Security proofs for device-independent randomness expansion

Carl A. Miller

University of Michigan, Ann Arbor

Ref: "Universal security for randomness expansion from the spot-checking protocol" (arXiv:1411.6608), with Yaoyun Shi.

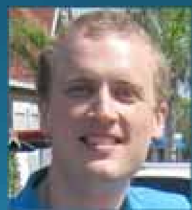

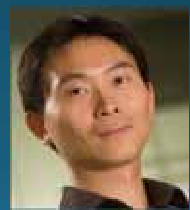

The question

Can we generate provable random numbers?

? → 10110111101101000010010001111101001001001001111010100 ....

Why it matters

Security of protocols like RSA breaks down if randomness is bad.
[Lenstra+ 12, Heninger+ 12]

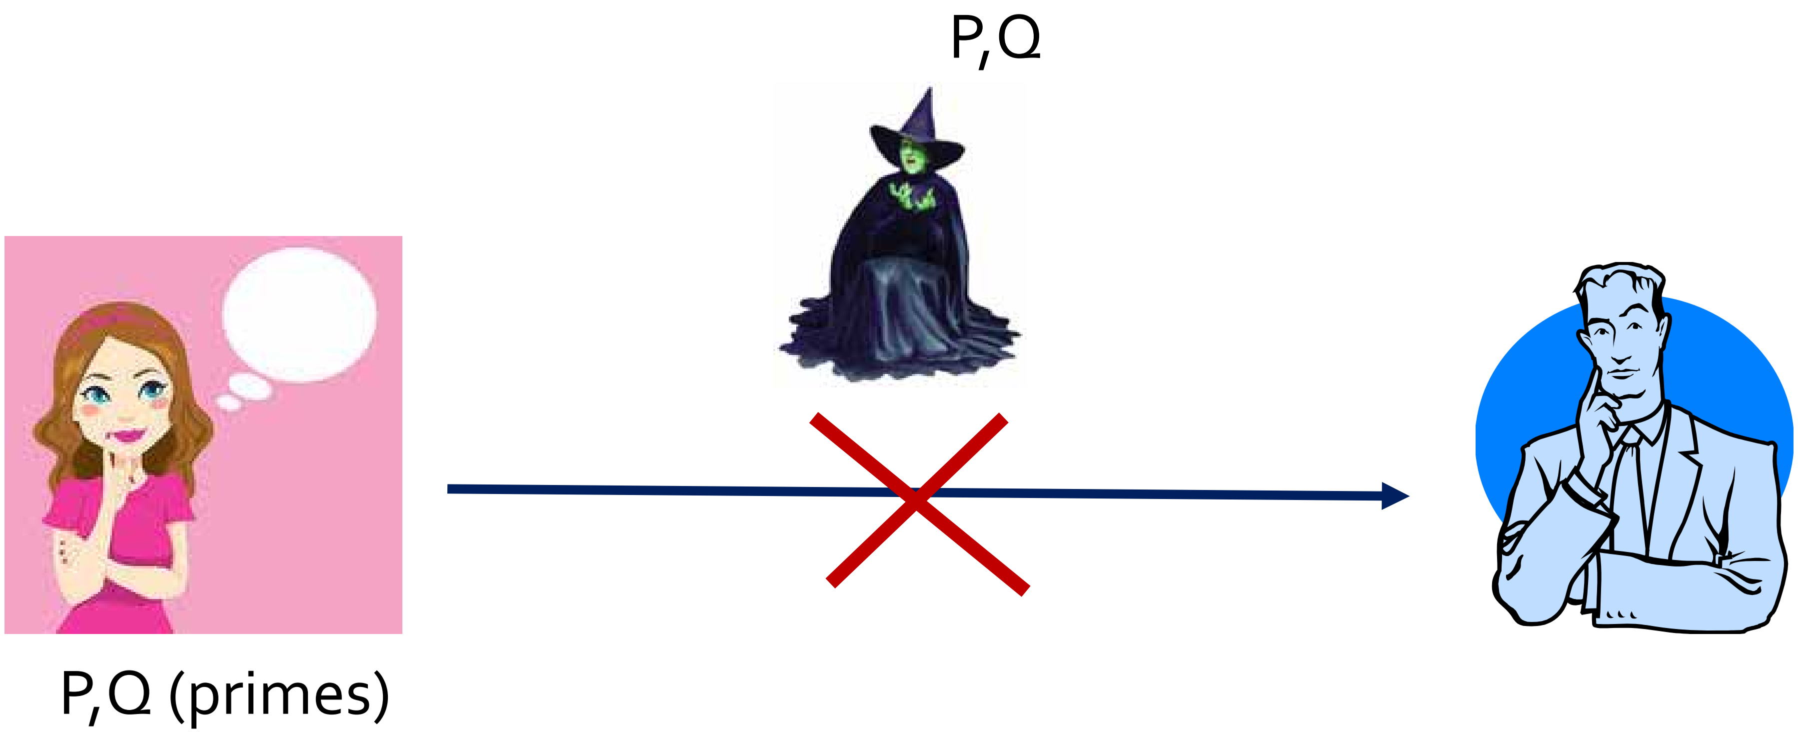

Existing solutions

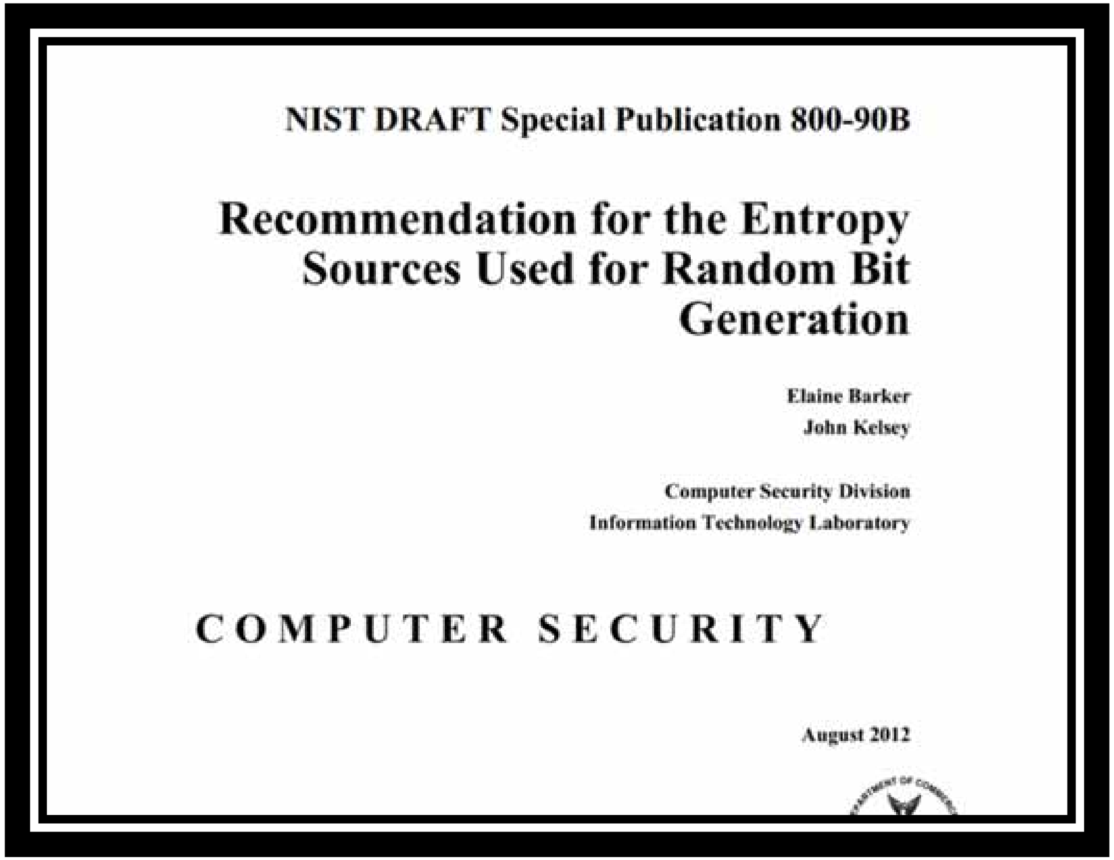

NIST DRAFT Special Publication 800-90B

Recommendation for the Entropy Sources Used for Random Bit Generation

Elaine Barker
John Kelsey

Computer Security Division
Information Technology Laboratory

COMPUTER SECURITY

August 2012

"[We assume] that the developer understands the behavior of the entropy source and has made a **good-faith effort** to produce a consistent source of entropy."

Can we generate randomness without assuming good faith?

Quantum random number generation

- Untrusted-device randomness expansion
- Untrusted-device randomness amplification
- Semi-device-independent random number generation.
- Contextuality-based randomness expansion.
- Randomness extraction.

Quantum random number generation

- Untrusted-device randomness expansion

Small uniform seed + untrusted device -> uniform randomness

0 0 1 1 1 0 1 1

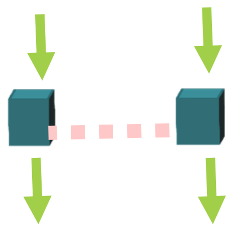

101011110110001001101100
1111011001101111011111111
10100001010001001111110
10101010111010101010 ....

Only assumption: Non-communication.

History

Timeline

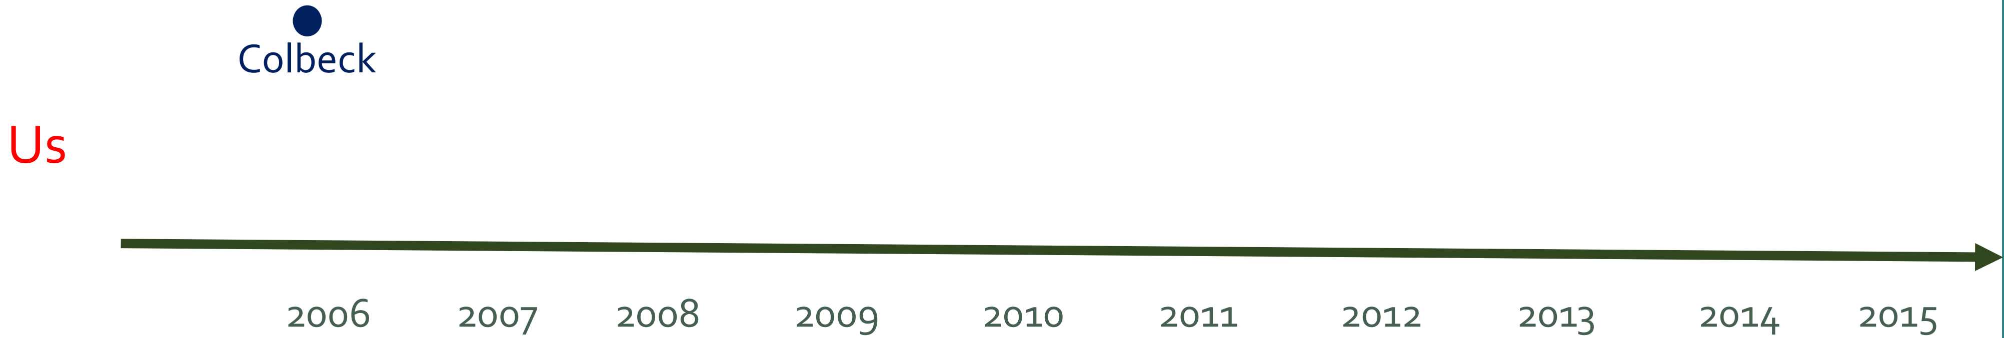

Colbeck proposed a protocol based on **repeating a nonlocal game.**

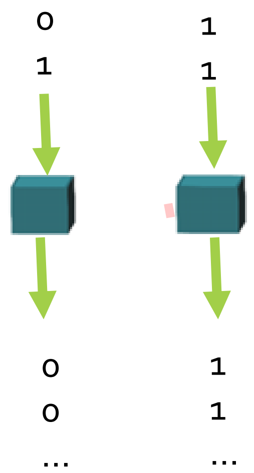

Timeline

Colbeck

Pironio+

Us

started

2006 2007 2008 2009 2010 2011 2012 2013 2014 2015

Pironio+: analysis & experiment.

We started working on the problem.

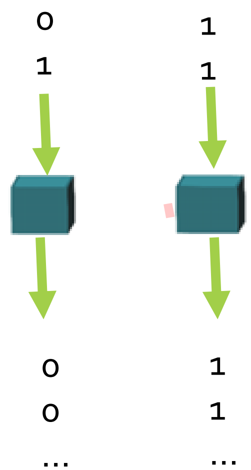

The challenge of the entangled adversary

Quantum information can be **locked** – accessible *only* to entangled adversaries. [E.g., DiVincenzo+ 04]

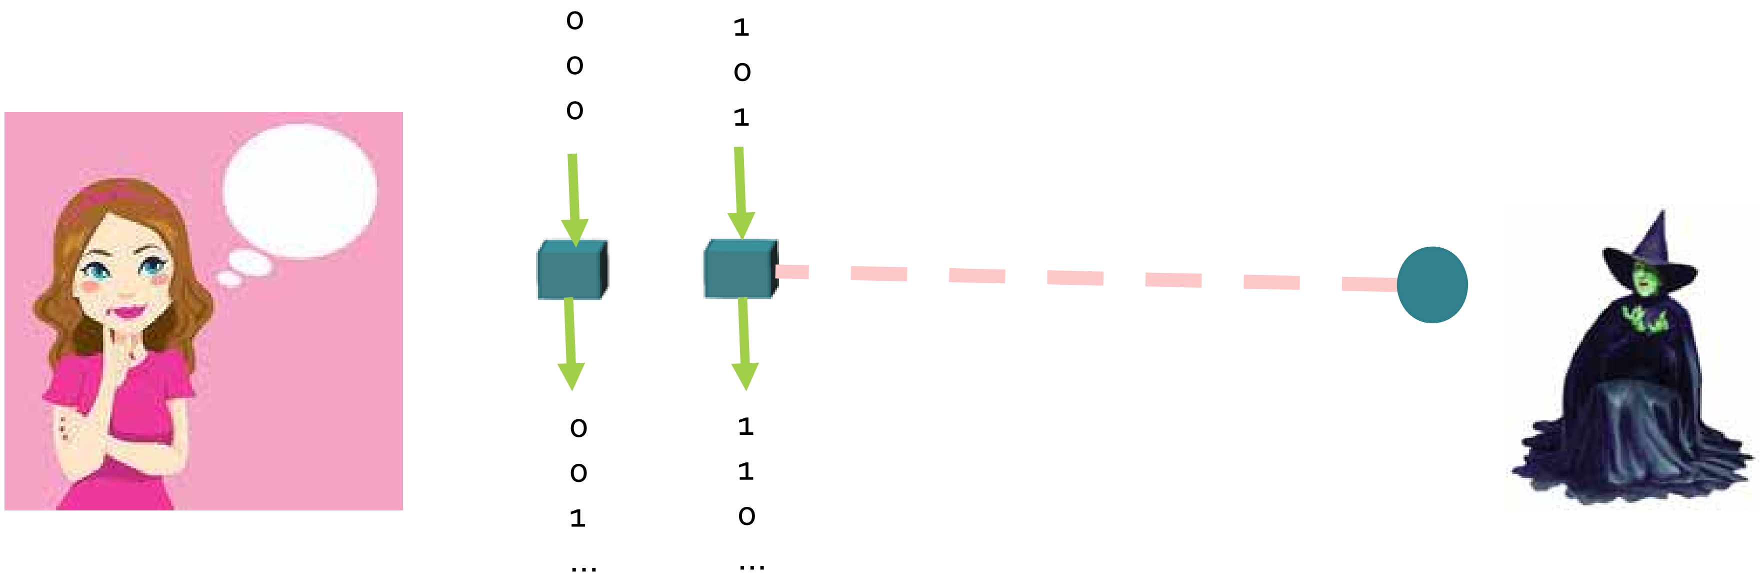

Timeline

Colbeck

Pironio+

Us

started

2006 2007 2008 2009 2010 2011 2012 2013 2014 2015

Pironio+: analysis & experiment.

We started working on the problem.

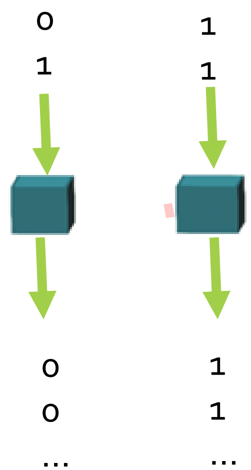

Timeline

Colbeck
Pironio+
Pironio+ Fehr+
Coudron+

Us
started

2006 2007 2008 2009 2010 2011 2012 2013 2014 2015

Pironio+, Fehr+, Coudron+ proved security against unentangled adversary.

Methods: Classical statistical arguments (e.g., Azuma's inequality).

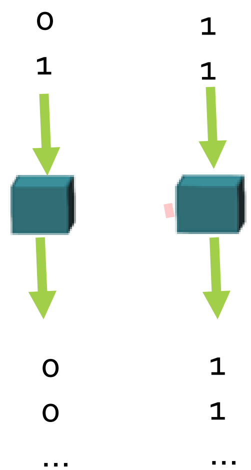

Timeline

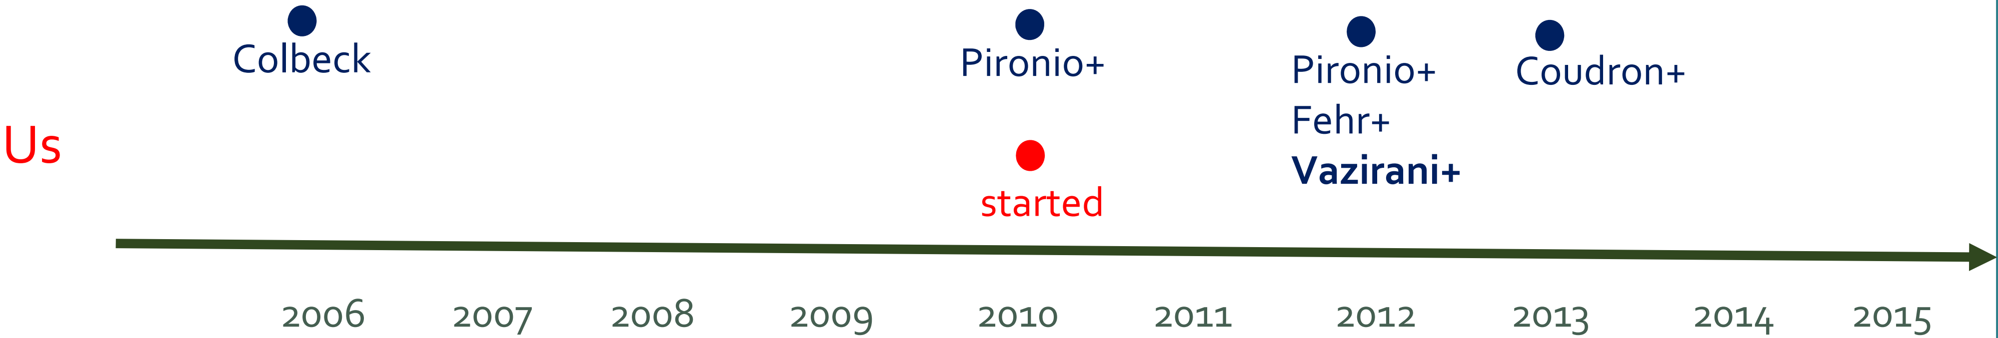

Vazirani-Vidick proved full security against an entangled adversary (no error tolerance).

Method: Insecurity implies communication paradox.

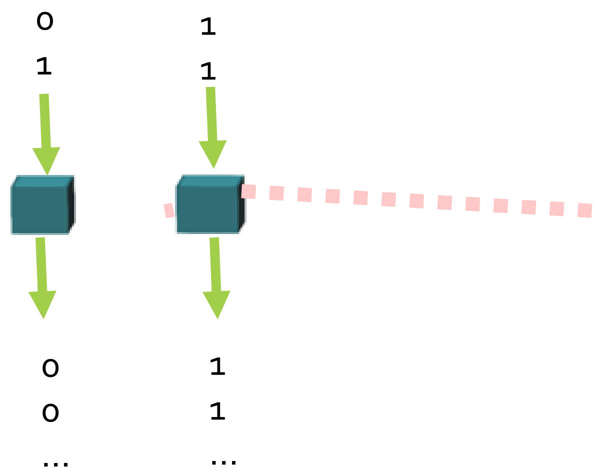

Timeline

Colbeck

Pironio+

Pironio+
Fehr+
Vazirani+

Coudron+

Us

started

Miller+

2006 2007 2008 2009 2010 2011 2012 2013 2014 2015

We proved full security,
with error-tolerance >1.4%.

Method: Constraints on
$X \mapsto Tr[X^{1+\varepsilon}]$
for adversary-output states.

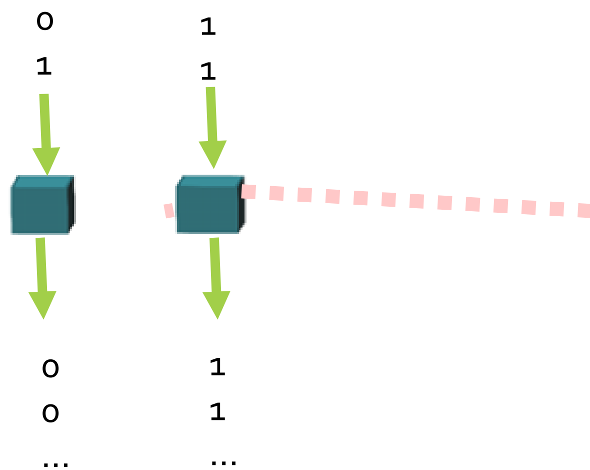

Timeline

Colbeck

Us

on+

2006 2007 2008 20

2015

Coudron+ and Chung+ proved **unbounded** rate of expansion.

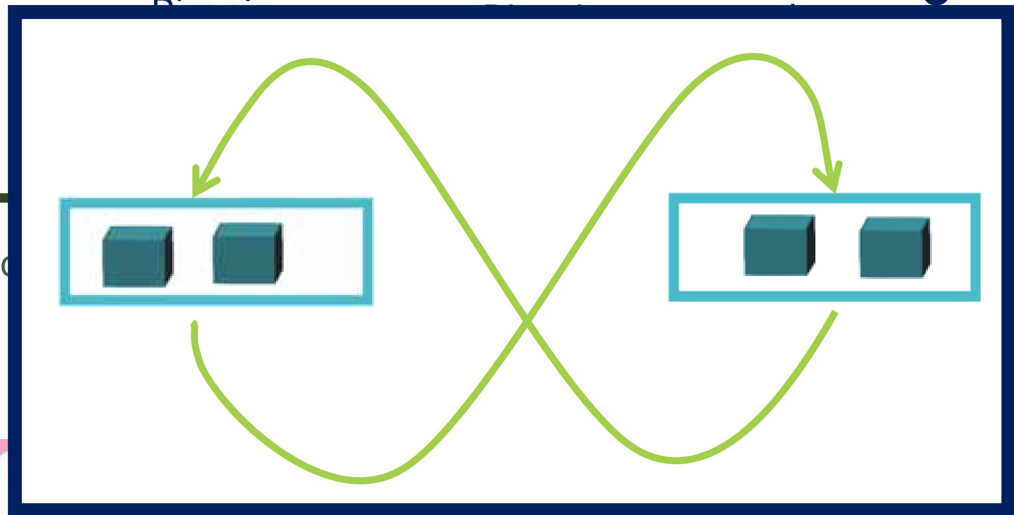

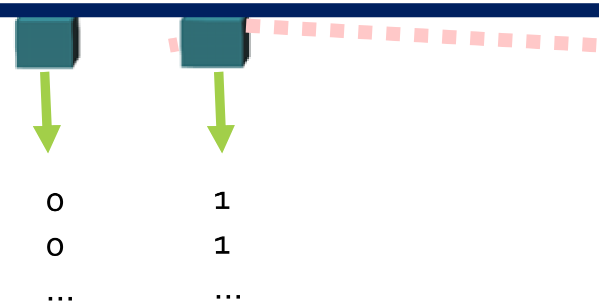

Timeline

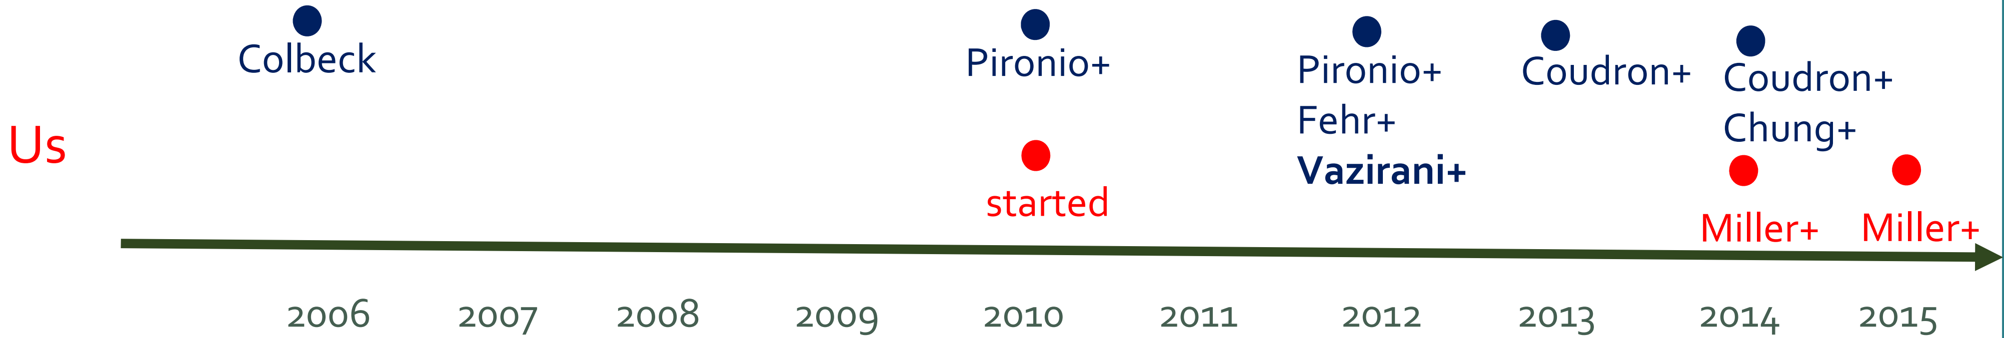

Coudron+ and Chung+ proved **unbounded** rate of expansion.

Our current paper: Covers all nonlocal games, maximal noise tolerance.

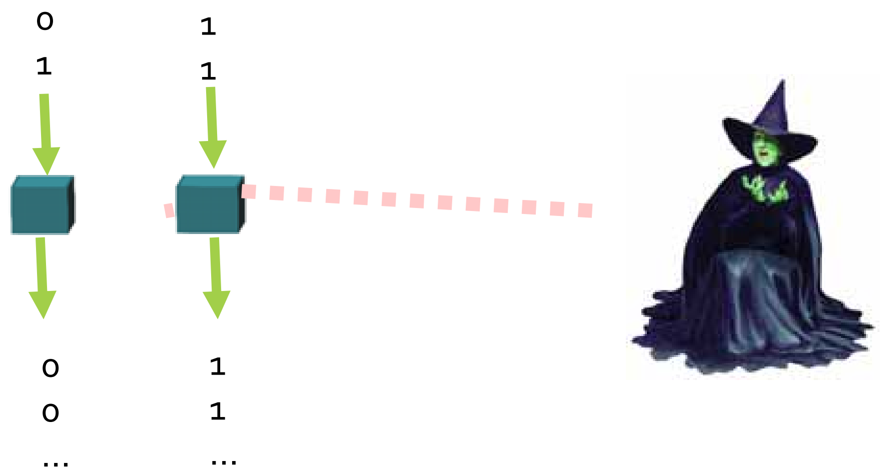

Current Result

The spot-checking protocol

(taken from Coudron-Vidick-Yuen 2013)

Let G = nonlocal game, a = fixed input.

1. Run the device N times. During “game rounds,” play G. Otherwise, just input a.

2. If the average score during game rounds was < C, abort.

3. Otherwise, apply randomness extractor.

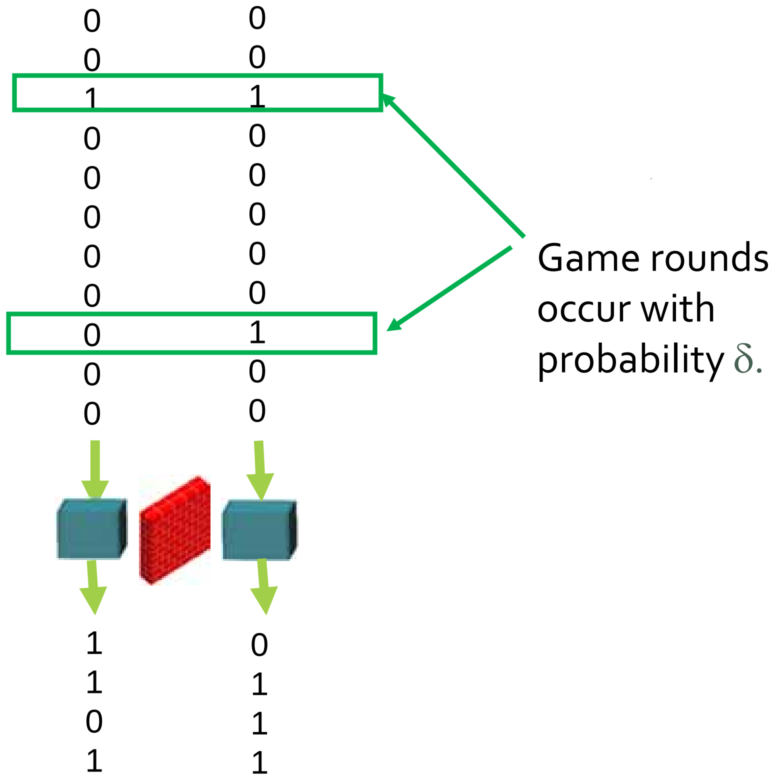

The spot-checking protoco

Let G = nonlocal game, a = fixed input.

1. Run the device N times. During "game rounds," play G. Otherwise, just input a.

2. If the average score during game rounds was < C, abort.

3. Otherwise, apply randomness extractor.

Let $W_{G,a}$ = optimal score among devices that are deterministic on a.

Thm (CVY 13): The protocol is secure against an unentangled adversary if $C > W_{G,a}$.

Thm (MS 15): The protocol is secure against any adversary if $C > W_{G,a}$.

Best possible!

How much randomness (MS 15)

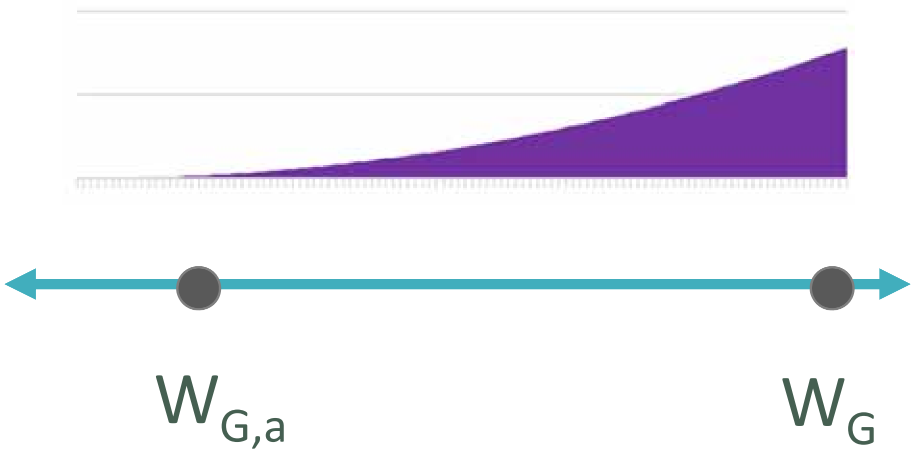

noise threshold vs.
# of random bits per round

$$y = \frac{2.88(x - W_{G,a})^2}{|output| - 1}$$

Proof Techniques

A Mathematical Preliminary

Consider the function $f(X) = \mathrm{Tr}\,[\,|X|^{1+\varepsilon}]$.

- If X is a density operator, f measures how deterministic X is. (Smaller = more random.)

- f is "almost" a norm on Hermitian operators.

A Mathematical Preliminary

The function Tr [$|X|^{1+\varepsilon}$] is uniformly convex. [Ball+ 94]

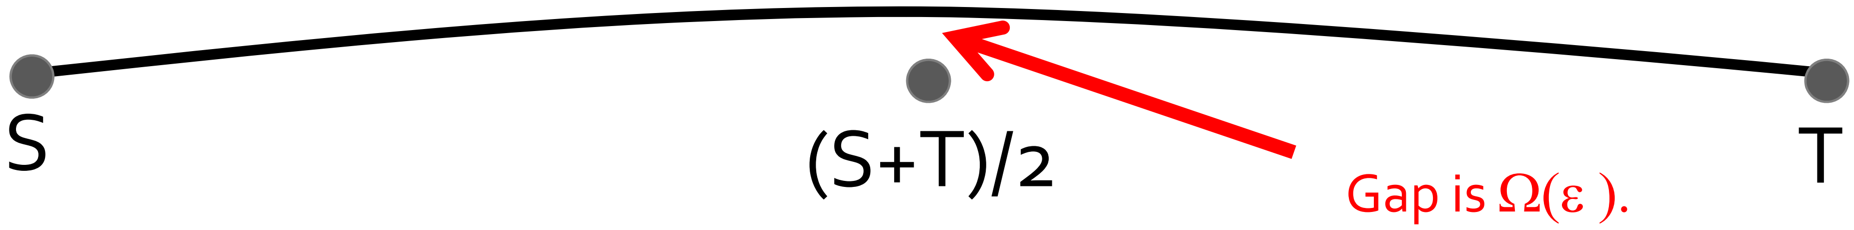

A Mathematical Preliminary

Consequence [MS 15]: Suppose $\phi \mapsto \phi'$ is the result of a binary measurement.

$$\phi' = \frac{\phi + U\phi U^*}{2}$$

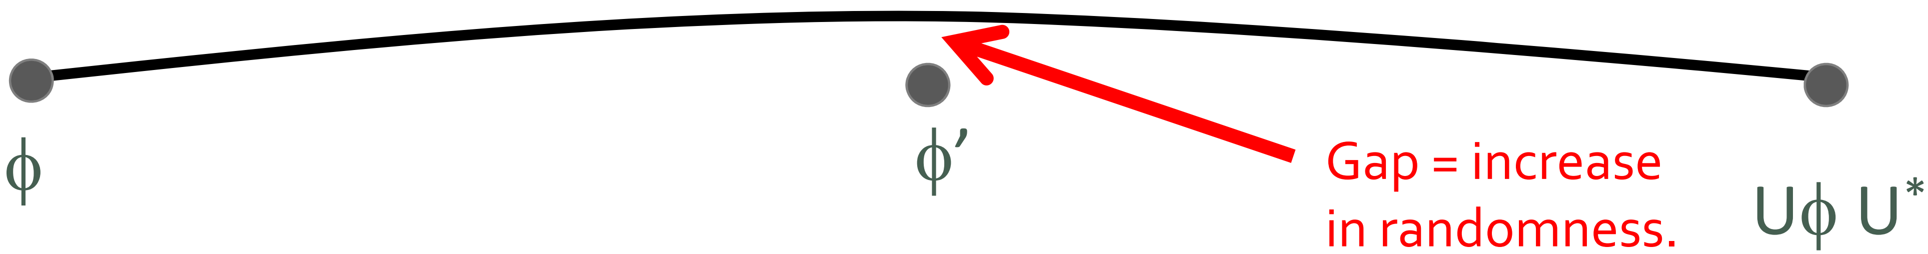

The more **disturbance** caused by a measurement, the more **randomness** it adds.

Call this the $(1+\varepsilon)$–**uncertainty principle.**

How do we prove security for this protocol?

Let G = nonlocal game, a = fixed input.

1. Run the device N times. During "game rounds," play G. Otherwise, just input a.
2. If the average score during game rounds was < C, abort.
3. Otherwise, apply randomness extractor.

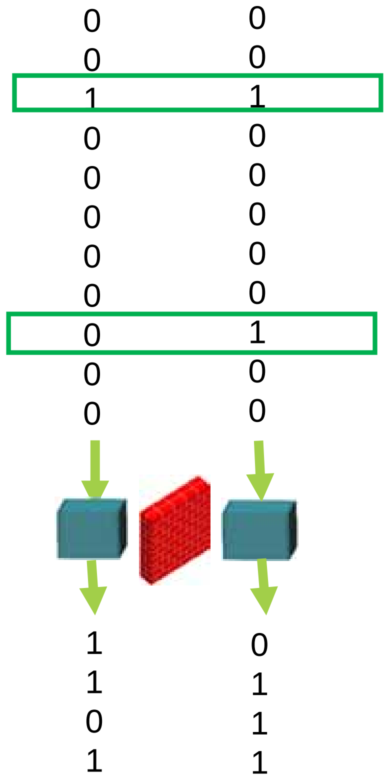

How do we prove security

Let G = nonlocal game, a = fixed input.

1. Run the device N times. During "game rounds," play G. Otherwise, just input a.
2. If the average score during game rounds was < C, abort.
3. Otherwise, apply randomness extractor.

A starting point:

Suppose π is a function such that any device satisfies

H (output | input = a) >= π (P (win))

Prop (easy)**:** In the **non-adversarial IID case,** the protocol produces at least π (C) N extractable bits.

π = "simple rate curve"

What about non-IID?

Compare:

* **von Neumann** entropy (H)
* **Renyi** entropy ($H_{1+\varepsilon}$).

$H_{1+\varepsilon}$ proves extractable bits in the non-IID case!
But it's hard to relate to the winning probability.

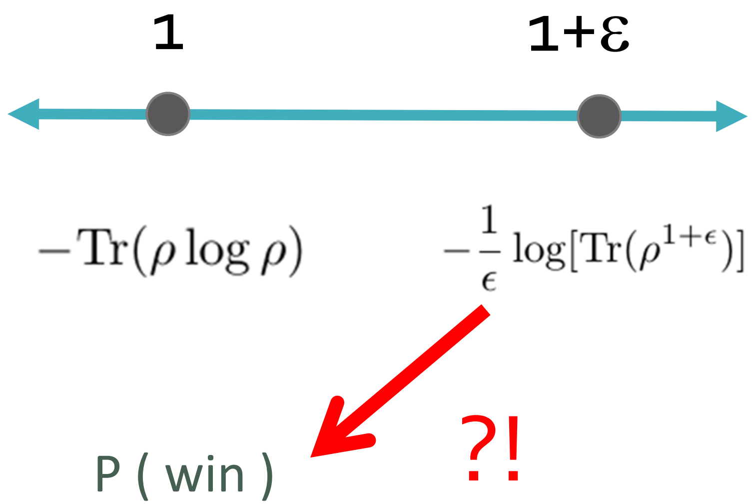

What about non-IID?

Compare:

* **von Neumann** entropy (H)
* **Renyi** entropy ($H_{1+\varepsilon}$).

$H_{1+\varepsilon}$ proves extractable bits in the non-IID case!
But it's hard to relate to the winning probability.

Def: the (1+ε)-**winning probability** of a device is

$$\frac{\mathrm{Tr}[\rho_{win}^{1+\epsilon}]}{\mathrm{Tr}[\rho^{1+\epsilon}]}$$

where ρ = adversary's state.

$-\mathrm{Tr}(\rho \log \rho)$ $\quad$ $-\frac{1}{\epsilon}\log[\mathrm{Tr}(\rho^{1+\epsilon})]$

P (win) $\quad$ $P_{1+\varepsilon}$ (win)

What about non-IID?

Def: π is a **strong rate curve** for the game G on input a if for all devices D,

$$H_{1+\varepsilon}\text{(output on input a | adversary)}$$

is greater than or equal to

$$\pi\,(P_{1+\varepsilon}\text{(win)}) - O_{\text{dev.-ind.}}(\varepsilon).$$

Thm [MS 15]: If π is a strong rate curve, then the spot-checking protocol produces $N \cdot \pi$ (C) extractable bits. (N = # of rounds, C = noise threshold.)

1	1+ε
$-\mathrm{Tr}(\rho \log \rho)$	$-\frac{1}{\epsilon}\log[\mathrm{Tr}(\rho^{1+\epsilon})]$
P (win)	$P_{1+\varepsilon}$ (win)

How do we prove strong rate curves?

Want: High $P_{1+\varepsilon}$ (win) implies high $H_{1+\varepsilon}$.

Create a new device by pre-measuring w/ input a.

If this brings the score down significantly, then a significant amount of state disturbance has occurred. **(1+ε)–uncertainty principle** says that randomness was generated!

So if $P_{1+\varepsilon}$(win) is significantly larger than $W_{G,a}$, we have randomness.

Pre-apply the measurement for input a.

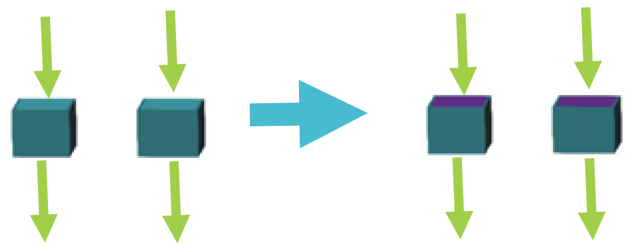

$P_{1+\varepsilon}$(win) vs. $P_{1+\varepsilon}(\text{win}) \leq W_{G,a}$

The universal rate curves

Thm: For any (G,a) the function

$$y = \frac{2.88(x - W_{G,a})^2}{|output| - 1}$$

is a strong rate curve.

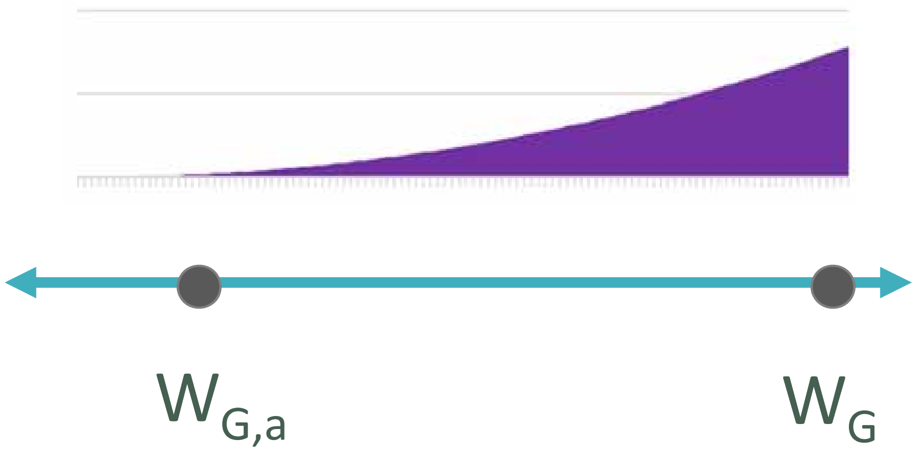

Example: The CHSH Game (2-player, binary)

Inputs	Score if outputs agree	Score if outputs disagree
00	1	0
01	1	0
10	1	0
11	0	1

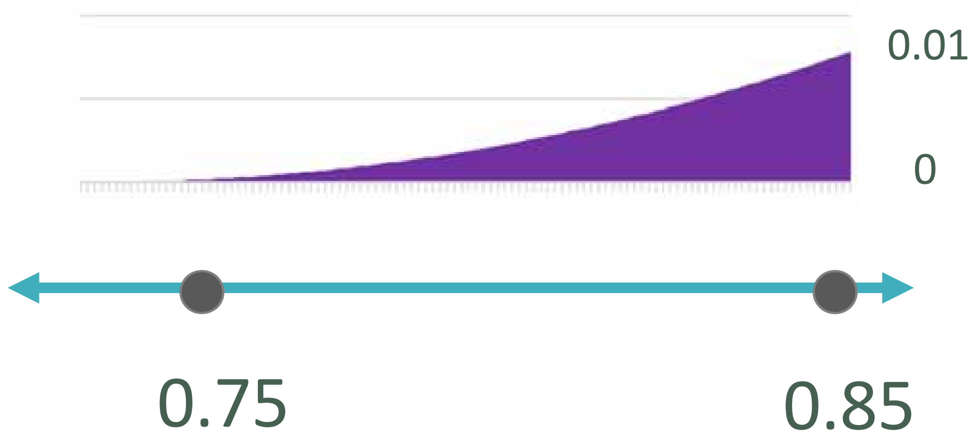

Best possible noise tolerance.

Alternate challenge: Increase the height!

Higher Rate Curves

Self-Testing with CHSH

The quantum device that achieves the optimal CHSH score is unique (state + measurements).

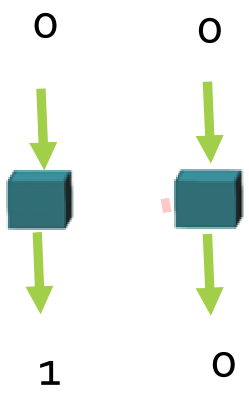

Inputs	Score if outputs agree	Score if outputs disagree
00	1	0
01	1	0
10	1	0
11	0	1

Self-Testing with CHSH

Why?

The only way to maximize the score on **each** input pair is to have a maximally entangled state with measurements at an angle of $\pi/8$ from one another:

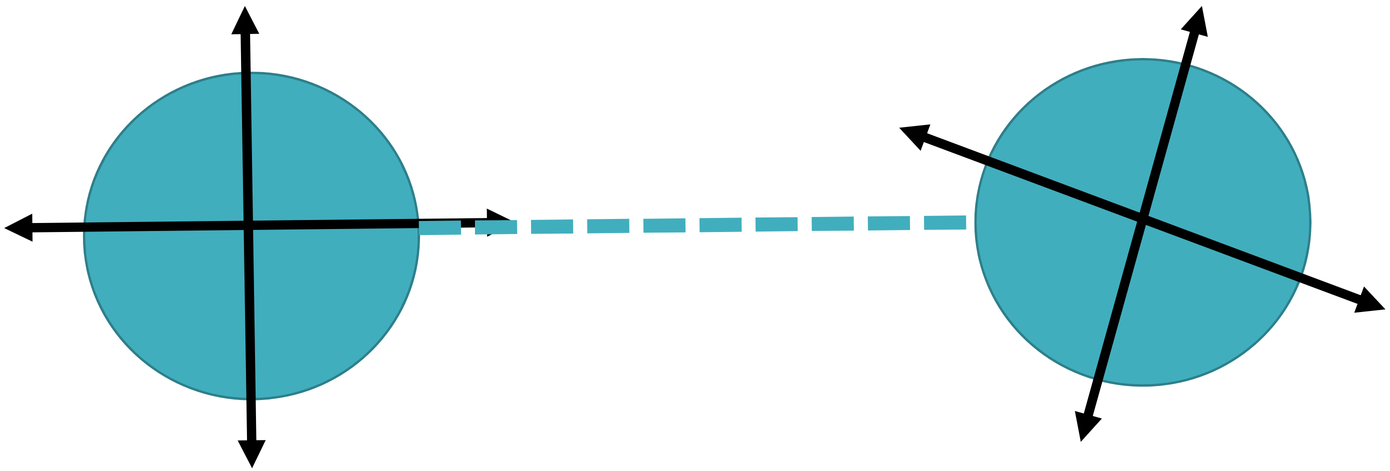

Self-Testing with CHSH

Why?

The only way to maximize the score on **each** input pair is to have a maximally entangled state with measurements at an angle of $\pi/8$ from one another:

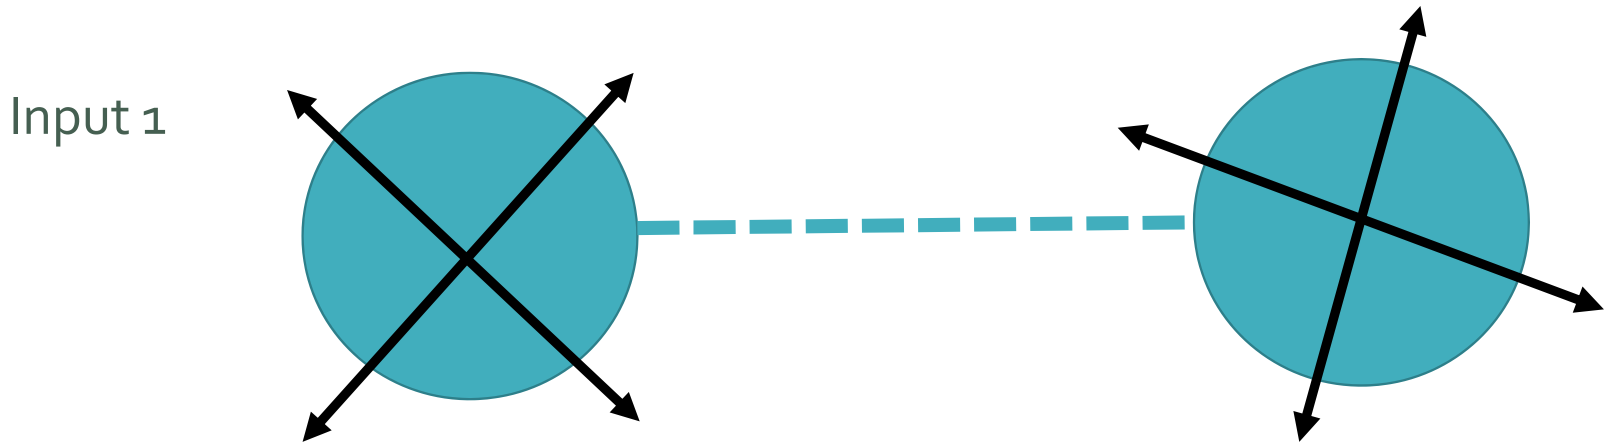

Self-Testing with CHSH

Why?

The only way to maximize the score on **each** input pair is to have a maximally entangled state with measurements at an angle of $\pi/8$ from one another:

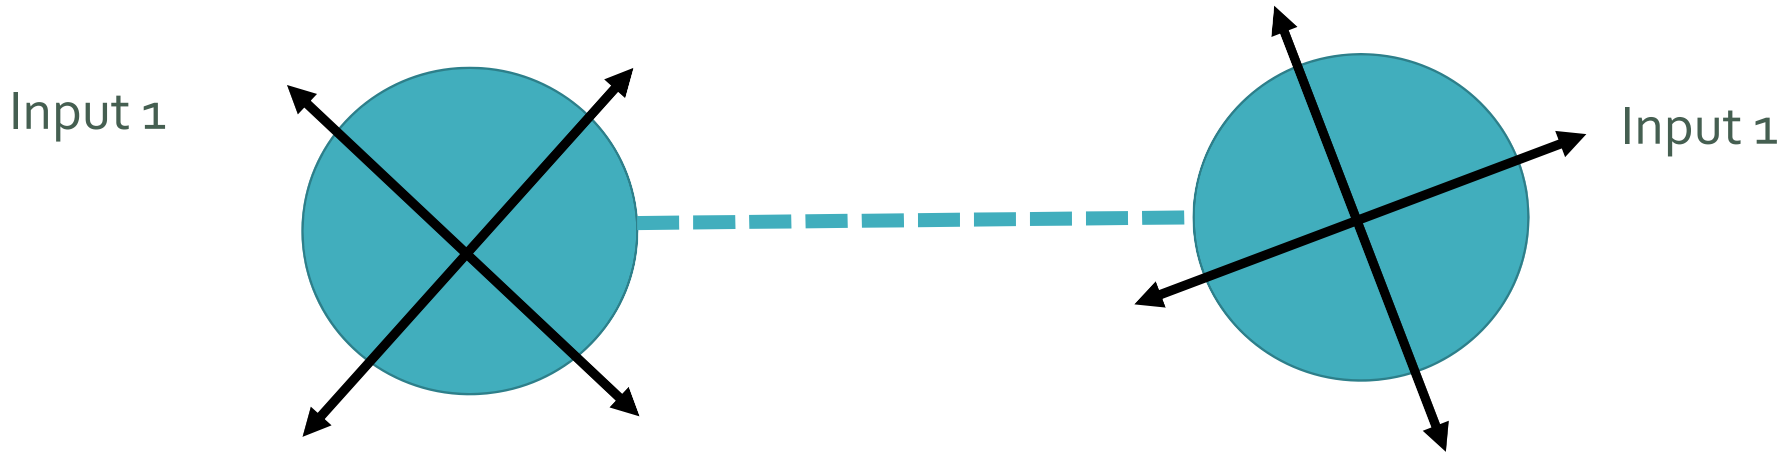

Self-Testing with CHSH

Every device w/ a near optimal score is approximately the same as the optimal one.

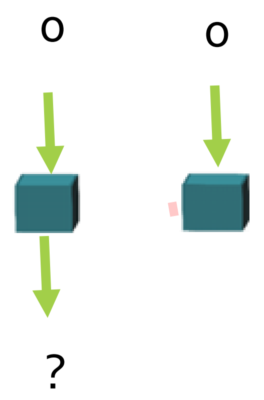

The optimal device gives a perfect coin flip on input 00!
This implies a simple rate curve which approaches 1.

Self-Testing with CHSH

We prove a strong rate curve for CHSH (MS 14):

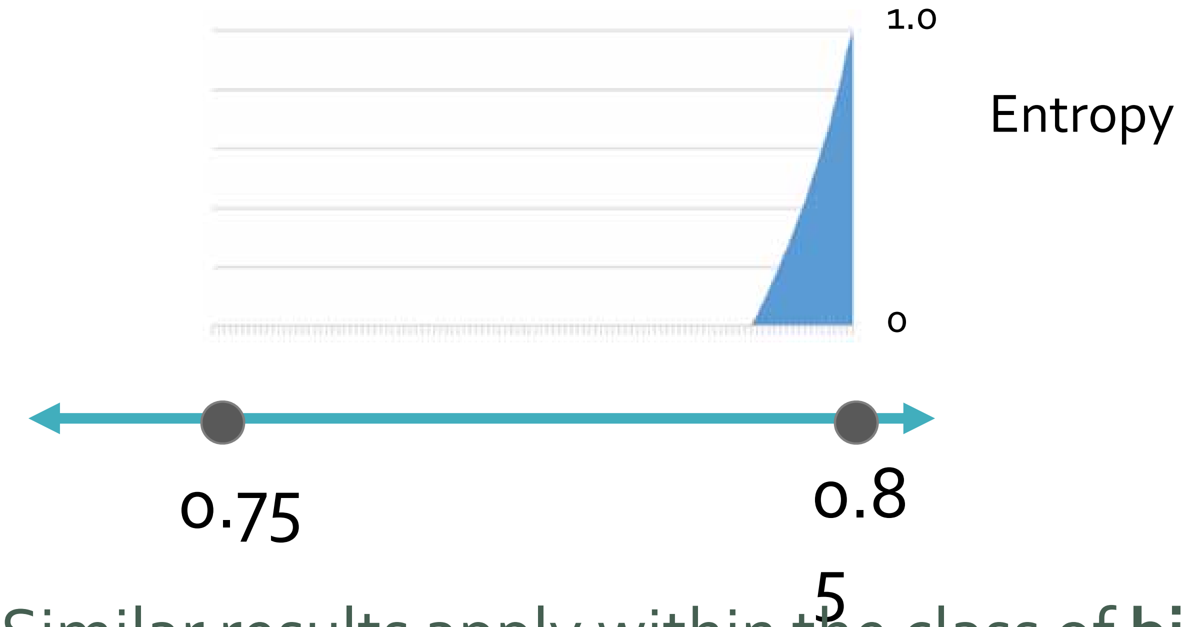

Similar results apply within the class of **binary XOR games.**

Application: QKD

Device-independent Quantum Key Distribution

Our proof can be adapted to give another proof of DI-QKD.

1. Do step 1 of the spot-checking protocol. Communicate to check score.

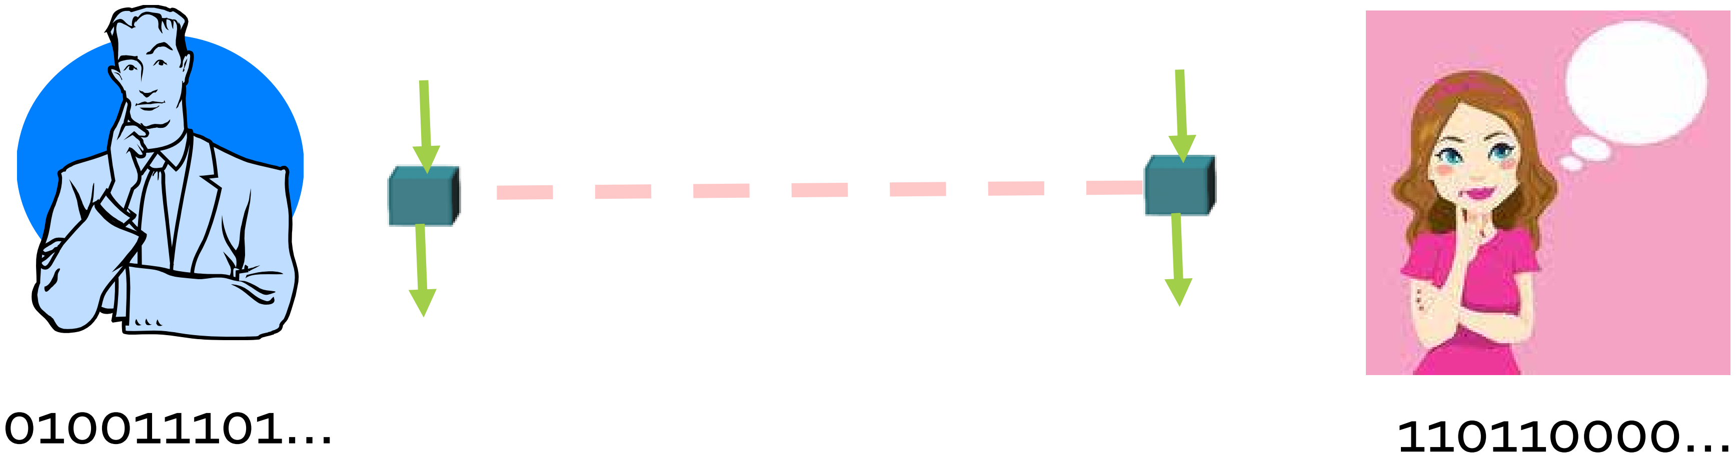

Device-independent Quantum Key Distribution

Our proof can be adapted to give another proof of DI-QKD.

1. Do step 1 of the spot-checking protocol. Communicate to check score.
2. Have Alice make an optimal guess at Bob's bits using her bits.

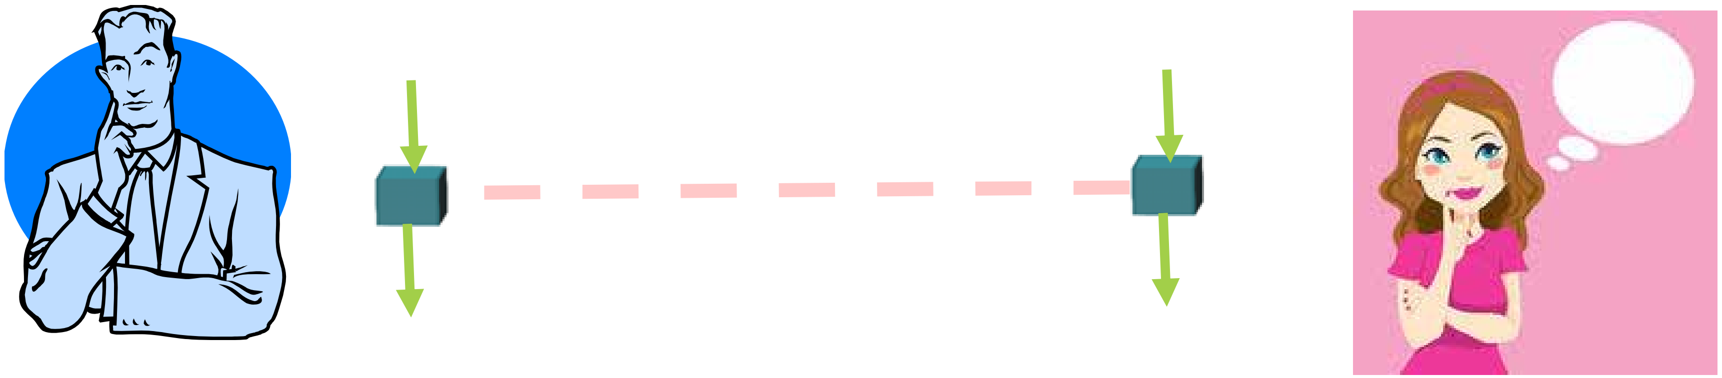

010011101...

011011100...

Device-independent Quantum Key Distribution

Our proof can be adapted to give another proof of DI-QKD.

1. Do step 1 of the spot-checking protocol. Communicate to check score.
2. Have Alice make an optimal guess at Bob's bits using her bits.
3. Perform information reconciliation.

Device-independent Quantum Key Distribution

Our proof can be adapted to give another proof of DI-QKD.

1. Do step 1 of the spot-checking protocol. Communicate to check score.
2. Have Alice make an optimal guess at Bob's bits using her bits.
3. Perform information reconciliation.
4. Perform randomness extraction.

This works if step 1 generates more entropy than is lost at step 3.

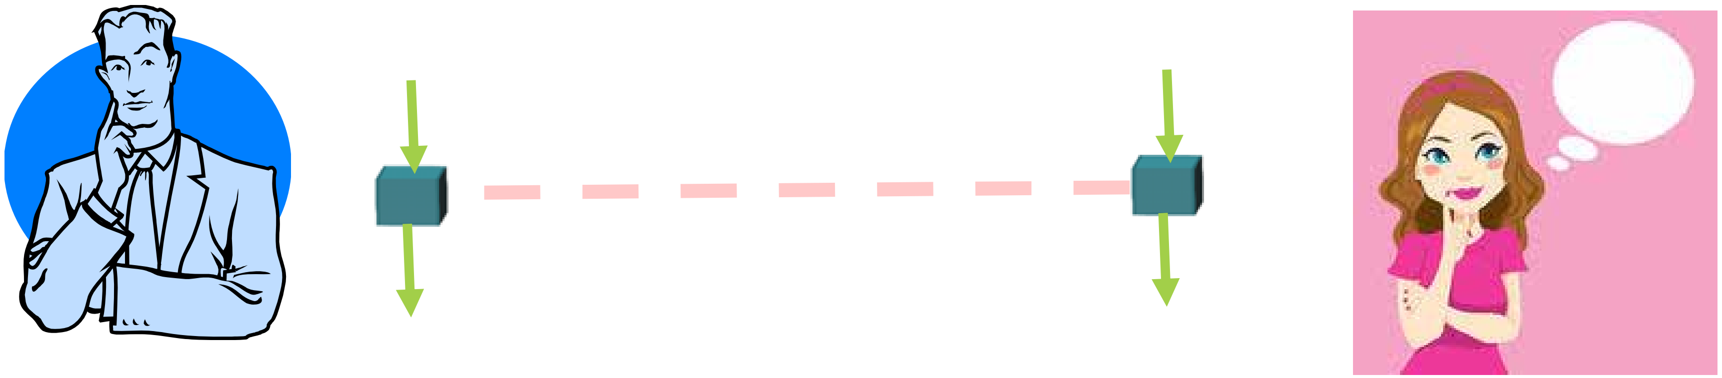

111011100... 111011100...

Open Problems

Prove the best possible rate curves

We have two families of rate curves, neither optimal. What are the best rate curves?

Can we match the classical-adversary rate curves?

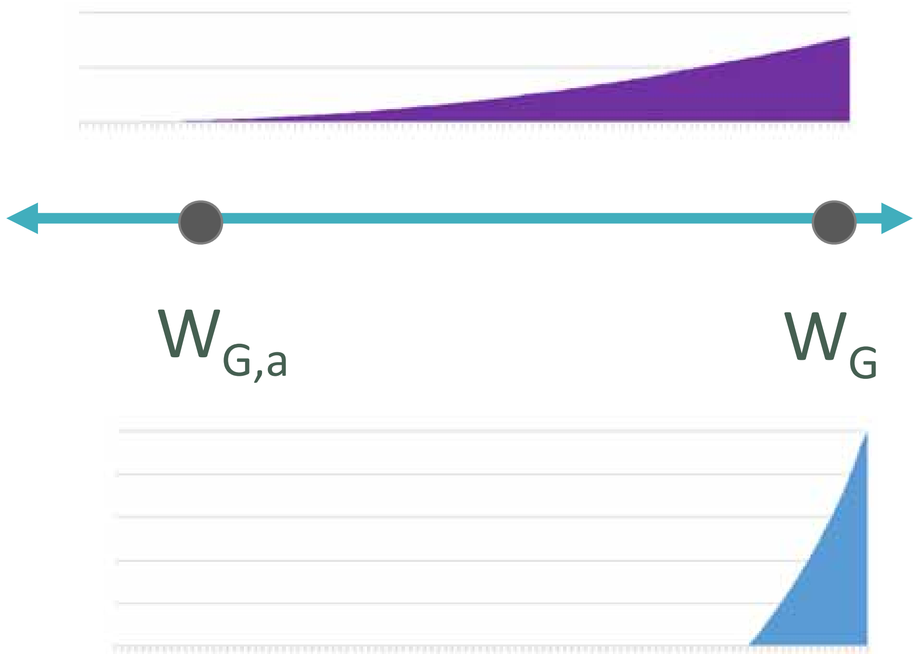

Parallel randomness expansion

Give inputs to the boxes all at once. Can we still verify randomness?

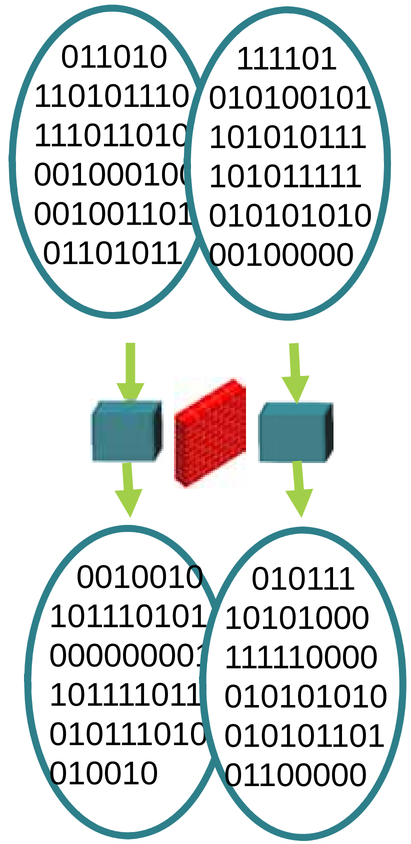

Security proofs for device-independent randomness expansion

Carl A. Miller

University of Michigan, Ann Arbor

Ref: "Universal security for randomness expansion from the spot-checking protocol" (arXiv:1411.6608), with Yaoyun Shi.