme):

C $\xrightarrow{\text{Enc(m)}}$ / $\xleftarrow{\text{m}}$ **A** $\longleftrightarrow$ **Nat**

Can be stateful
Can be unbounded

Then ∃ PPT **A'** that uses **Nat** to 'break' **C'** (e.g. inverting a OWF)

C' $\xrightarrow{f(x)}$ / $\xleftarrow{f^{-1}(f(x))}$ **A'** $\longleftrightarrow$ **Nat**

A' can no longer "undo" an interaction with Nat

Contradiction! (assuming security of the OWF vs PPT+**Nat**)

Our work: Cosmic Security (Informal)

Suppose PPT **A** uses **Nat** to 'break' **C** (e.g. breaking an encryption scheme):

$$\mathbf{C} \rightleftarrows \mathbf{A} \longleftrightarrow \mathbf{Nat}$$

Can be stateful
Can be unbounded

Minimal assumption: Nat won't just "shut down" after one use

Then ∃ PPT **A'** that uses **Nat** to 'break' **C'** (e.g. inverting a OWF)

$$\mathbf{C'} \underset{f^{-1}(f(x))}{\overset{f(x)}{\rightleftarrows}} \mathbf{A'} \longleftrightarrow \mathbf{Nat}$$

A' can no longer "undo" an interaction with Nat

Contradiction! (assuming security of the OWF vs PPT+**Nat**)

Roadmap

1. ~~Motivation~~ (*10min*)
2. Defining Cosmic Security (*15min*)
3. Properties of Cosmic Security: a Sanity Check
   a. Composition, Black-box reductions (*5min*)
4. Summary of Key Results
   a. Feasibilities and Impossibilities (20*min*)
5. Other Notions of Cosmic Security (10min)
6. Conclusion (5min)

Quick Comparisons

Relativized Reductions

- A *relativized reduction* gives attackers A^O access to some arbitrary oracle O
- O is modeled as a (perhaps uncomputable) function
- Cosmic reductions can be viewed as *relativized reductions* for *stateful, interactive oracles* O (in contrast to a *non-interactive, stateless* oracle).

Universal Composability [Canetti00]

- Cosmic security is *syntactically* similar to UC with unbounded environments
- *Semantically* very different: our notion is reduction-based & computational. (For instance, UC *security proofs* can rewind the environment [e.g. CLP10])

Roadmap

1. ~~Motivation~~ (*10min*)
2. **Defining Cosmic Security** **(*15min*)**
3. Properties of Cosmic Security: a Sanity Check
    a. Composition, Black-box reductions (*5min*)
4. Summary of Key Results
    a. Feasibilities and Impossibilities (20*min*)
5. Other Notions of Cosmic Security (10min)
6. Conclusion (5min)

Defining Cosmic Security

A $\leftrightarrow$ **Nat**

"Attacker"
uniform **PPT**

"Nature"
nonuniform **any choice of runtime**
Stateful!

"Cosmic Adversary" (A, Nat)

Cosmic Adversaries

$A \longleftrightarrow Nat$

"Attacker"
uniform **PPT**

"Nature"
nonuniform **any choice of runtime**
Stateful

"Cosmic Adversary" (A, Nat)

A classic attacker
that uses Nature
to break some scheme

Cosmic Adversaries

A ⟷ **Nat**

"Attacker"
uniform **PPT**

"Nature"
nonuniform **any choice of runtime**
Stateful

"Cosmic Adversary" (A, Nat)

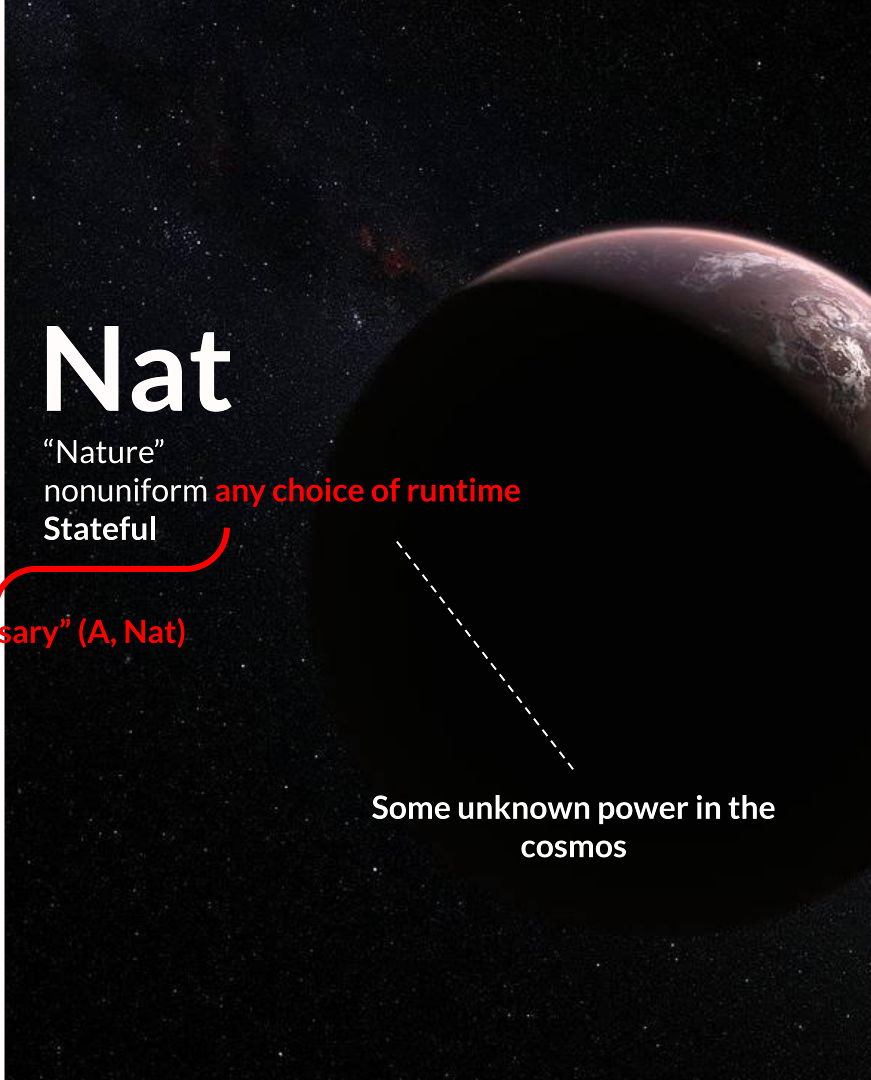

Some unknown power in the cosmos

Cosmic Security Game

C ⟷ **A** ⟷ **Nat**

C
"Challenger"
uniform **PPT**
outputs "win"/"lose"

A
"Attacker"
uniform **PPT**

Nat
"Nature"
nonuniform **any choice of runtime**
Stateful

"Cosmic Adversary" (A, Nat)

Cosmic Security Game

$C \longleftrightarrow A \longleftrightarrow$ **Nat**

"Challenger"
uniform **PPT**
outputs "win"/"lose"

"Attacker"
uniform **PPT**

"Nature"
nonuniform **any choice of runtime**
Stateful

"Cosmic Adversary" (A, Nat)

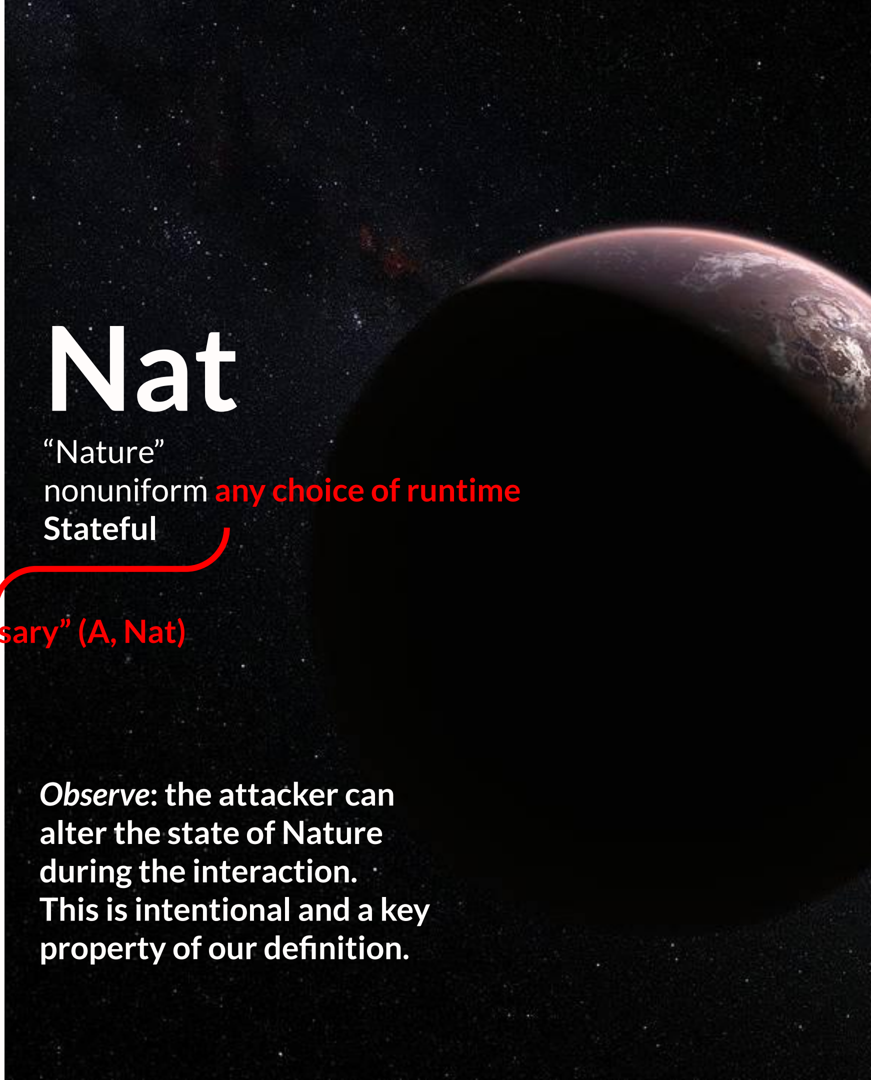

***Observe*: the attacker can alter the state of Nature during the interaction. This is intentional and a key property of our definition.**

Problem! Nature can be invoked many times, "*statefully*"

time

WIN. C ↔ A ↔

LOSE. C ↔ A ↔

LOSE. C ↔ A ↔

LOSE. C ↔ A ↔

LOSE. C ↔ A ↔

Useful adversaries win repeatedly.

No point in inverting a OWF only once.

time

WIN. $C \leftrightarrow A$

WIN. $C \leftrightarrow A$

WIN. $C \leftrightarrow A$

WIN. $C \leftrightarrow A$

WIN. $C \leftrightarrow A$

Classically, we always have the option of re-running A from scratch, with independent coins.

What is a weaker assumption we could make, that might still be useful?

We consider only (**A**, **Nat**) that "wins repeatedly" over time.

time

WIN. C $\leftrightarrow$ A $\leftrightarrow$

WIN. C $\leftrightarrow$ A $\leftrightarrow$

WIN. C $\leftrightarrow$ A $\leftrightarrow$

WIN. C $\leftrightarrow$ A $\leftrightarrow$

It doesn't matter that Nat is "stateful" if it only wins once.

robust winning

the past

↔

42

↔

C ↔ A ↔

C ↔ A ↔

WIN. C ↔ A ↔ Nat

time

(**A**, **Nat**) wins for **C** (that flips fresh coins) regardless of any interactions that **Nat** had in the past!

robust winning

42

the past

ρ

$C \leftrightarrow A \leftrightarrow$

$C \leftrightarrow A \leftrightarrow$

WIN. $C \leftrightarrow A \leftrightarrow$ Nat(ρ)

Interaction prefix ρ: a transcript of messages previously sent to Nat before the beginning of execution, including coins flipped by Nat.

(**A**, **Nat**) wins for **C** (that flips fresh coins) regardless of any interactions that **Nat** had in the past!

robust winning

42

C ↔ A ↔

C ↔ A ↔

the past

ρ

WIN. C ↔ A ↔ Nat(ρ)

Interaction prefix ρ: a transcript of messages previously sent to Nat before the beginning of execution, including coins flipped by Nat.

Definition: (A, Nat) has **robust advantage** a(.) for C, if ∀ *interaction prefixes* ρ, ∀ λ: Pr[(**A**, **Nat**(ρ)) wins **C**] ≥ a(λ)

robust winning

42

the past

ρ

Interaction prefix ρ: a transcript of messages previously sent to Nat before the beginning of execution, including coins flipped by Nat.

C ⟷ A ⟷

C ⟷ A ⟷

WIN. C ⟷ A ⟷ Nat(ρ)

Definition: (A, Nat) has **robust advantage** a(.) for C, if ∀ *interaction prefixes* ρ, ∀ λ: Pr[(**A**, **Nat**(ρ)) wins **C**] ≥ a(λ)

A weak notion!
Can win in a different way for each prefix.

Cosmic Security (Final)

∃ an **ε-cosmic reduction** from **C** to **C'** if ∀ PPT **A**, ∃ PPT **A'** s.t. ∀**Nat**:

Suppose (**A**, **Nat**) has **robust advantage** a(·) for **C**

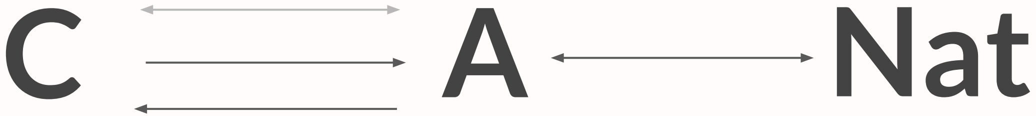

Cosmic Security (Final)

∃ an **ε-cosmic reduction** from **C** to **C’** if ∀ PPT **A**, ∃ PPT **A’** s.t. ∀**Nat**:

Suppose (**A**, **Nat**) has **robust advantage a(·)** for **C**

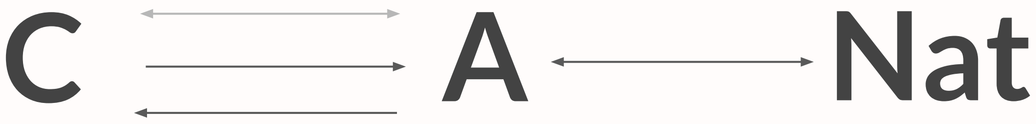

Then (**A’**, **Nat**) has **robust advantage ε(·,a(·))** for **C’**.

Cosmic Security (Final)

∃ an **ε-cosmic reduction** from **C** to **C'** if ∀ PPT **A**, ∃ PPT **A'** s.t. ∀**Nat**:

Suppose (**A**, **Nat**) has **robust advantage a(·)** for **C**

$$\begin{tikzcd}
C \arrow[r, leftrightarrow, shift left=2] \arrow[r, shift left=0] & A \arrow[l, shift left=2] \arrow[r, leftrightarrow] & Nat
\end{tikzcd}$$

Then (**A'**, **Nat**) has **robust advantage ε(·,a(·))** for **C'**.

A' can now use the fact that (A, Nat) wins "repeatedly".

$$\begin{tikzcd}
C' \arrow[r, shift left, "f(x)"] & A' \arrow[l, shift left, "f^{-1}(f(x))"] \arrow[r, leftrightarrow] & Nat
\end{tikzcd}$$

Cosmic Security (Final)

∃ an **ε-cosmic reduction** from **C** to **C'** if ∀ PPT **A**, ∃ PPT **A'** s.t. ∀**Nat**:

Suppose (**A**, **Nat**) has **robust advantage a(·)** for **C**

$$\begin{tikzcd} \mathbf{C} \arrow[r, leftrightarrow, shift left=2] \arrow[r] & \mathbf{A} \arrow[l, shift left=2] \arrow[r, leftrightarrow] & \mathbf{Nat} \end{tikzcd}$$

Then (**A'**, **Nat**) has ~~robust~~ **advantage ε(·,a(·))** for **C'**.

A' can now use the fact that (A, Nat) wins "repeatedly".

Equivalent definition!

$$\begin{tikzcd} \mathbf{C'} \arrow[r, shift left, "f(x)"] & \mathbf{A'} \arrow[l, shift left, "f^{-1}(f(x))"] \arrow[r, leftrightarrow] & \mathbf{Nat} \end{tikzcd}$$

Cosmic Security (Final)

∃ an **ε-cosmic reduction** from **C** to **C'** if ∀ PPT **A**, ∃ PPT **A'** s.t. ∀**Nat**:

Suppose (**A**, **Nat**) has **robust advantage a(·)** for **C**

C ⇄ **A** ⟷ **Nat**

Then (**A'**, **Nat**) has ~~robust~~ **advantage ε(·,a(·))** for **C'**.

Equivalent definition!

A' can now use the fact that (A, Nat) wins "repeatedly".

$$C' \underset{f^{-1}(f(x))}{\overset{f(x)}{\rightleftarrows}} A' \longleftrightarrow Nat$$

Rest of this talk:
this is the right definition to work on!

Roadmap

1. ~~Motivation~~ (*10min*)
2. ~~Defining Cosmic Security~~ (*15min*)
3. **Properties of Cosmic Security: a Sanity Check**
    a. **Composition, Black-box reductions** **(*5min*)**
4. Summary of Key Results
    a. Feasibilities and Impossibilities (20*min*)
5. Other Notions of Cosmic Security (10min)
6. Conclusion (5min)

Two Nice Properties

<u>***Def***</u>**:** *ε-cosmic reduction* from **C** to **C'**:

∀ PPT attackers **A**, ∃ PPT **A'** s.t. ∀ **Nat**:

(**A**, **Nat**) has *robust adv* a(.) for **C**

$\Downarrow$

(**A'**, **Nat**) has *robust adv* ε(.,a(.)) for **C'**.

What makes us think that cosmic security is a good definition?

1. <u>Composability</u>: Cosmic reductions are composable; that is, *if* **C** *reduces to* **C'** *and* **C'** *reduces to* **C''**, *then* **C** *reduces to* **C''**.

Two Nice Properties

<u>Def</u>: *ε-cosmic reduction* from **C** to **C'**:

∀ PPT attackers **A**, ∃ PPT **A'** s.t. ∀ **Nat**:

(**A**, **Nat**) has *robust adv* a(.) for **C**
⇓
(**A'**, **Nat**) has *robust adv* ε(.,a(.)) for **C'**.

What makes us think that
cosmic security is a good definition?

1. <u>Composability</u>: Cosmic reductions are composable; that is, *if* **C** *reduces to* **C'** *and* **C'** *reduces to* **C''**, *then* **C** *reduces to* **C''**.

follows nicely from definition!

Two Nice Properties

What makes us think that cosmic security is a good definition?

1. <u>Composability</u>: Cosmic reductions are composable; that is, *if* **C** *reduces to* **C'** *and* **C'** *reduces to* **C''**, *then* **C** *reduces to* **C''**.

2. <u>Dummy Lemma</u>: Regular cosmic reductions are *equivalent* to black-box cosmic reductions.

<u>Def</u>: *ε-cosmic reduction* from **C** to **C'**:
∀ PPT attackers **A**, ∃ PPT **A'** s.t. ∀ **Nat**:

(**A**, **Nat**) has *robust adv* a(.) for **C**

⇓

(**A'**, **Nat**) has *robust adv* ε(.,a(.)) for **C'**.

⇕

<u>Def</u>: *ε-cosmic black-box reduction* from **C** to **C'** if ∃ PPT **R** s.t. ∀ (**A**,**Nat**):

(**A**, **Nat**) has *robust adv* a(.) for **C**

⇓

($\mathbf{R^A}$, **Nat**) has *robust adv* ε(.,a(.)) for **C'**.

Two Nice Properties

What makes us think that
cosmic security is a good definition?

1. <u>Composability</u>: Cosmic reductions are composable; that is, *if* **C** *reduces to* **C'** *and* **C'** *reduces to* **C''**, *then* **C** *reduces to* **C''**.

2. <u>Dummy Lemma</u>: Regular cosmic reductions are *equivalent* to black-box cosmic reductions.

Let's quickly intuit this to understand cosmic security better

Def: *ε-cosmic reduction* from **C** to **C'**:
∀ PPT attackers **A**, ∃ PPT **A'** s.t. ∀ **Nat**:

(**A**, **Nat**) has *robust adv* a(.) for **C**

⇓

(**A'**, **Nat**) has *robust adv* ε(.,a(.)) for **C'**.

⇕

Def: *ε-cosmic black-box reduction* from **C** to **C'** if ∃ PPT **R** s.t. ∀ (**A**,**Nat**):

(**A**, **Nat**) has *robust adv* a(.) for **C**

⇓

($\mathbf{R^A}$, **Nat**) has *robust adv* ε(.,a(.)) for **C'**.

Two Nice Properties

What makes us think that cosmic security is a good definition?

1. Composability: Cosmic reductions are composable; that is, *if* **C** *reduces to* **C'** *and* **C'** *reduces to* **C''**, *then* **C** *reduces to* **C''**.

2. Dummy Lemma: Regular cosmic reductions are *equivalent* to black-box cosmic reductions.

$\underbrace{\hphantom{xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx}}$

Let's quickly intuit this to understand cosmic security better

Def: *ε-cosmic reduction* from **C** to **C'**:
∀ PPT attackers **A**, ∃ PPT **A'** s.t. ∀ **Nat**:

(**A**, **Nat**) has *robust adv* a(.) for **C**

⇓

(**A'**, **Nat**) has *robust adv* ε(.,a(.)) for **C'**.

Easy direction

Def: ε-cosmic *black-box reduction* from **C** to **C'** if ∃ PPT **R** s.t. ∀ (**A**,**Nat**):

(**A**, **Nat**) has *robust adv* a(.) for **C**

⇓

($\mathbf{R}^{\mathbf{A}}$, **Nat**) has *robust adv* ε(.,a(.)) for **C'**.

Two Nice Properties

What makes us think that cosmic security is a good definition?

1. Composability: Cosmic reductions are composable; that is, *if* **C** *reduces to* **C'** *and* **C'** *reduces to* **C''**, *then* **C** *reduces to* **C''**.

2. Dummy Lemma: Regular cosmic reductions are *equivalent* to black-box cosmic reductions.

Let's quickly intuit this to understand cosmic security better

Def: *ε-cosmic reduction* from **C** to **C'**:
∀ PPT attackers **A**, ∃ PPT **A'** s.t. ∀ **Nat**:

(**A**, **Nat**) has *robust adv* a(.) for **C**

$\Downarrow$

(**A'**, **Nat**) has *robust adv* ε(.,a(.)) for **C'**.

$\Downarrow$ want to show

Def: *ε-cosmic black-box reduction* from **C** to **C'** if ∃ PPT **R** s.t. ∀ (**A**,**Nat**):

(**A**, **Nat**) has *robust adv* a(.) for **C**

$\Downarrow$

(R^A, **Nat**) has *robust adv* ε(.,a(.)) for **C'**.

From Existential to Constructive Reductions

$$\mathbf{C} \longleftrightarrow \underbrace{\mathbf{A} \longleftrightarrow \mathbf{Nat}}_{\textit{robust adv}\ \mathrm{a}(.)\ \text{for}\ \mathbf{C}}$$

Def**:** *ε-cosmic reduction* from **C** to **C'**:
∀ PPT attackers **A**, ∃ PPT **A'** s.t. ∀ **Nat**:

(**A**, **Nat**) has *robust adv* a(.) for **C**

⇓

(**A'**, **Nat**) has *robust adv* ε(.,a(.)) for **C'**.

⇓ want to show

Def**:** *ε-cosmic black-box reduction*
from **C** to **C'** if ∃ PPT **R** s.t. ∀ (**A**,**Nat**):

(**A**, **Nat**) has *robust adv* a(.) for **C**

⇓

($\mathbf{R}^{\mathbf{A}}$, **Nat**) has *robust adv* ε(.,a(.)) for **C'**.

From Existential to Constructive Reductions

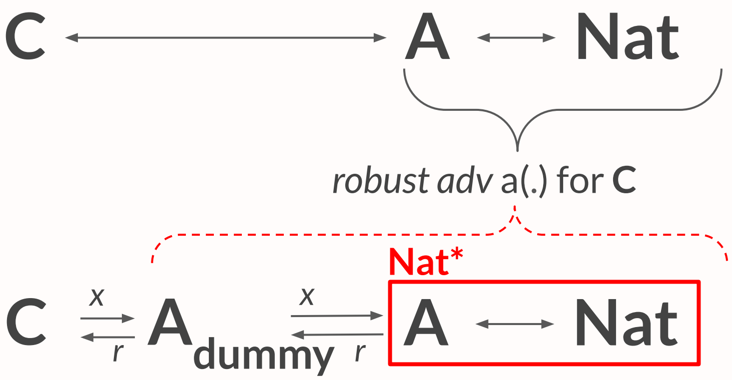

***Def*:** ε-*cosmic reduction* from **C** to **C'**:

∀ PPT attackers **A**, ∃ PPT **A'** s.t. ∀ **Nat**:

(**A**, **Nat**) has *robust adv* a(.) for **C**

⇩

(**A'**, **Nat**) has *robust adv* ε(.,a(.)) for **C'**.

want to show

***Def*:** ε-*cosmic black-box reduction*

from **C** to **C'** if ∃ PPT **R** s.t. ∀ (**A**,**Nat**):

(**A**, **Nat**) has *robust adv* a(.) for **C**

⇩

($\mathbf{R}^{\mathbf{A}}$, **Nat**) has *robust adv* ε(.,a(.)) for **C'**.

From Existential to Constructive Reductions

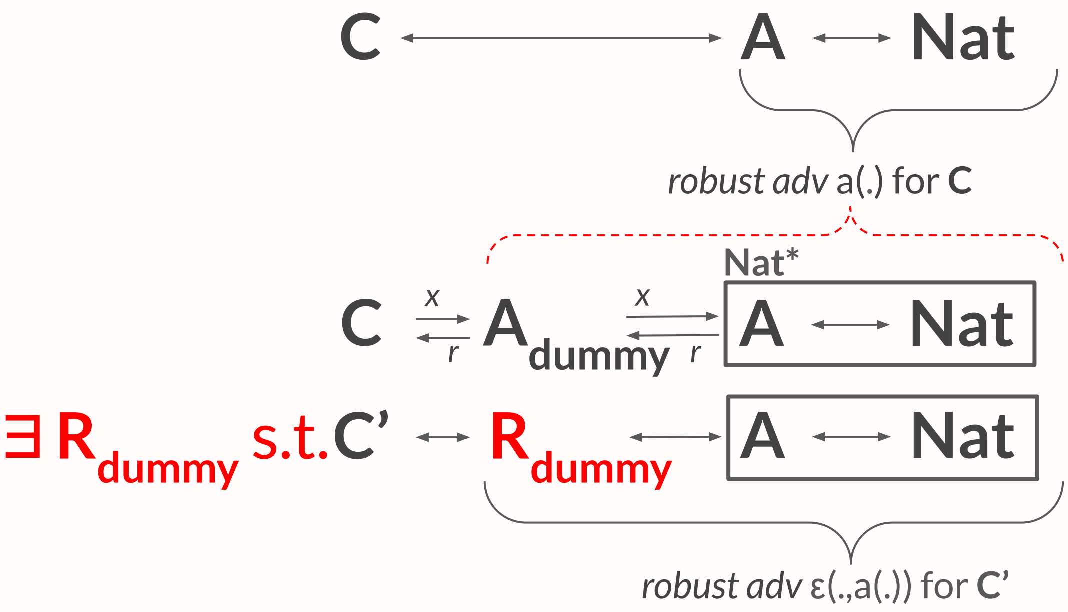

Def: *ε-cosmic reduction* from **C** to **C'**:
∀ PPT attackers **A**, ∃ PPT **A'** s.t. ∀ **Nat**:

(**A**, **Nat**) has *robust adv* a(.) for **C**

⇓

(**A'**, **Nat**) has *robust adv* ε(.,a(.)) for **C'**.

want to show

Def: *ε-cosmic black-box reduction* from **C** to **C'** if ∃ PPT **R** s.t. ∀ (**A**,**Nat**):

(**A**, **Nat**) has *robust adv* a(.) for **C**

⇓

($\mathbf{R}^{\mathbf{A}}$, **Nat**) has *robust adv* ε(.,a(.)) for **C'**.

From Existential to Constructive Reductions

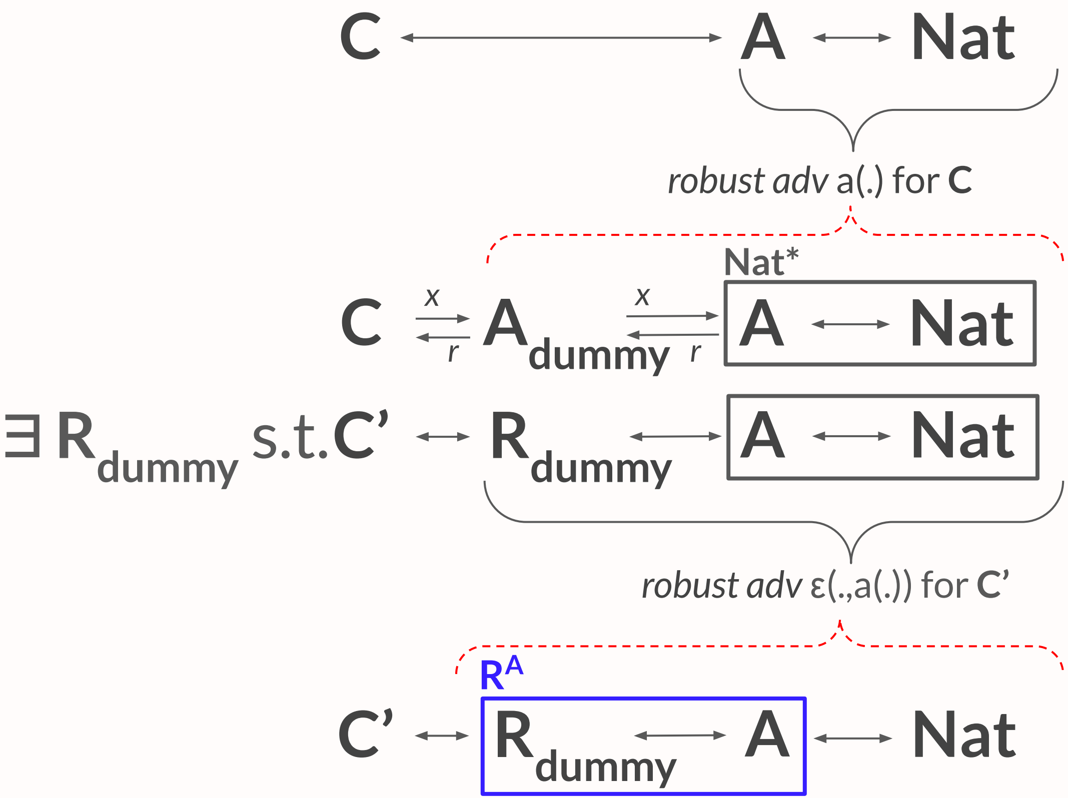

Def: *ε-cosmic reduction* from **C** to **C'**:
∀ PPT attackers **A**, ∃ PPT **A'** s.t. ∀ **Nat**:

(**A**, **Nat**) has *robust adv* a(.) for **C**

⇓

(**A'**, **Nat**) has *robust adv* ε(.,a(.)) for **C'**.

want to show

Def: *ε-cosmic black-box reduction* from **C** to **C'** if ∃ PPT **R** s.t. ∀ (**A**,**Nat**):

(**A**, **Nat**) has *robust adv* a(.) for **C**

⇓

($\mathbf{R^A}$, **Nat**) has *robust adv* ε(.,a(.)) for **C'**.

From Existential to Constructive Reductions

Now,
need to argue
robustness.
Nontrivial;
argued in paper!

$\exists\, \mathbf{R}_{\mathbf{dummy}}$ s.t.

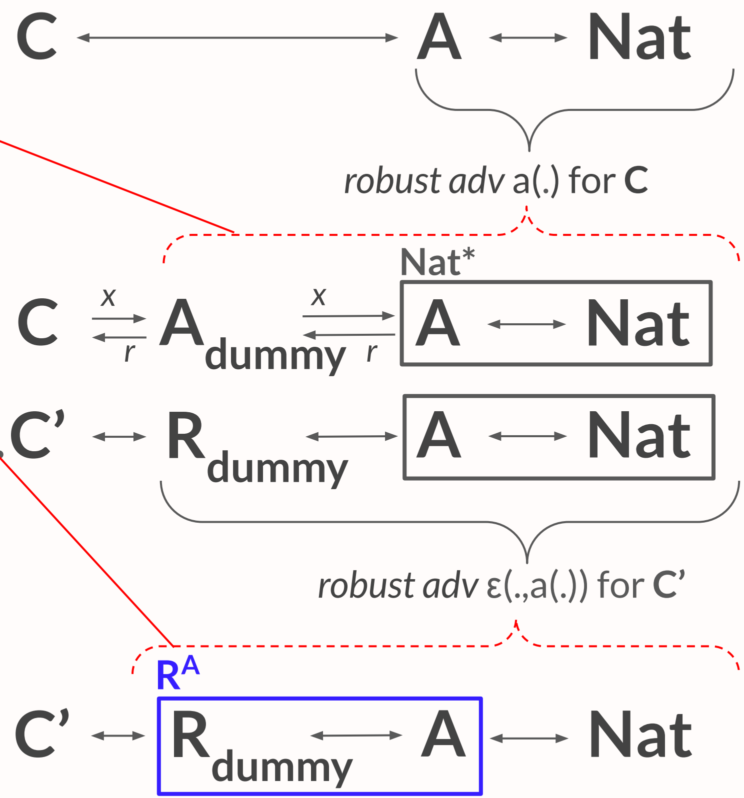

***Def*:** *ε-cosmic reduction* from **C** to **C'**:
∀ PPT attackers **A**, ∃ PPT **A'** s.t. ∀ **Nat**:

(**A**, **Nat**) has *robust adv* a(.) for **C**
⇓
(**A'**, **Nat**) has *robust adv* ε(.,a(.)) for **C'**.

want to show

***Def*:** *ε-cosmic black-box reduction* from **C** to **C'** if ∃ PPT **R** s.t. ∀ (**A**,**Nat**):

(**A**, **Nat**) has *robust adv* a(.) for **C**
⇓
($\mathbf{R}^{\mathbf{A}}$, **Nat**) has *robust adv* ε(.,a(.)) for **C'**.

From Existential to Constructive Reductions

$\mathbf{C} \longleftrightarrow \mathbf{A} \longleftrightarrow \mathbf{Nat}$

$\underbrace{\mathbf{A} \longleftrightarrow \mathbf{Nat}}$

robust adv a(.) for **C**

Dummy Lemma:
Suffices to show a cosmic reduction for $\mathbf{A}_{\mathbf{dummy}}$

$\mathbf{C} \underset{r}{\overset{x}{\rightleftarrows}} \mathbf{A}_{\mathbf{dummy}} \underset{r}{\overset{x}{\rightleftarrows}} \overset{\text{Nat*}}{\boxed{\mathbf{A} \longleftrightarrow \mathbf{Nat}}}$

$\exists\, \mathbf{R}_{\mathbf{dummy}}$ s.t. $\mathbf{C'} \longleftrightarrow \mathbf{R}_{\mathbf{dummy}} \longleftrightarrow \boxed{\mathbf{A} \longleftrightarrow \mathbf{Nat}}$

robust adv ε(.,a(.)) for **C'**

$\mathbf{C'} \longleftrightarrow \overset{\mathbf{R^A}}{\boxed{\mathbf{R}_{\mathbf{dummy}} \longleftrightarrow \mathbf{A}}} \longleftrightarrow \mathbf{Nat}$

Def: ε-*cosmic reduction* from **C** to **C'**:
∀ PPT attackers **A**, ∃ PPT **A'** s.t. ∀ **Nat**:

(**A**, **Nat**) has *robust adv* a(.) for **C**

⇓

(**A'**, **Nat**) has *robust adv* ε(.,a(.)) for **C'**.

⇓ ⇑

Def: ε-*cosmic black-box reduction*
from **C** to **C'** if ∃ PPT **R** s.t. ∀ (**A**,**Nat**):

(**A**, **Nat**) has *robust adv* a(.) for **C**

⇓

($\mathbf{R^A}$, **Nat**) has *robust adv* ε(.,a(.)) for **C'**.

Recap So Far: A Necessary, Natural Definition

Why do we like Cosmic Security?

1. Acknowledges the Unknown: allows Nature to be *stateful* in a way previously not considered, but that we are now compelled to acknowledge.

Moreover, the definition is Well-Behaved.

2. Composability: Cosmic reductions are composable; that is, *if* **C** *reduces to* **C'** *and* **C'** *reduces to* **C''**, *then* **C** *reduces to* **C''**.

3. Dummy Lemma: Regular cosmic reductions are *equivalent* to black-box cosmic reductions.

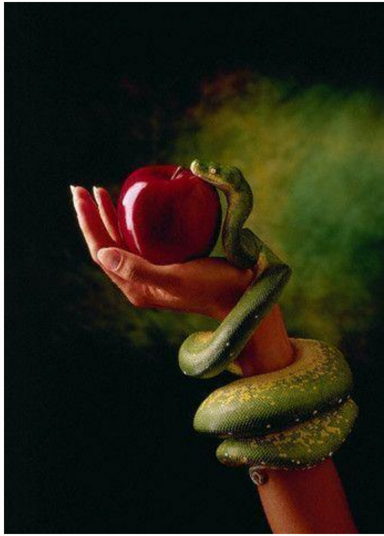

Recap So Far: A Necessary, Natural Definition

Why do we like Cosmic Security?

1. <u>Acknowledges the Unknown</u>: allows Nature to be *stateful* in a way previously not considered, but that we are now compelled to acknowledge.

Moreover, the definition is Well-Behaved.

2. <u>Composability</u>: Cosmic reductions are composable; that is, *if* **C** *reduces to* **C'** *and* **C'** *reduces to* **C''**, *then* **C** *reduces to* **C''**.

3. <u>Dummy Lemma</u>: Regular cosmic reductions are *equivalent* to black-box cosmic reductions.

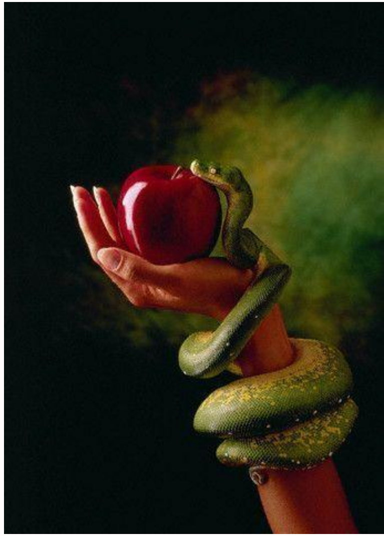

NEXT UP, The Real Test:
What Can We Build?

Roadmap

1. ~~Motivation~~ *(10min)*
2. ~~Defining Cosmic Security~~ *(15min)*
3. ~~Properties of Cosmic Security: a Sanity Check~~
    a. ~~Composition, Black-box reductions~~ *(5min)*
4. **Summary of Key Results**
    a. **Feasibilities and Impossibilities** ***(20min)***
5. Other Notions of Cosmic Security (10min)
6. Conclusion (5min)

Key Results: Feasibilities and Impossibilities

<u>Thm 1</u> (*Feasibility*): classical 1-shot straight-line black-box reductions imply cosmic reductions.

<u>Corollaries</u>: PRFs/SKE/Commitments/Witness Indistinguishability/PRG Length Extension

> *Not surprising, since it doesn't matter that Nature is stateful: a 1-shot reduction only uses Nature once.*

Key Results:
Feasibilities and Impossibilities

<u>Thm 1</u> (*Feasibility*): classical 1-shot straight-line black-box reductions imply cosmic reductions.

<u>Corollaries</u>: PRFs/SKE/Commitments/Witness Indistinguishability/PRG Length Extension

> *Not surprising, since it doesn't matter that Nature is stateful: a 1-shot reduction only uses Nature once.*

New 1-shot straight-line black-box proof for WI! (See Paper)

Key Results:
Feasibilities and Impossibilities

<u>Thm 1</u> (*Feasibility*): classical 1-shot straight-line black-box reductions imply cosmic reductions.

<u>Corollaries</u>: PRFs/SKE/Commitments/Witness Indistinguishability/PRG Length Extension

> Not surprising, since it doesn't matter that Nature is stateful: a 1-shot reduction only uses Nature once.

What about reductions that use the attacker multiple times?

Key Results:
Feasibilities and Impossibilities

Thm 1 (*Feasibility*): classical 1-shot straight-line black-box reductions imply cosmic reductions.

Corollaries: PRFs/SKE/Commitments/Witness Indistinguishability/PRG Length Extension

> Not surprising, since it doesn't matter that Nature is stateful: a 1-shot reduction only uses Nature once.

Thm 2 (*Impossibility*): Hardness amplification of arbitrary weak OWFs via direct product, using only black-box access to the OWF, is impossible.

Thm 3 (*Impossibility*): A Goldreich Levin Theorem, where the reduction has only black-box access to the OWF, is impossible.

Key Results: Feasibilities and Impossibilities

Thm 1 (*Feasibility*): classical 1-shot straight-line black-box reductions imply cosmic reductions.

Corollaries: PRFs/SKE/Commitments/Witness Indistinguishability/PRG Length Extension

> *Not surprising, since it doesn't matter that Nature is stateful: a 1-shot reduction only uses Nature once.*

Thm 2 (*Impossibility*): Hardness amplification of arbitrary weak OWFs via direct product, using only black-box access to the OWF, is impossible.

Thm 3 (*Impossibility*): A Goldreich Levin Theorem, where the reduction has only black-box access to the OWF, is impossible.

> *Teaser: the cosmic adversary can notice when it is sent correlated inputs.*

Key Results:
Feasibilities and Impossibilities

<u>Thm 2</u> (*Impossibility of Hardness Amplification*):
Suppose there is an ε-cosmic black reduction from the OWF security of $g^n(x_1 \ldots x_n) = (g(x_1), \ldots, g(x_n))$ to the OWF security of $g(x)$ that uses only black-box access to g, and that works for any function g. Then, there exists a negligible function μ such that $\varepsilon(\lambda, a) \leq a + \mu(\lambda)$.

<u>Thm 3</u> (*Impossibility of a Goldreich Levin Theorem*):
Suppose there is an ε-cosmic black-box reduction from the security of the hardcore predicate $h(x,r) = \langle x,r \rangle$ w.r.t. $f(x, r) = (g(x), r)$ to the OWF security of g that uses only black-box access to g and that works for any function g. Then, there is a negligible function μ such that $\varepsilon(\lambda, a) \leq \mu(\lambda)$ for all a.

Key Results: Feasibilities and Impossibilities

<u>**Thm 1**</u> (*Feasibility*): classical 1-shot straight-line black-box reductions imply cosmic reductions.

<u>**Corollaries**</u>: PRFs/SKE/Commitments/Witness Indistinguishability/PRG Length Extension

> Not surprising, since it doesn't matter that Nature is stateful: a 1-shot reduction only uses Nature once.

<u>**Thm 2**</u> (*Impossibility*): Hardness amplification of arbitrary weak OWFs via direct product, using only black-box access to the OWF, is impossible.

<u>**Thm 3**</u> (*Impossibility*): A Goldreich Levin Theorem, where the reduction has only black-box access to the OWF, is impossible.

> Teaser: the cosmic adversary can notice when it is sent correlated inputs.

Claim: we need new techniques to build advanced cosmic cryptography!

Key Results:
Feasibilities and Impossibilities

Thm 1 (*Feasibility*): classical 1-shot straight-line black-box reductions imply cosmic reductions.

Corollaries: PRFs/SKE/Commitments/Witness Indistinguishability/PRG Length Extension

> Not surprising, since it doesn't matter that Nature is stateful: a 1-shot reduction only uses Nature once.

Thm 2 (*Impossibility*): Hardness amplification of arbitrary weak OWFs via direct product, using only black-box access to the OWF, is impossible.

Thm 3 (*Impossibility*): A Goldreich Levin Theorem, where the reduction has only black-box access to the OWF, is impossible.

> Teaser: the cosmic adversary can notice when it is sent correlated inputs.

Thm 4 (*Feasibility*): Hardness amplification is possible for "re-randomizable" OWFs.

Key Results: Feasibilities and Impossibilities

Thm 1 (*Feasibility*): classical 1-shot straight-line black-box reductions imply cosmic reductions.

Corollaries: PRFs/SKE/Commitments/Witness Indistinguishability/PRG Length Extension

> *Not surprising, since it doesn't matter that Nature is stateful: a 1-shot reduction only uses Nature once.*

Thm 2 (*Impossibility*): Hardness amplification of arbitrary weak OWFs via direct product, using only black-box access to the OWF, is impossible.

Thm 3 (*Impossibility*): A Goldreich Levin Theorem, where th… cess to the OWF, is impossible.

"A First Stab" at feasibility in the cosmic setting

> *Teaser: the cosmic adversary can notice when it is sent correlated i…uts.*

Thm 4 (*Feasibility*): Hardness amplification is possible for "re-randomizable" OWFs.

Plain Cosmic Security

Claim: a non-trivial setting (for feasibility),
since we rule out some classical approaches.

Plain Cosmic Security

Claim: a non-trivial setting (for feasibility),
since we rule out some classical approaches.

Later: “Relaxed” Cosmic Security

Suppose Nature isn’t fully stateful; the only state it keeps is the “time”, or the number of queries it has received.

Plain Cosmic Security

Claim: a non-trivial setting (for feasibility),
since we rule out some classical approaches.

Later: “Relaxed” Cosmic Security

Suppose Nature isn’t fully stateful; the only state it keeps is the “time”, or the number of queries it has received.

<u>Thm 5</u> (*Informal*): Non-adaptive straight-line black-box reductions give cosmic reductions for Natures that evolve over time (but are otherwise stateless).

Plain Cosmic Security

Claim: a non-trivial setting (for feasibility),
since we rule out some classical approaches.

Later: “Relaxed” Cosmic Security

Suppose Nature isn’t fully stateful; the only state it keeps is the “ti[illegible]” number of queries it has received.

Table this for now.

<u>Thm 5</u> (*Informal*): Non-adaptive straight-line black-box reductions give cosmic reductions for Natures that evolve over time (but are otherwise stateless).

Part I, Takeaways

1. At the end of the day, we want to build secure cryptosystems for "real-world attackers."

2. We currently model "real-world attackers" as PPT or QPT algorithms – and security proofs depend on the fact that we can "take a piece of code" and "re-run" it or "rewind" it. *Is this realistic?*

3. Our work: such "physical Church-Turing assumptions" are not necessary for a theory of cryptography.

Intermission

Next Up:
Walkthrough of our results

Key Results: Feasibilities and Impossibilities

Thm 1 (*Feasibility*): classical 1-shot straight-line black-box reductions imply cosmic reductions.

Corollaries: PRFs/SKE/Commitments/Witness Indistinguishability/PRG Length Extension

> Not surprising, since it doesn't matter that Nature is stateful: a 1-shot reduction only uses Nature once.

Thm 2 (*Impossibility*): Hardness amplification of arbitrary weak OWFs via direct product, using only black-box access to the OWF, is impossible.

Thm 3 (*Impossibility*): A Goldreich Levin Theorem, where the reduction has only black-box access to the OWF, is impossible.

> Teaser: the cosmic adversary can notice when it is sent correlated inputs.

Thm 4 (*Feasibility*): Hardness amplification is possible for "re-randomizable" OWFs.

Key Results: Feasibilities and Impossibilities

<u>Thm 1</u> (*Feasibility*): classical 1-shot straight-line black-box reductions imply cosmic reductions.

<u>**Corollaries**</u>: PRFs/SKE/Commitments/Witness Indistinguishability/PRG Length Extension

> Not surprising, since it doesn't matter that Nature is stateful: a 1-shot reduction only uses Nature once.

<u>Thm 2</u> (*Impossibility*): Hardness amplification of arbitrary weak OWFs via direct product, using only black-box access to the OWF, is impossible.

<u>Thm 3</u> (*Impossibility*): A Goldreich Levin Theorem, where the reduction has only black-box access to the OWF, is impossible.

> Teaser: the cosmic adversary can notice when it is sent correlated inputs.

<u>Thm 4</u> (*Feasibility*): Hardness amplification is possible for "re-randomizable" OWFs.

Basic Feasibility

<u>Thm 1</u> (*Feasibility*): classical 1-shot straight-line black-box reductions imply cosmic reductions.

1-shot straightline black-box **R**

if WIN $\mathbf{C} \longleftrightarrow \boxed{\mathbf{A}}$

then WIN $\mathbf{C'} \longleftrightarrow \mathbf{R}^{\boxed{A}}$

Basic Feasibility

Thm 1 (*Feasibility*): classical 1-shot straight-line black-box reductions imply cosmic reductions.

1-shot straightline black-box **R**

if WIN $C \longleftrightarrow \boxed{A}$

then WIN $C' \longleftrightarrow R^{\boxed{A}}$

Cosmic reduction $\forall \rho$

if WIN $C \longleftrightarrow A \longleftrightarrow \mathrm{Nat}(\rho)$

Basic Feasibility

Thm 1 (*Feasibility*): classical 1-shot straight-line black-box reductions imply cosmic reductions.

1-shot straightline black-box **R**

if WIN $C \longleftrightarrow \boxed{A}$

then WIN $C' \longleftrightarrow R^{\boxed{A}}$

Cosmic reduction $\forall \rho$

if WIN $C \longleftrightarrow A \longleftrightarrow \mathrm{Nat}(\rho)$

then WIN $C \longleftrightarrow \boxed{A \longleftrightarrow \mathrm{Nat}(\rho)}$

Basic Feasibility

Thm 1 (*Feasibility*): classical 1-shot straight-line black-box reductions imply cosmic reductions.

1-shot straightline black-box **R**

if WIN $C \longleftrightarrow \boxed{A}$

then WIN $C' \longleftrightarrow R^{\boxed{A}}$

Cosmic reduction $\forall \rho$

if WIN $C \longleftrightarrow A \longleftrightarrow \mathrm{Nat}(\rho)$

then WIN $C \longleftrightarrow \boxed{A \longleftrightarrow \mathrm{Nat}(\rho)}$

then WIN $C' \longleftrightarrow R^{\boxed{A \longleftrightarrow \mathrm{Nat}(\rho)}}$

Basic Feasibility

Thm 1 (*Feasibility*): classical 1-shot straight-line black-box reductions imply cosmic reductions.

1-shot straightline black-box **R**

if WIN $C \longleftrightarrow \boxed{A}$

then WIN $C' \longleftrightarrow R^{\boxed{A}}$

Cosmic reduction $\forall \rho$

if WIN $C \longleftrightarrow A \longleftrightarrow \mathrm{Nat}(\rho)$

then WIN $C \longleftrightarrow \boxed{A \longleftrightarrow \mathrm{Nat}(\rho)}$

then WIN $C' \longleftrightarrow R^{\boxed{A \longleftrightarrow \mathrm{Nat}(\rho)}}$

then WIN $C' \longleftrightarrow R^{A} \longleftrightarrow \mathrm{Nat}(\rho)$

Basic Feasibility

Thm 1 (*Feasibility*): classical 1-shot straight-line black-box reductions imply cosmic reductions.

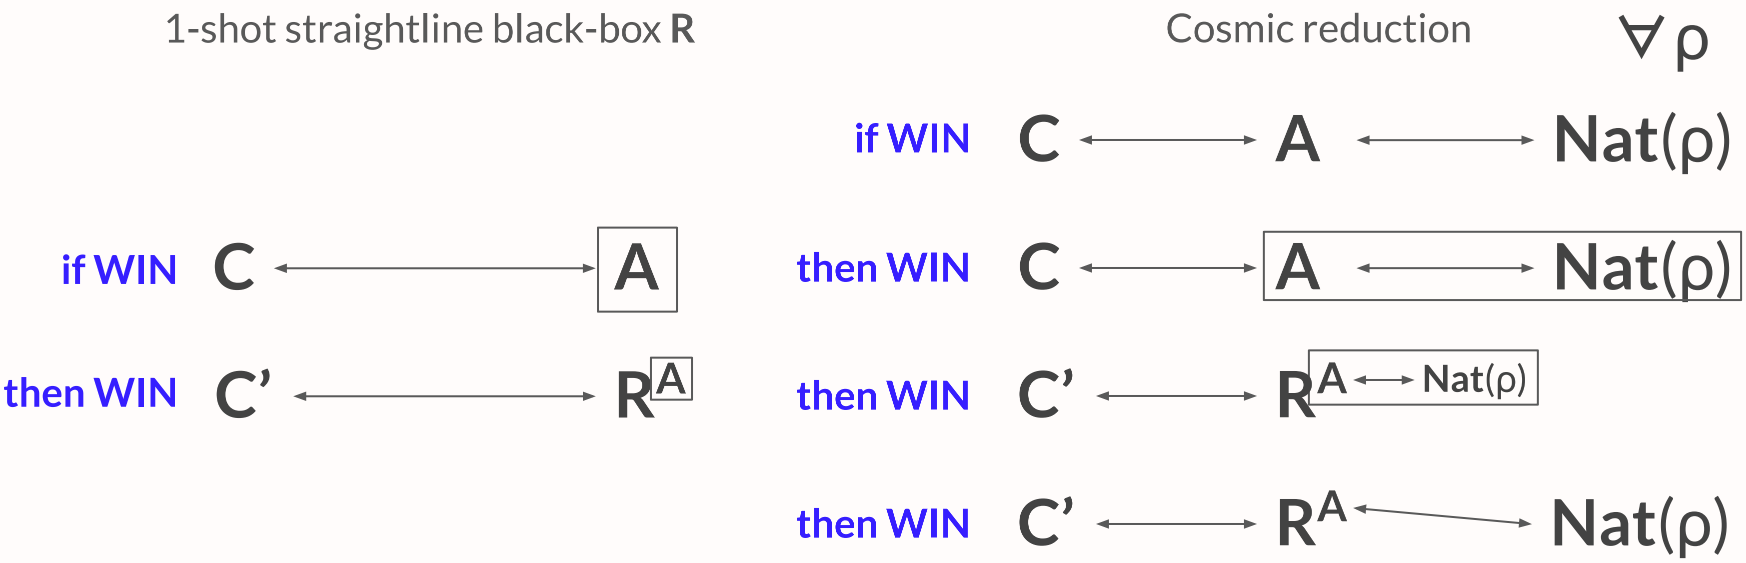

Corollaries: PRFs/SKE/Commitments/Witness Indistinguishability/PRG Length Extension

Key Results:
Feasibilities and Impossibilities

<u>**Thm 1**</u> (*Feasibility*): classical 1-shot straight-line black-box reductions imply cosmic reductions.

<u>**Corollaries**</u>: PRFs/SKE/Commitments/Witness Indistinguishability/PRG Length Extension

> Not surprising, since it doesn't matter that Nature is stateful: a 1-shot reduction only uses Nature once.

<u>Thm 2</u> (*Impossibility*): Hardness amplification of arbitrary weak OWFs via direct product, using only black-box access to the OWF, is impossible.

<u>**Thm 3**</u> (*Impossibility*): A Goldreich Levin Theorem, where the reduction has only black-box access to the OWF, is impossible.

> Teaser: the cosmic adversary can notice when it is sent correlated inputs.

<u>**Thm 4**</u> (*Feasibility*): Hardness amplification is possible for "re-randomizable" OWFs.

Key Results:
Feasibilities and Impossibilities

<u>Thm 2</u> (*Impossibility of Hardness Amplification*):
Suppose there is an ε-cosmic black reduction from the OWF security of $g^n(x_1 \ldots x_n) = (g(x_1), \ldots, g(x_n))$ to the OWF security of $g(x)$ that uses only black-box access to g, and that works for any function g.
Then, there exists a negligible function μ such that $\varepsilon(\lambda, a) \leq a + \mu(\lambda)$.

Recall: Classical Hardness Amplification

$g(x)$

Weak OWF Inverter

for $i=1\ldots n$, $y_i = g(U_\lambda)$

for random r, $y_r = g(x)$

if $g(x_r) == g(x)$,

output $x' = x_r$

else, repeat up to

$2\lambda n^2 p(\lambda)$ times

x'

Win with Pr ≥ n/2λ

$y_1, \ldots, y_n$

Strong OWF Inverter

for $g^n = (g(x_1), \ldots, g(x_n))$

$x_1, \ldots, x_n$

Win with non-negl Pr

Recall: Classical Hardness Amplification

$g(x)$

Weak OWF Inverter

for $i=1\ldots n$, $y_i = g(U_\lambda)$

for random r, $y_r = g(x)$

if $g(x_r) == g(x)$,

output $x' = x_r$

else, repeat up to

$2\lambda n^2 p(\lambda)$ *times*

x'

embed g(x) in random location
in uniformly generated image of g^n

$y_1, \ldots, y_n$

$x_1, \ldots, x_n$

repeat!

Strong OWF Inverter

for $g^n = (g(x_1), \ldots, g(x_n))$

with non-negl probability,
all n are preimages

High Level Idea: Cosmic Hardness Amplification

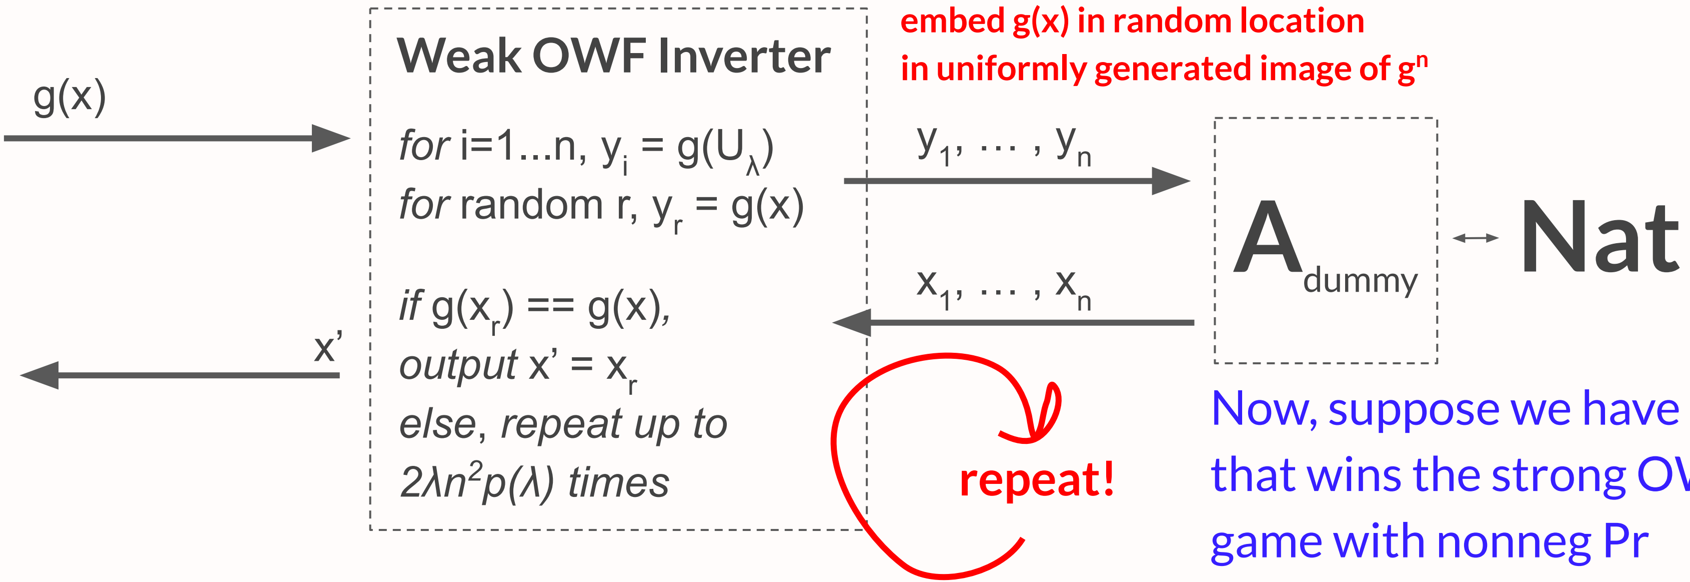

Now, suppose we have Nat that wins the strong OWF game with nonneg Pr

High Level Idea: Cosmic Hardness Amplification

$g(x)$

Weak OWF Inverter

for $i=1\ldots n$, $y_i = g(U_\lambda)$

for random r, $y_r = g(x)$

if $g(x_r) == g(x)$,

output $x' = x_r$

else, repeat up to

$2\lambda n^2 p(\lambda)$ *times*

x'

embed $g(x)$ in random location
in uniformly generated image of g^n

$y_1, \ldots, y_n$

$x_1, \ldots, x_n$

repeat!

A_{dummy} $\leftrightarrow$ **Nat**

Problem: Nat is only guaranteed to invert on uniformly generated input $y_r = g(U_\lambda)$ from a fresh instance of C

High Level Idea: Cosmic Hardness Amplification

$g(x)$

Weak OWF Inverter

for $i=1\ldots n$, $y_i = g(U_\lambda)$

for random r, $y_r = g(x)$

if $g(x_r) == g(x)$,

output $x' = x_r$

else, repeat up to

$2\lambda n^2 p(\lambda)$ times

x'

embed g(x) in random location
in uniformly generated image of g^n

$y_1, \ldots, y_n$

$x_1, \ldots, x_n$

repeat!

$A_{dummy} \leftrightarrow$ **Nat**

Since Nat is stateful,
if it saw g(x) once already,
$y_r = g(x)$ no longer looks like
part of a fresh challenge!

High Level Idea: Cosmic Hardness Amplification

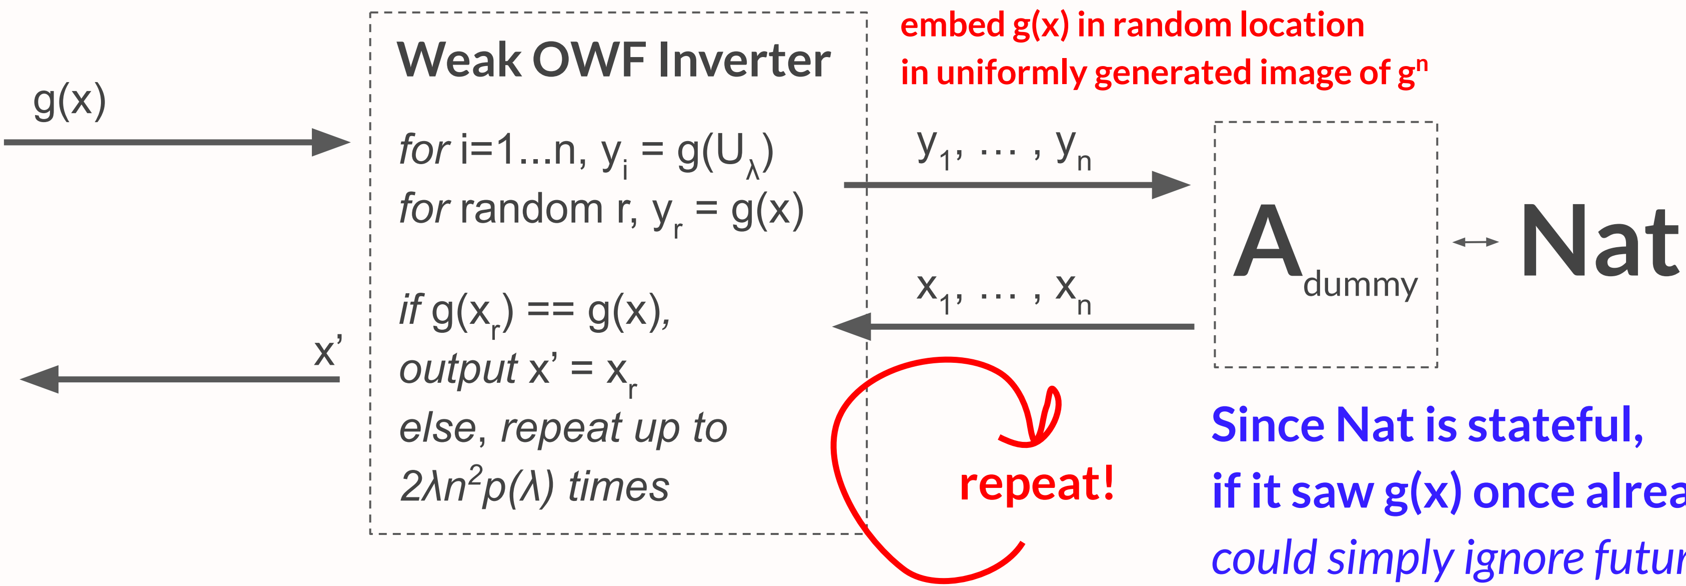

Since Nat is stateful,
if it saw g(x) once already, *it could simply ignore future queries that contain g(x), whilst still being robust winning.*

Metareduction

Suppose we did have a reduction that works with such a Nat (R uses oracle access to g)...

$g(x)$

Weak OWF Inverter

for $i=1\ldots n$, $y_i = g(U_\lambda)$
for random r, $y_r = g(x)$

if $g(x_r) == g(x)$,
output $x' = x_r$
else, repeat up to
$2\lambda n^2 p(\lambda)$ *times*

x'

embed $g(x)$ in random location in uniformly generated image of g^n

$y_1, \ldots, y_n$

$x_1, \ldots, x_n$

$A_{dummy} \leftrightarrow$ **Nat**

repeat!

Since Nat is stateful, if it saw g(x) once already, *it could simply ignore future queries that contain g(x), whilst still being robust winning.*

Claim: can simulate Nat "easily" s.t. the reduction loses at most a(.) advantage overall.

Metareduction

Suppose we did have a reduction that works with such a Nat (R uses oracle access to g)…

$g(x)$

Weak OWF Inverter

for $i=1\ldots n$, $y_i = g(U_\lambda)$
for random r, $y_r = g(x)$

if $g(x_r) == g(x)$,
output $x' = x_r$
else, repeat up to $2\lambda n^2 p(\lambda)$ *times*

x'

embed g(x) in random location in uniformly generated image of g^n

$y_1, \ldots, y_n$

$x_1, \ldots, x_n$

repeat!

$\mathbf{Sim^g}$

- Suppose g is a random length-tripling function
- Then, images of g are "sparse". Reduction won't find other images except by invoking the oracle, which Sim controls
- **Thus, Sim only needs to guess "one response", namely that for the challenge g(x).**
- Sim simply returns ⊥ when asked to invert g(x), whereas Nat returns $g^{-1}(g(x))$ with Pr a(.). Thus, Pr[Diverge] <= a(.)

Claim: can simulate Nat "easily" s.t. the reduction loses at most a(.) advantage overall.

Thus, Black-Box Cosmic Hardness Amplification is *Impossible*.

Classical reductions that use an attacker repeatedly on correlated inputs may fail, if the attacker notices the correlation and halts.

We may need to hide the correlation.

Key Results:
Feasibilities and Impossibilities

<u>Thm 1</u> (*Feasibility*): classical 1-shot straight-line black-box reductions imply cosmic reductions.

<u>Corollaries</u>: PRFs/SKE/Commitments/Witness Indistinguishability/PRG Length Extension

> *Not surprising, since it doesn't matter that Nature is stateful: a 1-shot reduction only uses Nature once.*

<u>Thm 2</u> (*Impossibility*): Hardness amplification of arbitrary weak OWFs via direct product, using only black-box access to the OWF, is impossible.

<u>Thm 3</u> (*Impossibility*): A Goldreich Levin Theorem, where the reduction has only black-box access to the OWF, is impossible. Same High Level Idea!

> *Teaser: the cosmic adversary can notice when it is sent correlated inputs.*

<u>Thm 4</u> (*Feasibility*): Hardness amplification is possible for "re-randomizable" OWFs.

Key Results:
Feasibilities and Impossibilities

<u>Thm 3</u> (*Impossibility of a Goldreich Levin Theorem*):
Suppose there is an ε-cosmic black-box reduction from the security of the hardcore predicate $h(x,r) = \langle x,r \rangle$ w.r.t. $f(x, r) = (g(x), r)$ to the OWF security of g that uses only black-box access to g and that works for any function g. Then, there is a negligible function μ such that $\varepsilon(\lambda, a) \le \mu(\lambda)$ for all a.

Recall: Goldreich Levin Theorem

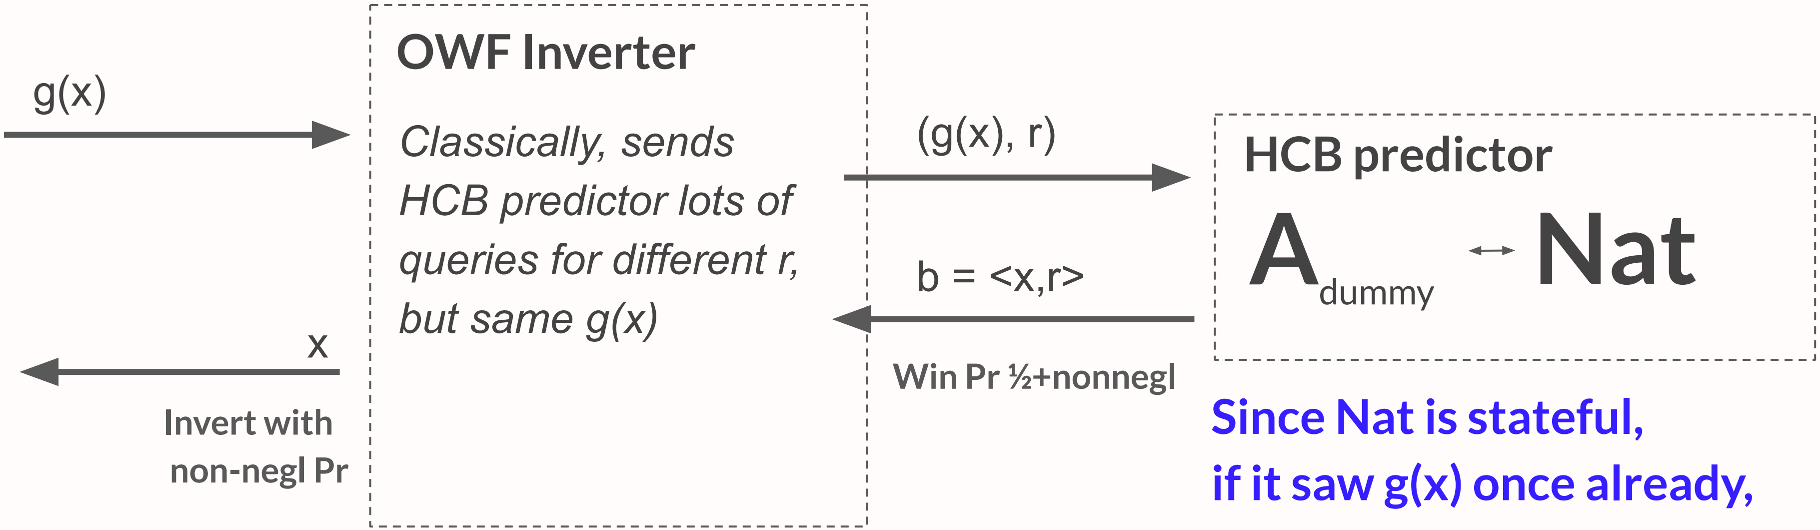

Since Nat is stateful,
if it saw g(x) once already,
it could *just ignore future queries that contain* g(x).

Recall: Goldreich Levin Theorem

g(x)

OWF Inverter

Classically, sends HCB predictor lots of queries for different r, but same g(x)

x

Invert with non-negl Pr

(g(x), r)

b = <x,r>

Win Pr ½+nonnegl

HCB predictor

$A_{\text{dummy}} \leftrightarrow \text{Nat}$

Since Nat is stateful,
if it saw g(x) once already,
it could *just ignore future queries that contain* g(x).

Claim: can simulate Nat "easily" s.t. the reduction loses at most 1/2 advantage overall. (How?)

Thus, a Goldreich Levin Theorem is *Impossible*.

Takeaway: classical techniques fail.

To get around it, need techniques that are non-black-box in the OWF.

Key Results:
Feasibilities and Impossibilities

<u>Thm 1</u> (*Feasibility*): classical 1-shot straight-line black-box reductions imply cosmic reductions.

<u>Corollaries</u>: PRFs/SKE/Commitments/Witness Indistinguishability/PRG Length Extension

> Not surprising, since it doesn't matter that Nature is stateful: a 1-shot reduction only uses Nature once.

<u>Thm 2</u> (*Impossibility*): Hardness amplification of arbitrary weak OWFs via direct product, using only black-box access to the OWF, is impossible.

<u>Thm 3</u> (*Impossibility*): A Goldreich Levin Theorem, where th … cess to the OWF, is impossible.

“A First Stab” at feasibility in the cosmic setting

> Teaser: the cosmic adversary can notice when it is sent correlated inputs.

<u>Thm 4</u> (*Feasibility*): Hardness amplification is possible for “re-randomizable” OWFs.

Feasibility: Hardness Amplification

Suppose g(x) is <u>re-randomizable</u>.

<u>**Def:**</u> A one-way function g is **re-randomizable** if $\exists$ PPT *rand*(.), *recover*(.) s.t.

$$\forall x, \{r \leftarrow \{0,1\}^\lambda : rand(g(x), r)\} \equiv \{x' \leftarrow \{0,1\}^\lambda : g(x')\}$$

$$\forall x, r, \text{ let } y \leftarrow rand(g(x), r),\ x' \leftarrow recover(g^{-1}(y), r) \text{, then } g(x) = g(x').$$

Denote **g(x)** = *rand*(g(x),U_λ) and **x'** = *recover*(x', r), where r is the previous seed used.

Feasibility: Hardness Amplification

Suppose $g(x)$ is re-randomizable.

$g(x)$

Weak OWF Inverter

for $i=1\ldots n$, $y_i = g(U_\lambda)$
for random r, $y_r = g(x)$

if $g(x_r) == g(x)$,
output $x' = x_r$
else, *repeat up to*
$2\lambda n^2 p(\lambda)$ *times*

x'

embed g(x) in random location
in uniformly generated image of g^n

$y_1, \ldots, y_n$

$x_1, \ldots, x_n$

$A_{\text{dummy}} \leftrightarrow \text{Nat}(\rho)$

repeat!

Feasibility: Hardness Amplification

Suppose g(x) is <u>re-randomizable</u>.

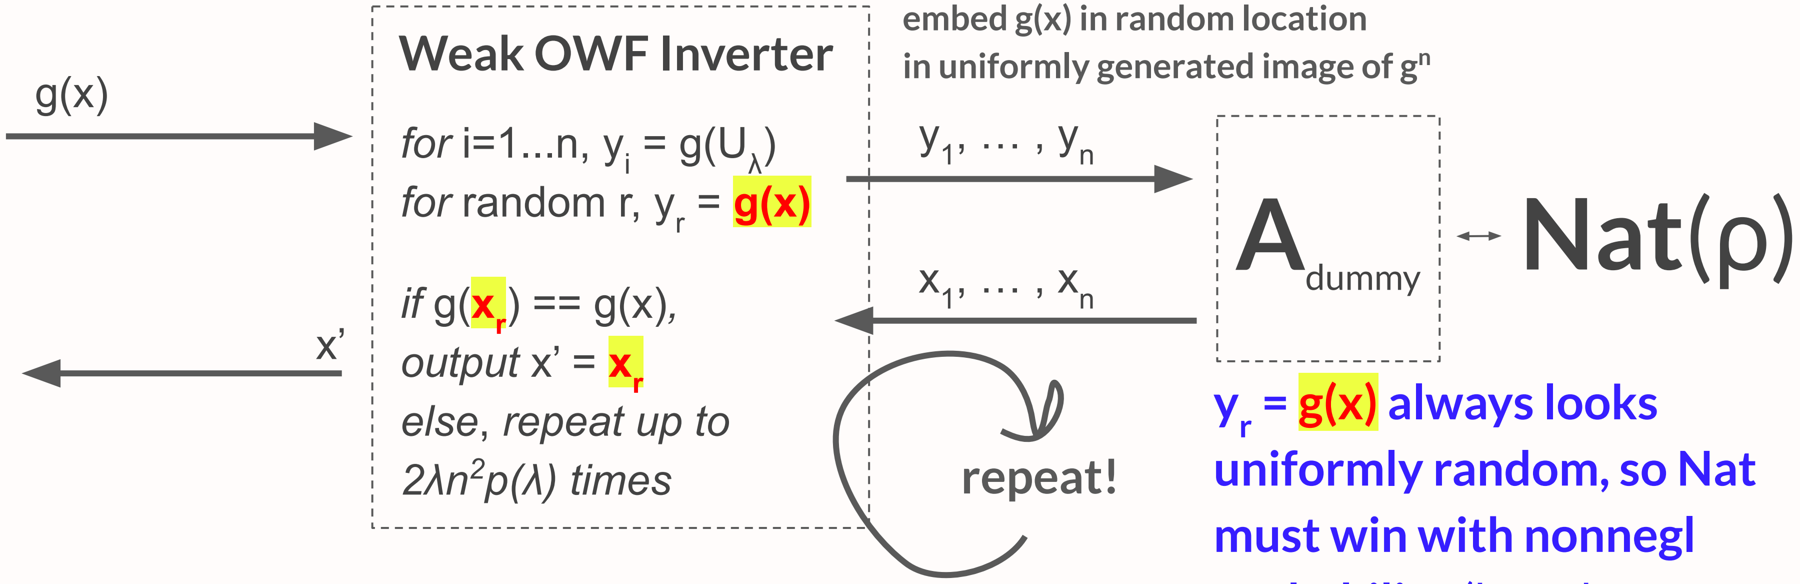

y_r = g(x) always looks uniformly random, so Nat must win with nonnegl probability (by robust winning).

We Can Go Beyond Single-Shot Straight-Line Reductions!

Key Point: We need to hide any correlation between queries in the view of the adversary. For example, by re-randomizing **g(x)**.

How Far Can We Go?

Open Problem:

Even though a Goldreich-Levin Theorem (that is black-box in the OWF) is impossible, can we build cosmic PRGs from OWFs?

For now…

…let’s climb a different mountain.

Roadmap

1. ~~Motivation~~ *(10min)*
2. ~~Defining Cosmic Security~~ *(15min)*
3. ~~Properties of Cosmic Security: a Sanity Check~~
    a. ~~Composition, Black-box reductions~~ *(5min)*
4. ~~Summary of Key Results~~
    a. ~~Feasibilities and Impossibilities (20min)~~
5. **Other Notions of Cosmic Security (10min)**
6. Conclusion (5min)

Recall: Why did this fail?

Because Nat can adjust future behavior based on prior game outcomes.

$g(x)$

Weak OWF Inverter

for $i=1\ldots n$, $y_i = g(U_\lambda)$
for random r, $y_r = g(x)$

if $g(x_r) == g(x)$,
output $x' = x_r$
else, *repeat up to* $2\lambda n^2 p(\lambda)$ *times*

x'

embed $g(x)$ in random location
in uniformly generated image of g^n

$y_1, \ldots, y_n$

$x_1, \ldots, x_n$

repeat!

$A_{dummy} \leftrightarrow \mathrm{Nat}(\rho)$

Since Nat is stateful,
if it saw $g(x)$ once already, it
could simply ***ignore future queries that contain*** $g(x)$.

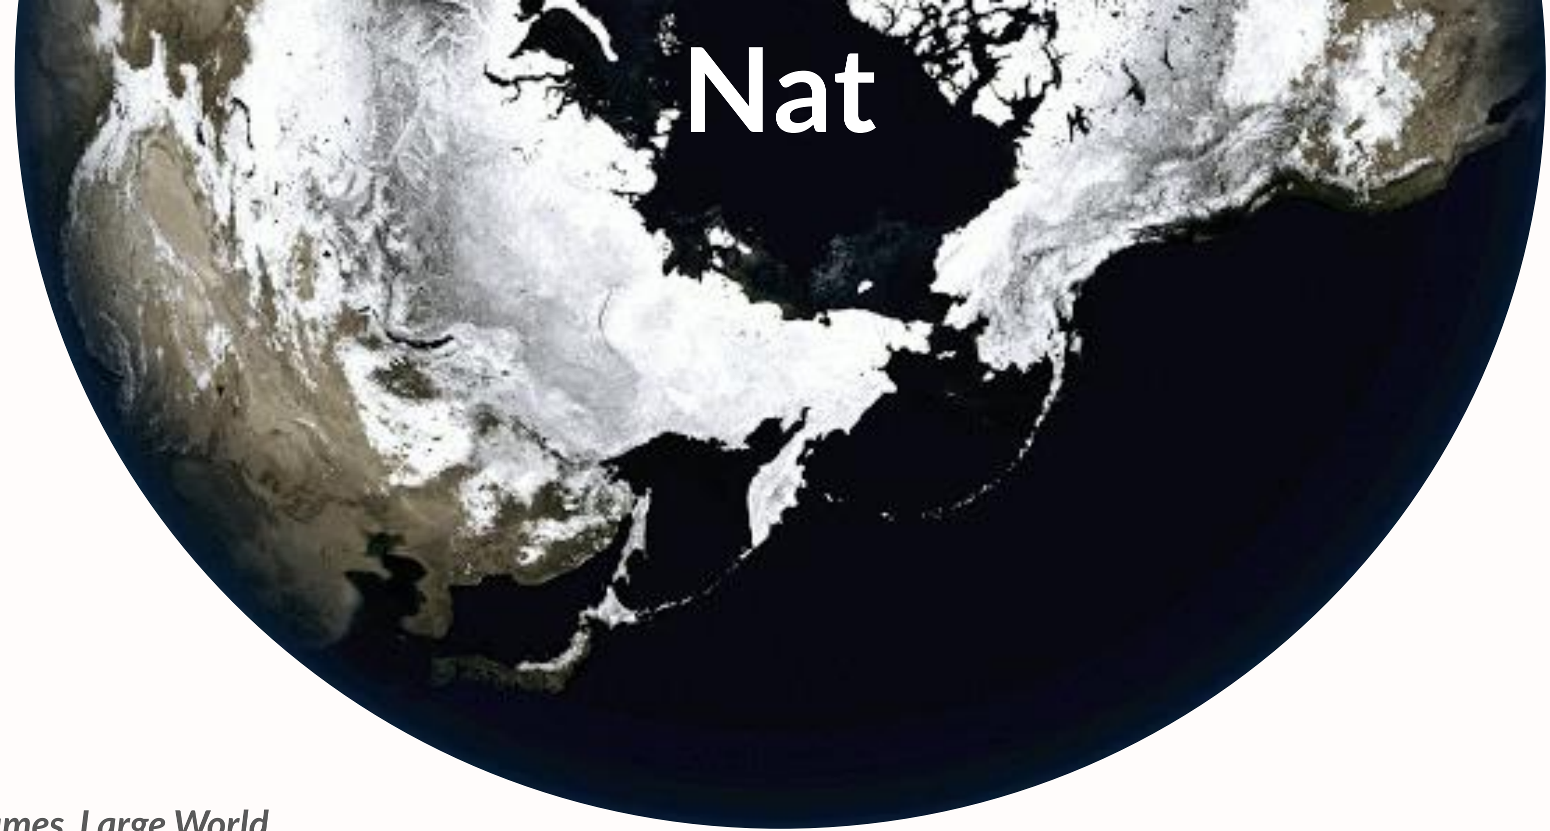

__Small Games, Large World__

It may be presumptuous to think that **C** or **A** can *influence* the future behavior of **Nat**.

$$\updownarrow$$
A
$$\updownarrow$$
C

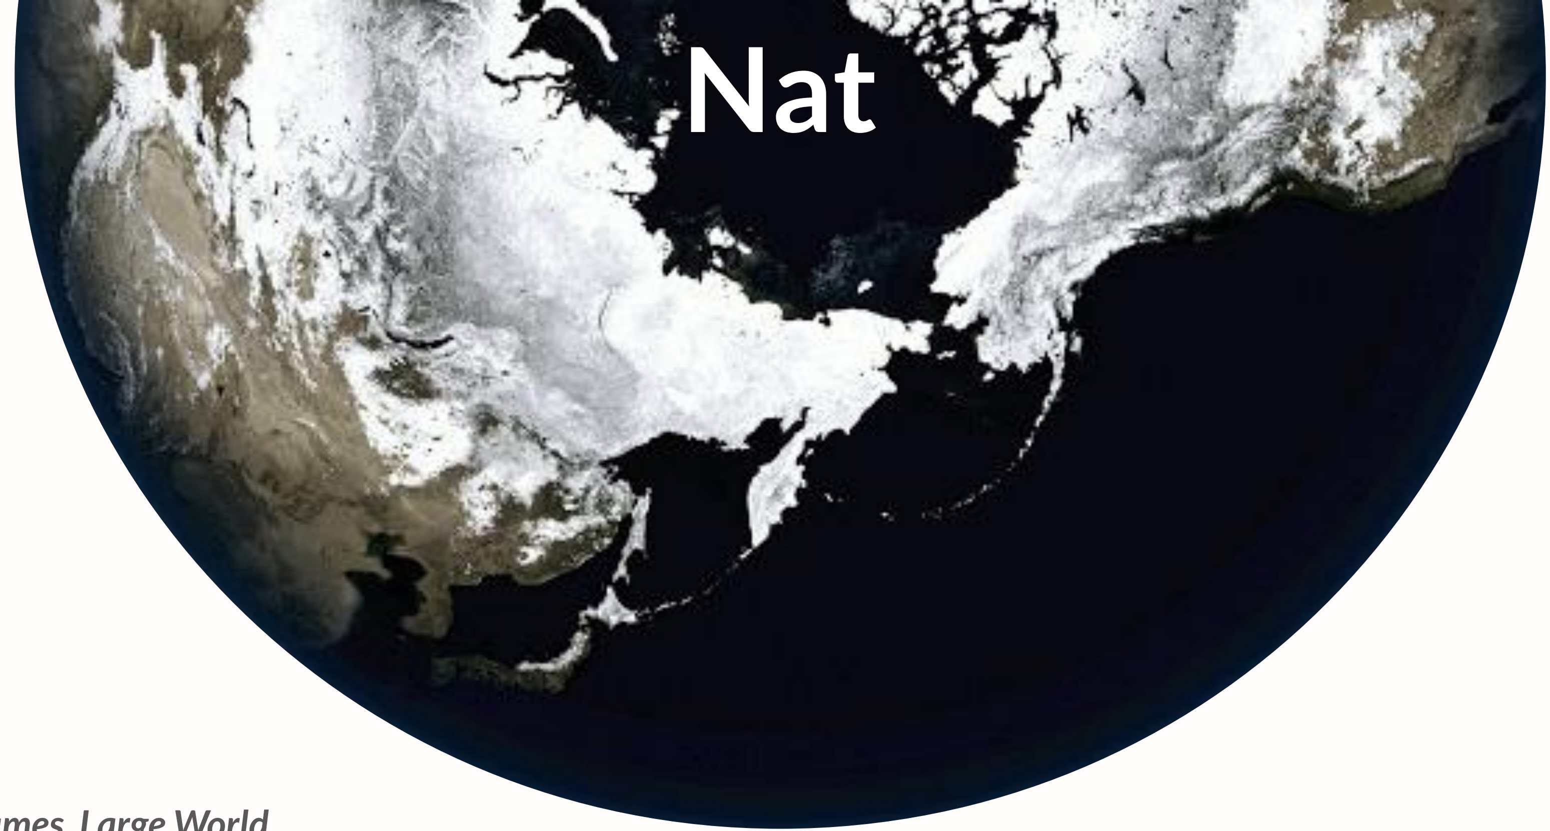

Small Games, Large World

It may be presumptuous to think that **C** or **A** can *influence* the future behavior of **Nat**.

A

C

What if Nat plays every game independently, the same way? “Restartable”, and *classic*.

But maybe **Nat** evolves over time (# of queries it has received)...

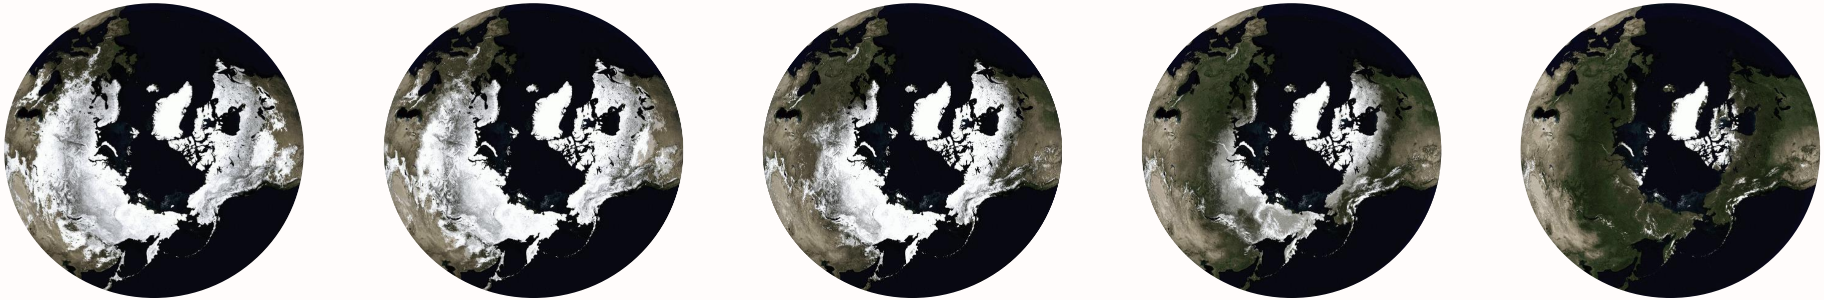

But maybe **Nat** evolves over time (# of queries it has received)...

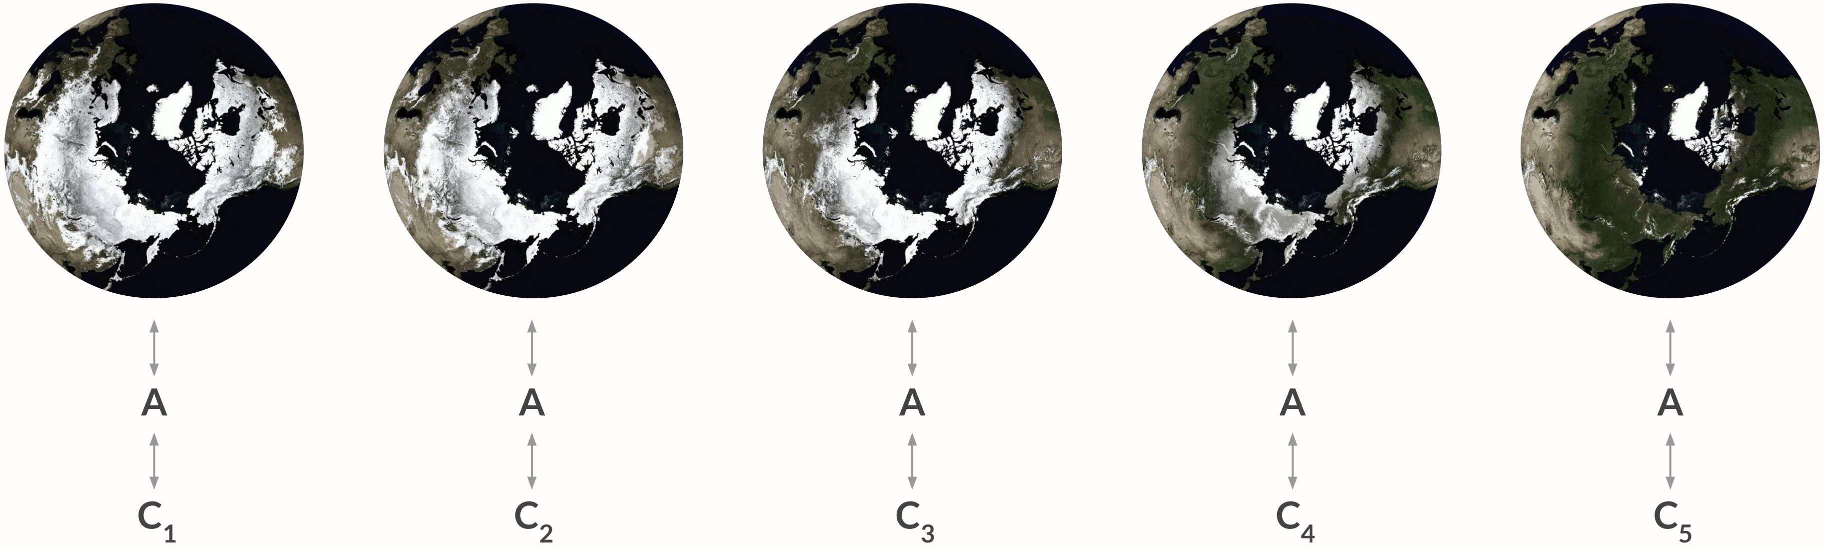

...the same way regardless of any interaction we have with it.

Let's formalize it.

Def: (**A**, **Nat**) is μ-weakly restartable if ∃ **Sim** s.t. ∀λ, ∀C , ∀ interaction prefixes ρ,

$$C \leftrightarrow A \leftrightarrow \mathrm{Nat}(\rho) \overset{\mu(\lambda)}{\approx} C \leftrightarrow \mathrm{Sim}(\overbrace{|\rho|}^{\text{Number of messages in } \rho})$$

statistical distance

"Simulator"
any distribution

***In other words, the behavior of* (A, Nat) *in the view of any C is pre-programmed and depends only on* |ρ|.**

Regular Cosmic Adversaries

Weakly-Restartable Cosmic Adversaries

(**A**, **Nat**) is now a sequence of attackers **Sim**(0), **Sim**(1), **Sim**(2), …

$$
\begin{tikzcd}
\mathbf{C'} \arrow[r, leftrightarrow] & \mathbf{R} \arrow[r, leftrightarrow] \arrow[dr, leftrightarrow] \arrow[ddr, leftrightarrow] \arrow[dddr, leftrightarrow] \arrow[ddddr, leftrightarrow] & \mathbf{Sim}(0) \\
& & \mathbf{Sim}(r) \\
& & \mathbf{Sim}(2r) \\
& & \mathbf{Sim}(3r) \\
& & \mathbf{Sim}(4r)
\end{tikzcd}
$$

Weakly-Restartable Cosmic Adversaries

(**A**, **Nat**) is now a sequence of attackers $\mathbf{A_1}$ $\mathbf{A_2}$ $\mathbf{A_3}$ ...

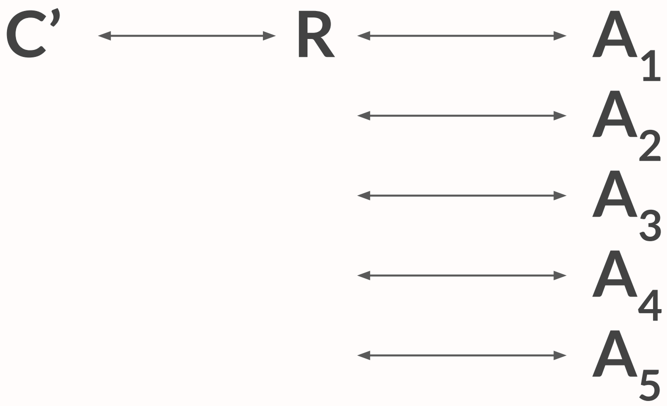

Weakly-Restartable Cosmic Adversaries

(**A**, **Nat**) is now a sequence of attackers **A_1 A_2 A_3 ...** (informal)

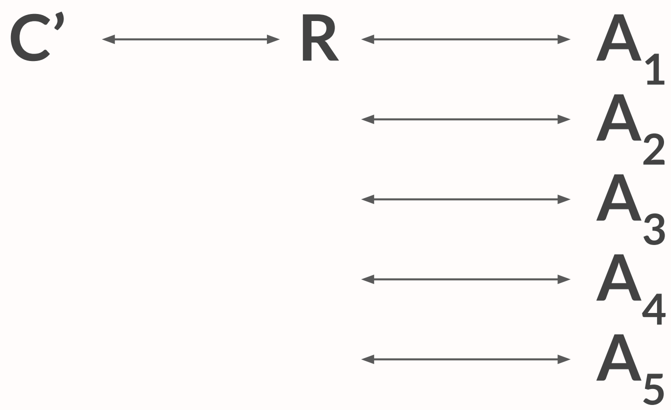

Theorem: Suppose there is a **non-adaptive** (straight-line black-box) reduction $\mathbf{R}_{\text{classic}}$ from **C** to **C'**.
Then there is a cosmic reduction from **C** to **C'**, assuming (**A**, **Nat**) is weakly restartable.

Corollaries: Hardcore Bits from OWFs, Hardness Amplification.

<u>Theorem:</u> Suppose there is a **non-adaptive** (straight-line black-box) reduction $\mathbf{R}_{\text{classic}}$ from **C** to **C'**. Then there is a cosmic reduction from **C** to **C'**, assuming (**A**, **Nat**) is weakly restartable.

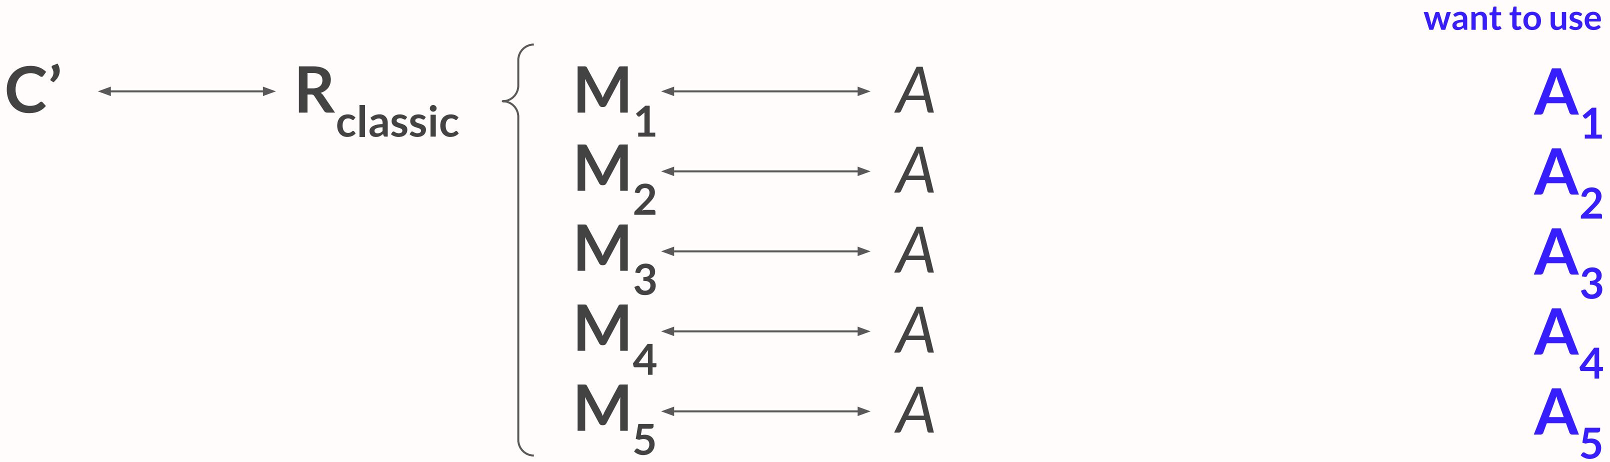

Theorem: Suppose there is a **non-adaptive** (straight-line black-box) reduction $\mathbf{R}_{\text{classic}}$ from **C** to **C'**.
Then there is a cosmic reduction from **C** to **C'**, assuming (**A**, **Nat**) is weakly restartable.

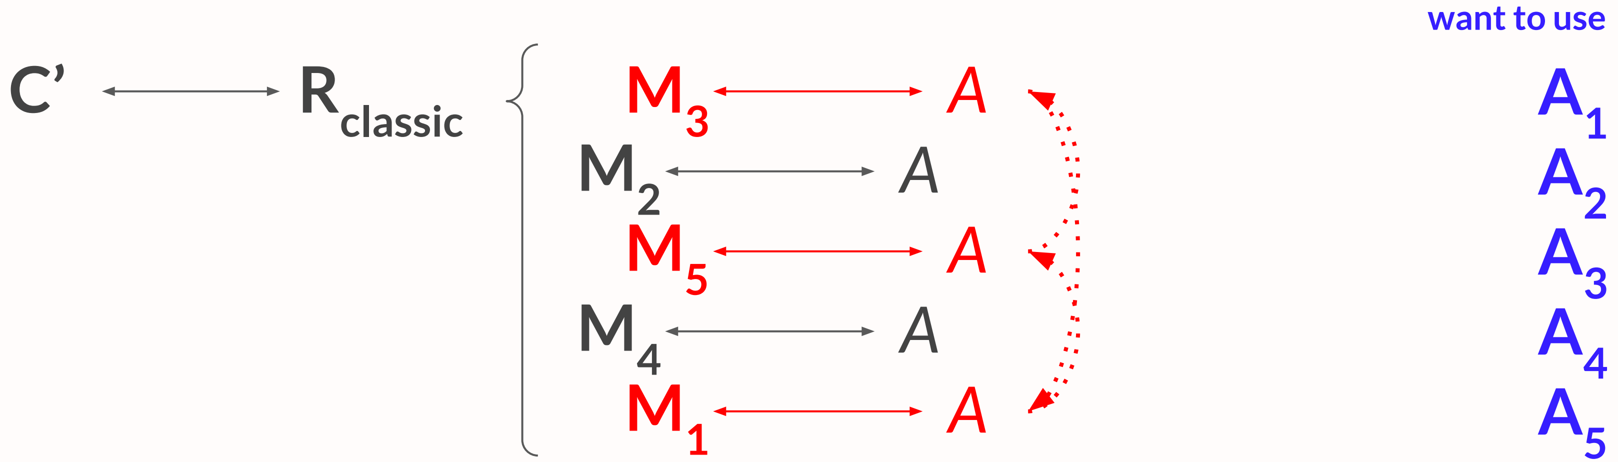

nonadaptive: can permute the order in any way

Theorem: Suppose there is a non-adaptive (straight-line black-box) reduction $\mathbf{R}_{\text{classic}}$ from **C** to **C'**.
Then there is a cosmic reduction from **C** to **C'**, assuming (**A**, **Nat**) is weakly restartable.

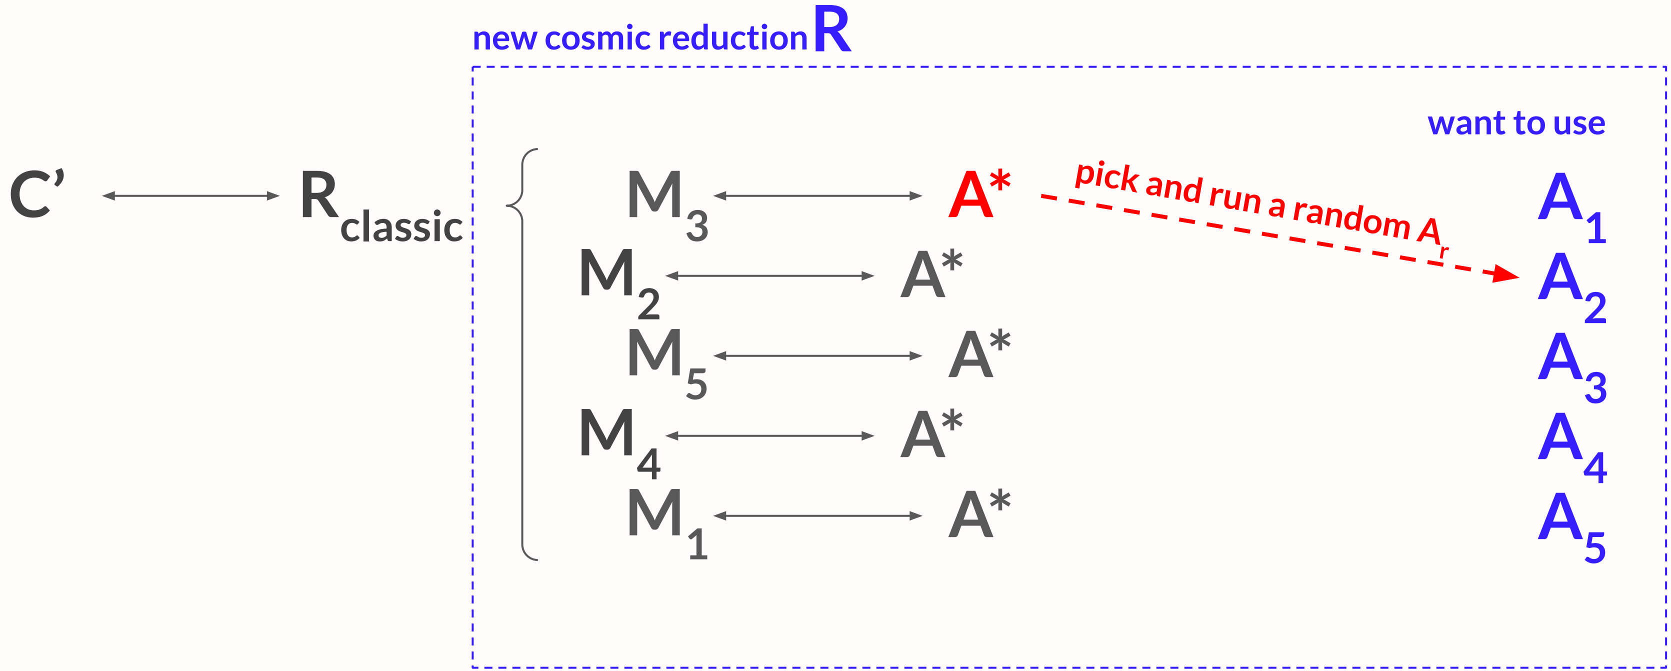

Theorem: Suppose there is a non-adaptive (straight-line black-box) reduction $\mathbf{R}_{\text{classic}}$ from **C** to **C'**.
Then there is a cosmic reduction from **C** to **C'**, assuming (**A**, **Nat**) is weakly restartable.

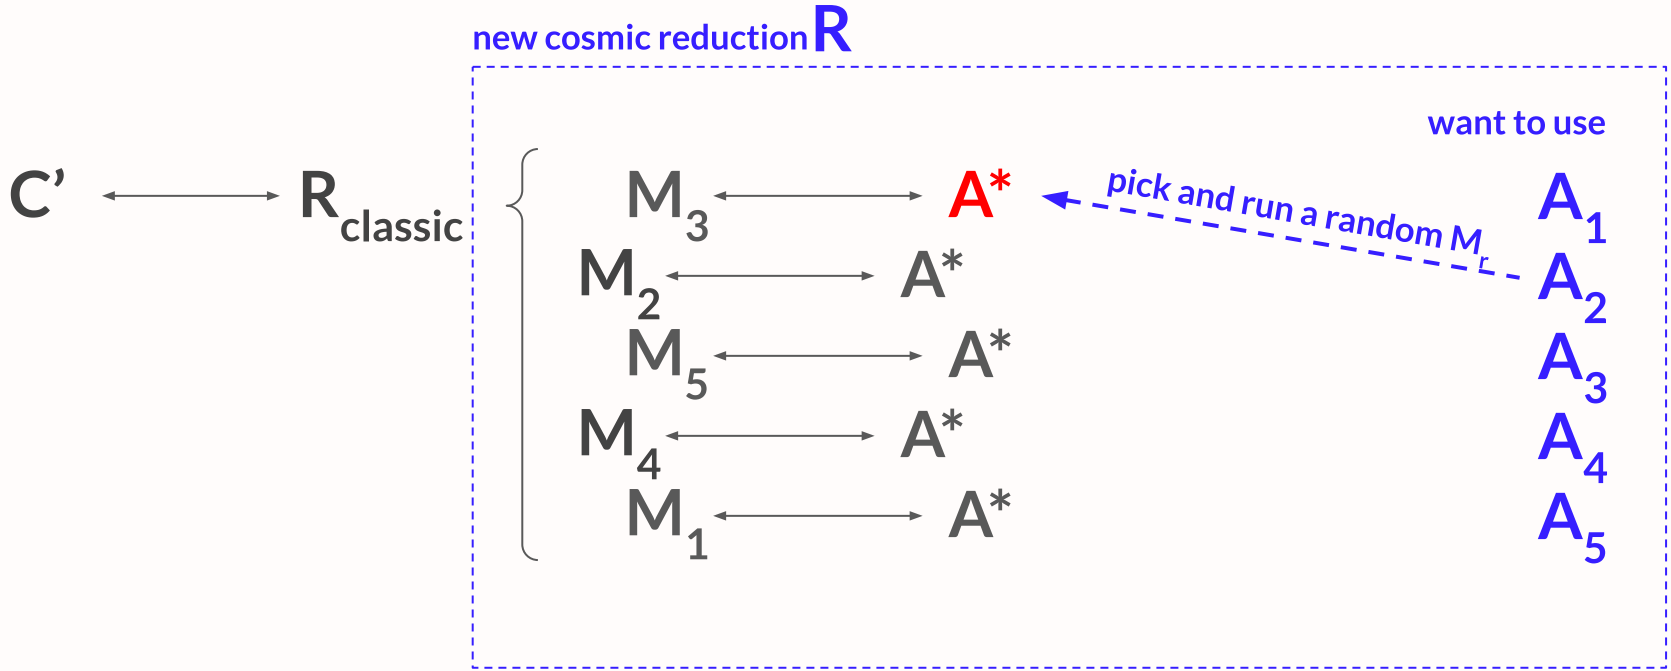

Theorem: Suppose there is a non-adaptive (straight-line black-box) reduction $\mathbf{R}_{\text{classic}}$ from **C** to **C'**. Then there is a cosmic reduction from **C** to **C'**, assuming (**A**, **Nat**) is weakly restartable.

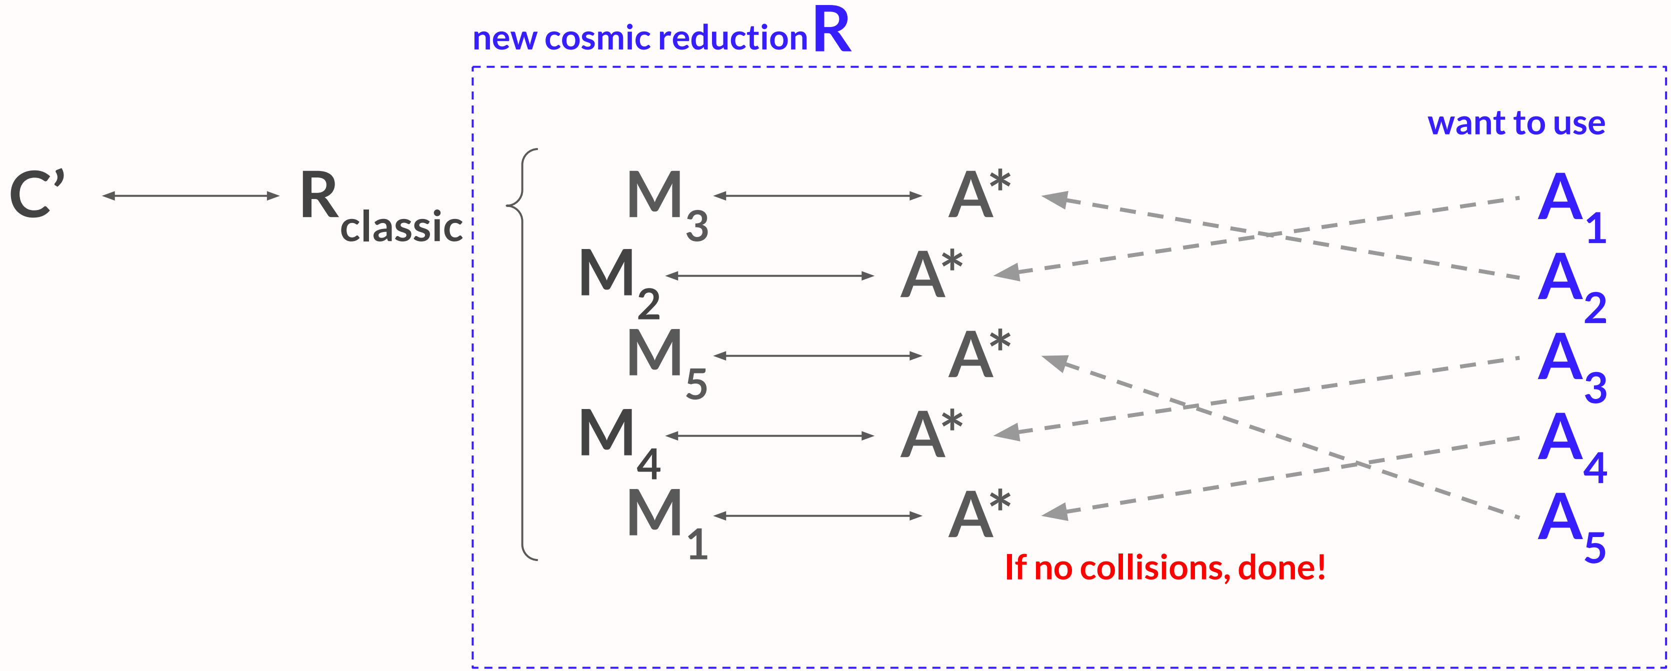

Theorem: Suppose there is a non-adaptive (straight-line black-box) reduction $\mathbf{R}_{\text{classic}}$ from **C** to **C'**. Then there is a cosmic reduction from **C** to **C'**, assuming (**A**, **Nat**) is weakly restartable.

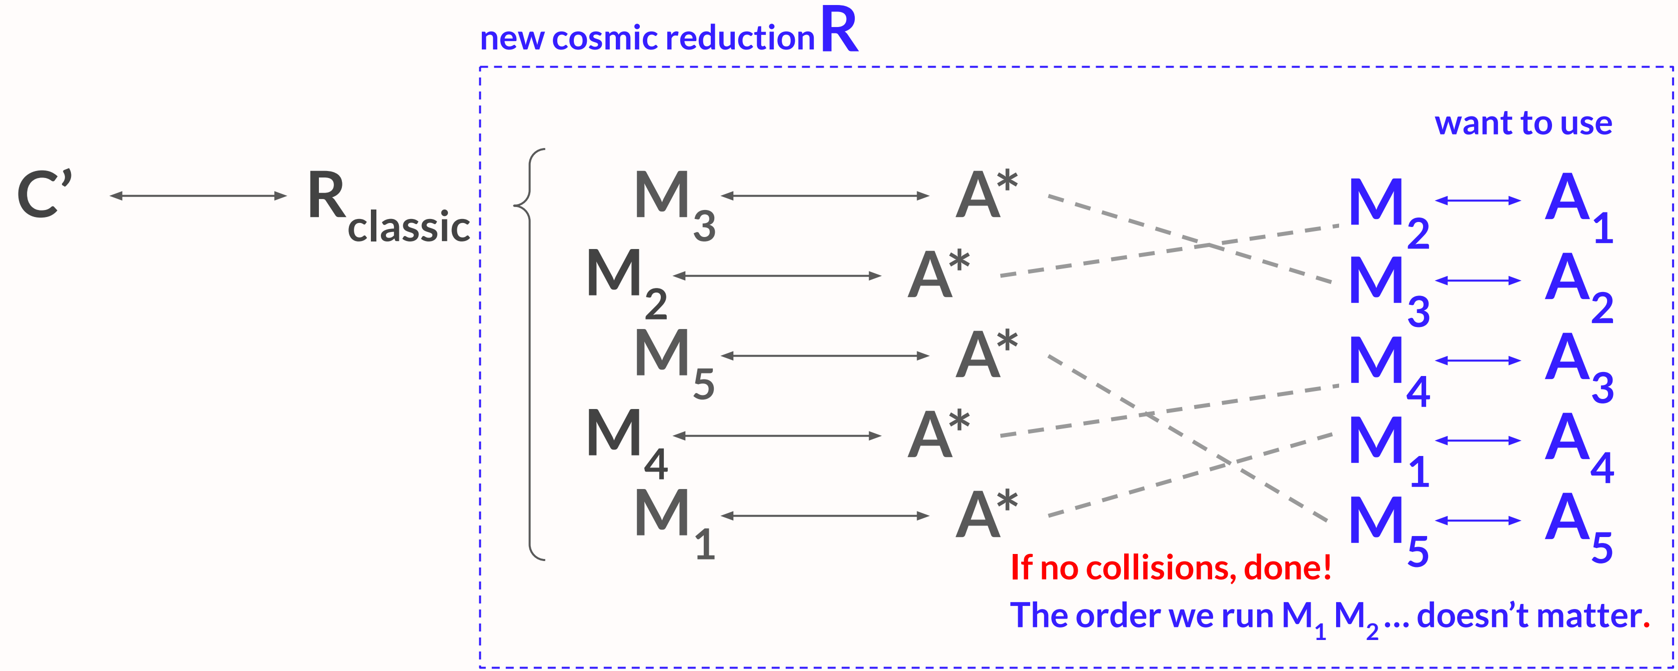

Theorem: Suppose there is a non-adaptive (straight-line black-box) reduction $\mathbf{R}_{\mathbf{classic}}$ from **C** to **C'**. Then there is a cosmic reduction from **C** to **C'**, assuming (**A**, **Nat**) is weakly restartable.

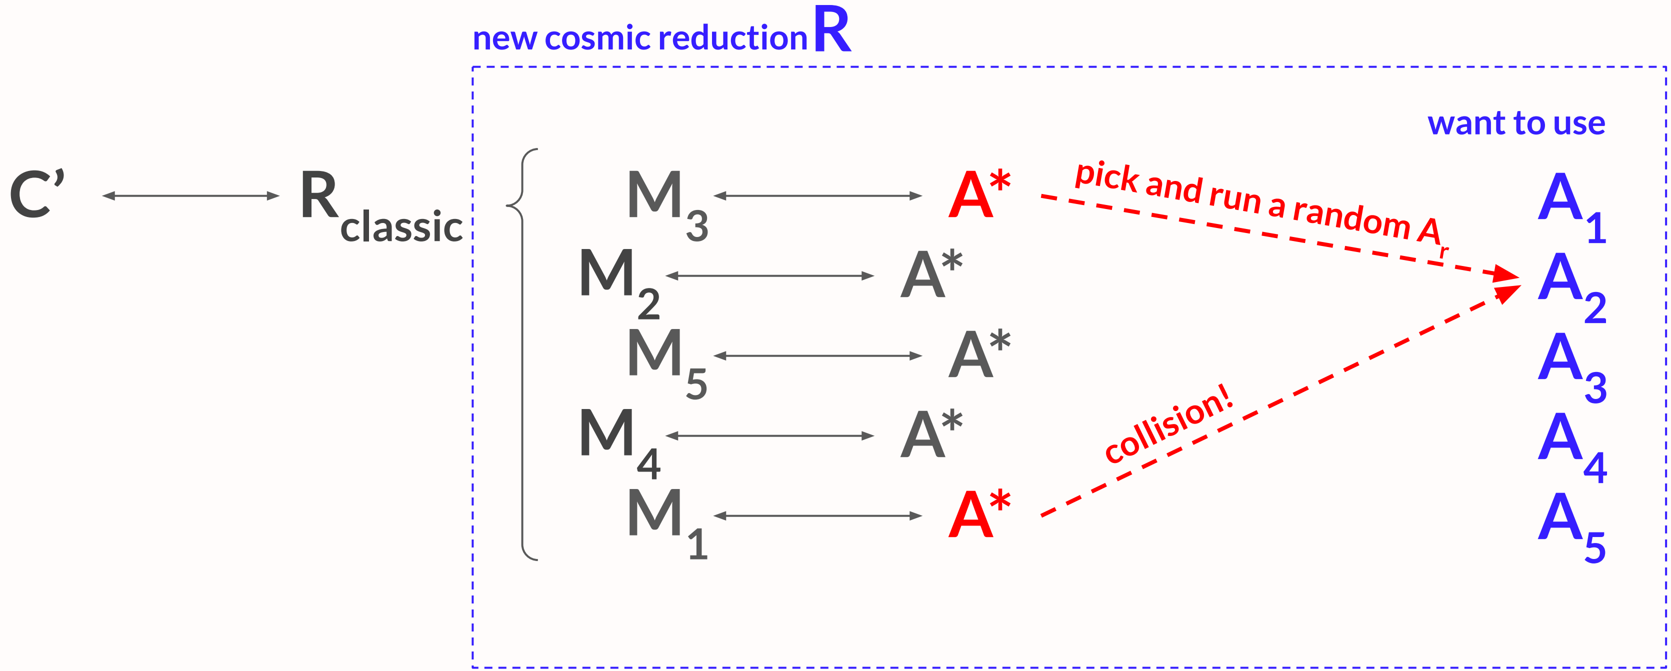

<u>Theorem:</u> Suppose there is a non-adaptive (straight-line black-box) reduction $\mathbf{R}_{\text{classic}}$ from **C** to **C'**. Then there is a cosmic reduction from **C** to **C'**, assuming (**A**, **Nat**) is weakly restartable.

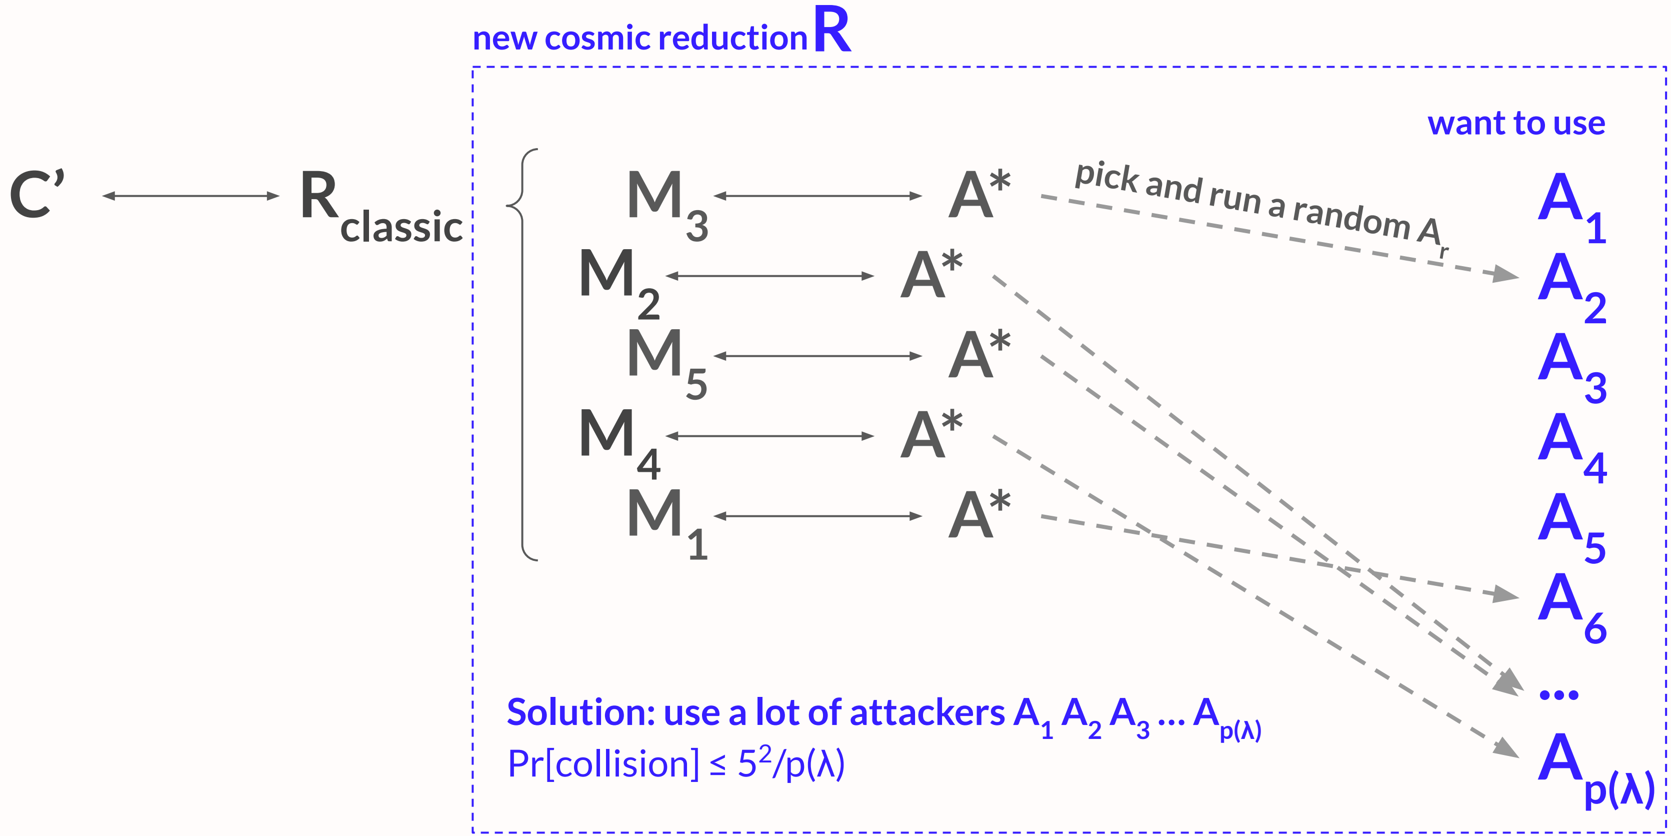

Theorem: Suppose there is a non-adaptive (straight-line black-box) reduction $\mathbf{R}_{\text{classic}}$ from **C** to **C'**. Then there is a cosmic reduction from **C** to **C'**, assuming (**A**, **Nat**) is weakly restartable.

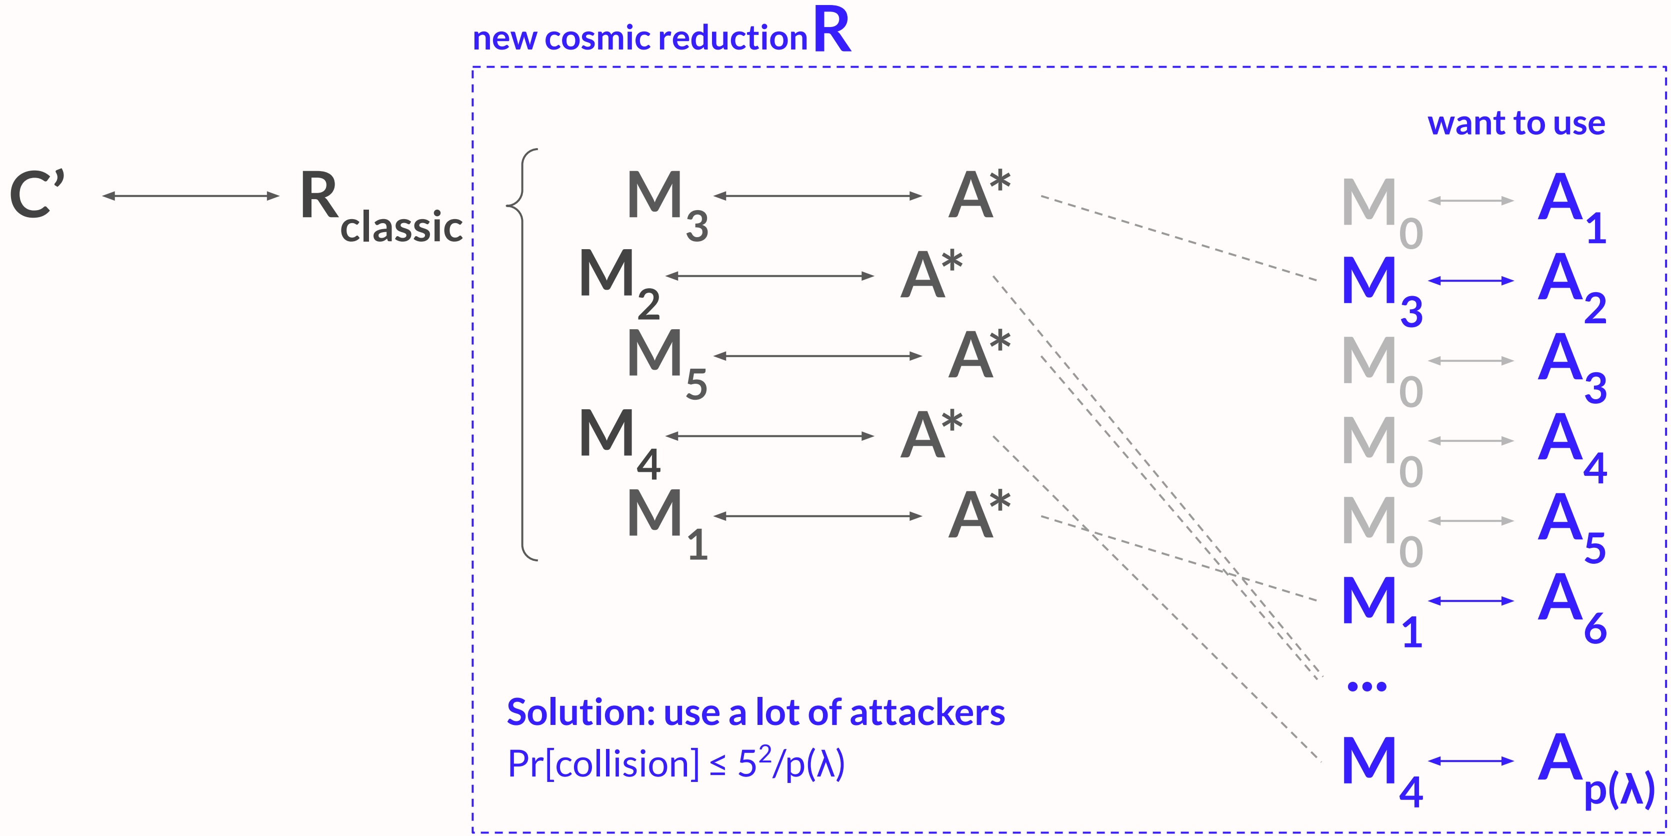

Theorem: Suppose there is a non-adaptive (straight-line black-box) reduction $\mathbf{R}_{classic}$ from **C** to **C'**. Then there is a cosmic reduction from **C** to **C'**, assuming (**A**, **Nat**) is weakly restartable.

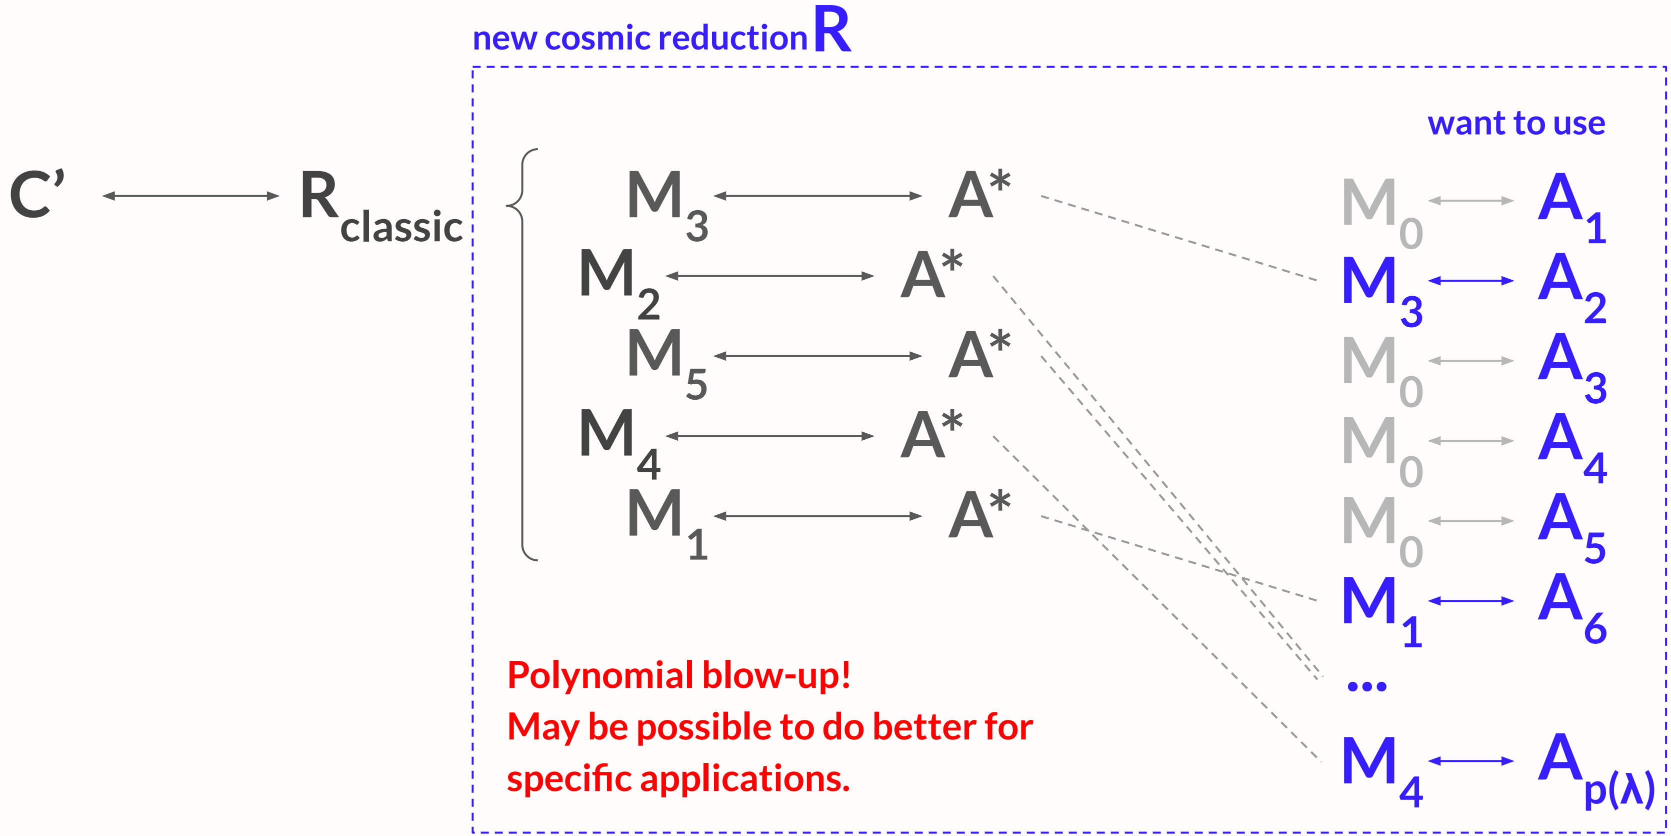

Theorem: Suppose there is a non-adaptive (straight-line black-box) reduction $\mathbf{R}_{\text{classic}}$ from **C** to **C'**. Then there is a cosmic reduction from **C** to **C'**, assuming (**A**, **Nat**) is weakly restartable.

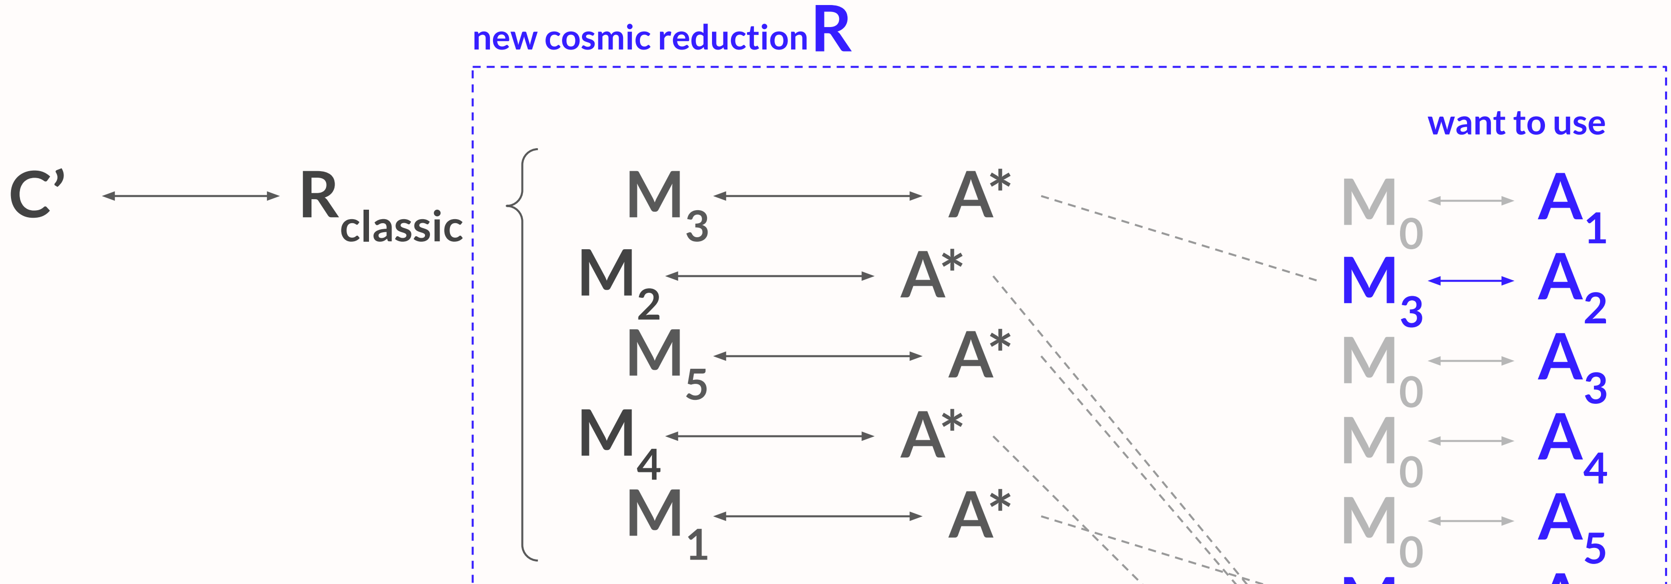

Polynomial blow-up!
May be possible to do better for specific applications.

Roadmap

1. ~~Motivation~~ *(10min)*
2. ~~Defining Cosmic Security~~ *(15min)*
3. ~~Properties of Cosmic Security: a Sanity Check~~
    a. ~~Composition, Black-box reductions~~ *(5min)*
4. ~~Summary of Key Results~~
    ~~a. Feasibilities and Impossibilities (20min)~~
5. ~~Other Notions of Cosmic Security (10min)~~
6. **Conclusion (5min)**

Let’s wrap it up.

In Sum

It is perhaps surprising that classical non-adaptive reductions “work” for the cosmic weakly-restartable “A_1 A_2 A_3” model.

Implication: we can treat slightly stateful Natures (that keep only time as state) as stateless!
(if we don’t make adaptive queries)

In Greater Sum

For fully stateful Natures,

- Several black-box classical reductions that run the adversary repeatedly on correlated queries cannot have cosmic equivalents.
- So far, feasibility for cosmic reductions is limited to either re-randomizing correlated queries, or sticking to one-shot reductions.

In Greater, Greater Sum

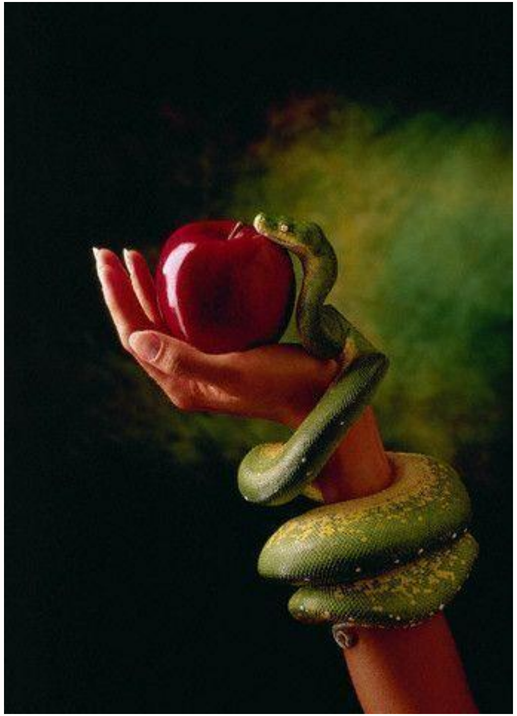

Biggest Takeaway: we should consider stateful adversaries that may behave differently each time its run, in order to minimize "unprovable" assumptions on the attacker.

What’s Next?
An Unexplored Universe.

- **PRGs from OWFs?**
- **MPC?**
- **New techniques to deal with a stateful Cosmos?**

What's Next?
An Unexplored Universe.

- PRGs from OWFs?
- MPC?
- New techniques to deal with a stateful Cosmos?

Thank You!

Image Credits: NASA/ESA/Hubble


Extra Slides