Post-Quantum Succinct Arguments: Breaking the Quantum Rewinding Barrier

Fermi Ma

Princeton → Simons & Berkeley

joint work with
Alessandro Chiesa, Nicholas Spooner, and Mark Zhandry

Why are quantum computers a threat to cryptography?

Why are quantum computers a threat to cryptography?

To answer this, recall how cryptographers *prove* security.

Fundamental formula of cryptography

Crypto security proof = Computational hardness assumption + Reduction

Fundamental formula of cryptography

Crypto security proof	=	Computational hardness assumption	+	Reduction

Ex: invert one-way function, factoring, discrete log, lattice problems, etc.

Fundamental formula of cryptography

Crypto security proof = Computational hardness assumption + Reduction

Ex: invert one-way function, factoring, discrete log, lattice problems, etc.

Any *efficient* attack on the protocol
→ Break underlying hardness assumption

Why are quantum computers a threat?

Why are quantum computers a threat?

Crypto security proof = Computational hardness assumption + Reduction

Ex: invert one-way function, ~~factoring~~, ~~discrete log~~, lattice problems, etc.

Simple answer: Shor's algorithm breaks widely-used hardness assumptions

Post-quantum cryptography
(classical crypto secure against quantum attack)

Crypto security proof	=	Computational hardness assumption	+	Reduction

Minimum requirement for *post-quantum* crypto:
hard problem must resist quantum attacks

Post-quantum cryptography

(classical crypto secure against quantum attack)

Crypto security proof = Computational hardness assumption + Reduction

Minimum requirement for *post-quantum* crypto:
hard problem must resist quantum attacks

Fortunately, we have candidate hard problems.

Post-quantum cryptography

(classical crypto secure against quantum attack)

Crypto security proof = *Post-quantum* hardness assumption + Reduction

Ex: lattice problems, isogenies, etc.

Minimum requirement for *post-quantum* crypto:
hard problem must resist quantum attacks

Fortunately, we have candidate hard problems.

Post-quantum cryptography

(classical crypto secure against quantum attack)

Crypto security proof	=	*Post-quantum* hardness assumption	+	Reduction

Common misconception:

Post-quantum assumptions are all we need for post-quantum cryptography.

Post-quantum cryptography
(classical crypto secure against quantum attack)

Crypto security proof	=	*Post-quantum* hardness assumption	+	Reduction

Common misconception:

Post-quantum assumptions are all we need for post-quantum cryptography.

Key point:
the *security reduction* must be *quantum-compatible*!

Post-quantum cryptography

(classical crypto secure against quantum attack)

Post-quantum security proof	=	*Post-quantum* hardness assumption	+	*Quantum-compatible* reduction

Classical reduction: Any *classical* attack on the protocol → (classical) attack on the assumption

Post-quantum cryptography

(classical crypto secure against quantum attack)

Post-quantum security proof	=	*Post-quantum* hardness assumption	+	*Quantum-compatible* reduction

Classical reduction: Any *classical* attack on the protocol → (classical) attack on the assumption

We need: Any *quantum* attack on the protocol → (quantum) attack on the assumption

Some classical reductions are quantum-compatible, but problems arise if the reduction *rewinds the adversary*.

Some classical reductions are quantum-compatible, but problems arise if the reduction *rewinds the adversary*.

Ex: midway through an execution, the reduction saves the adversary’s state and runs it on *multiple challenges*.

Some classical reductions are quantum-compatible, but problems arise if the reduction *rewinds the adversary*.

Ex: midway through an execution, the reduction saves the adversary's state and runs it on *multiple challenges*.

Some classical reductions are quantum-compatible, but problems arise if the reduction *rewinds the adversary*.

Ex: midway through an execution, the reduction saves the adversary's state and runs it on *multiple challenges*.

Some classical reductions are quantum-compatible, but problems arise if the reduction *rewinds the adversary*.

Ex: midway through an execution, the reduction saves the adversary's state and runs it on *multiple challenges*.

Some classical reductions are quantum-compatible, but problems arise if the reduction *rewinds the adversary*.

Ex: midway through an execution, the reduction saves the adversary's state and runs it on *multiple challenges*.

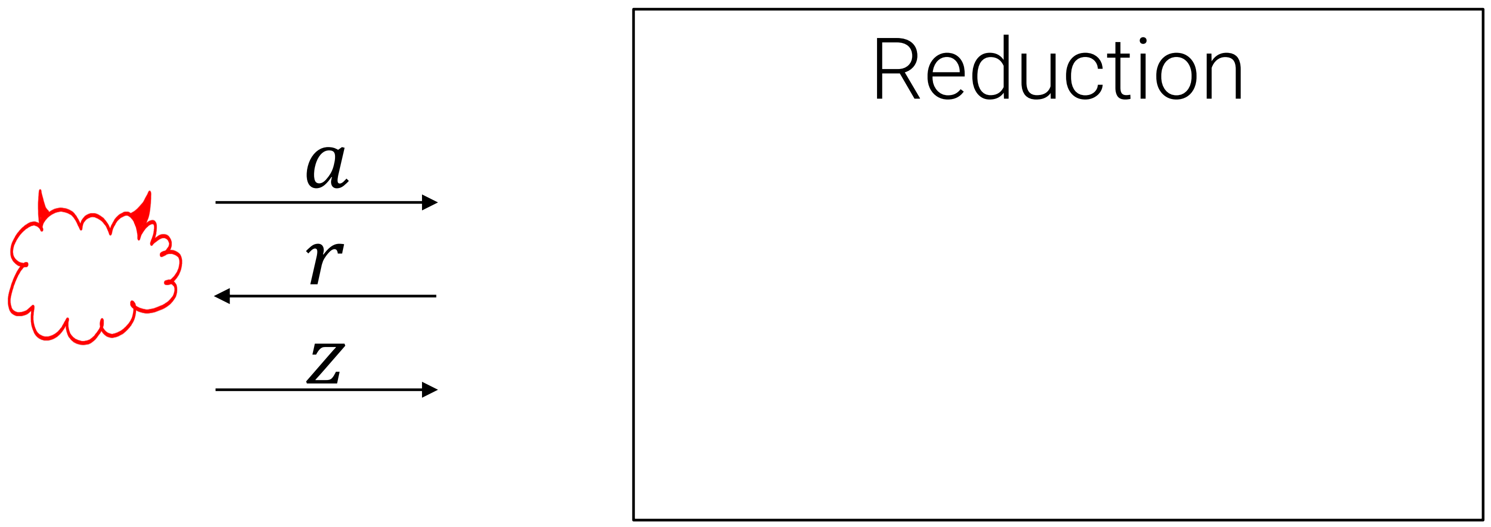

Some classical reductions are quantum-compatible, but problems arise if the reduction *rewinds the adversary*.

Ex: midway through an execution, the reduction saves the adversary's state and runs it on *multiple challenges*.

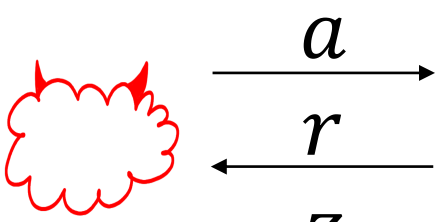

Reduction

1) Record (a, r, z)

Some classical reductions are quantum-compatible, but problems arise if the reduction *rewinds the adversary*.

Ex: midway through an execution, the reduction saves the adversary's state and runs it on *multiple challenges*.

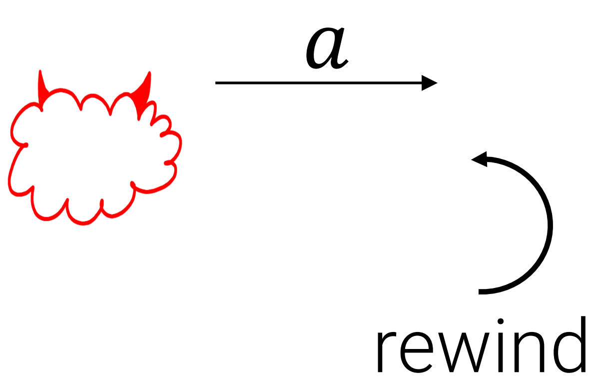

Reduction

1) Record (a, r, z)
2) Rewind

Some classical reductions are quantum-compatible, but problems arise if the reduction *rewinds the adversary*.

Ex: midway through an execution, the reduction saves the adversary's state and runs it on *multiple challenges*.

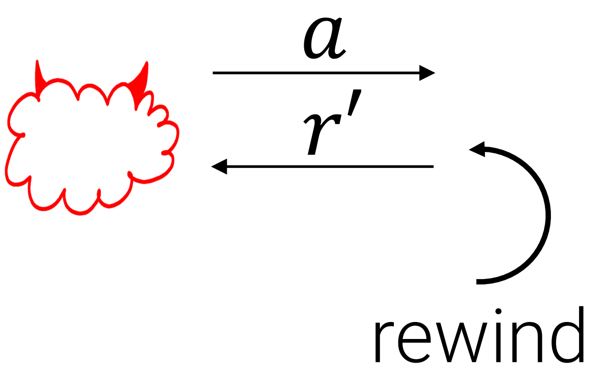

Reduction

1) Record (a, r, z)
2) Rewind

Some classical reductions are quantum-compatible, but problems arise if the reduction *rewinds the adversary*.

Ex: midway through an execution, the reduction saves the adversary's state and runs it on *multiple challenges*.

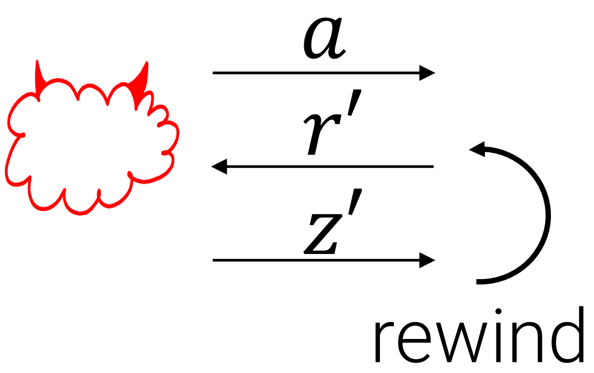

Reduction

1) Record (a, r, z)
2) Rewind

Some classical reductions are quantum-compatible, but problems arise if the reduction *rewinds the adversary*.

Ex: midway through an execution, the reduction saves the adversary's state and runs it on *multiple challenges*.

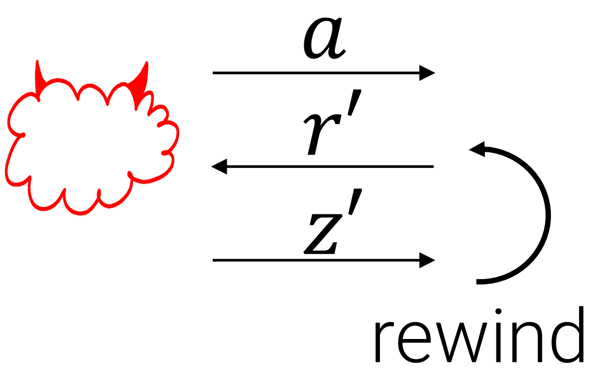

Reduction

1) Record (a, r, z)
2) Rewind
3) Record (a, r', z')

Some classical reductions are quantum-compatible, but problems arise if the reduction *rewinds the adversary*.

Ex: midway through an execution, the reduction saves the adversary's state and runs it on *multiple challenges*.

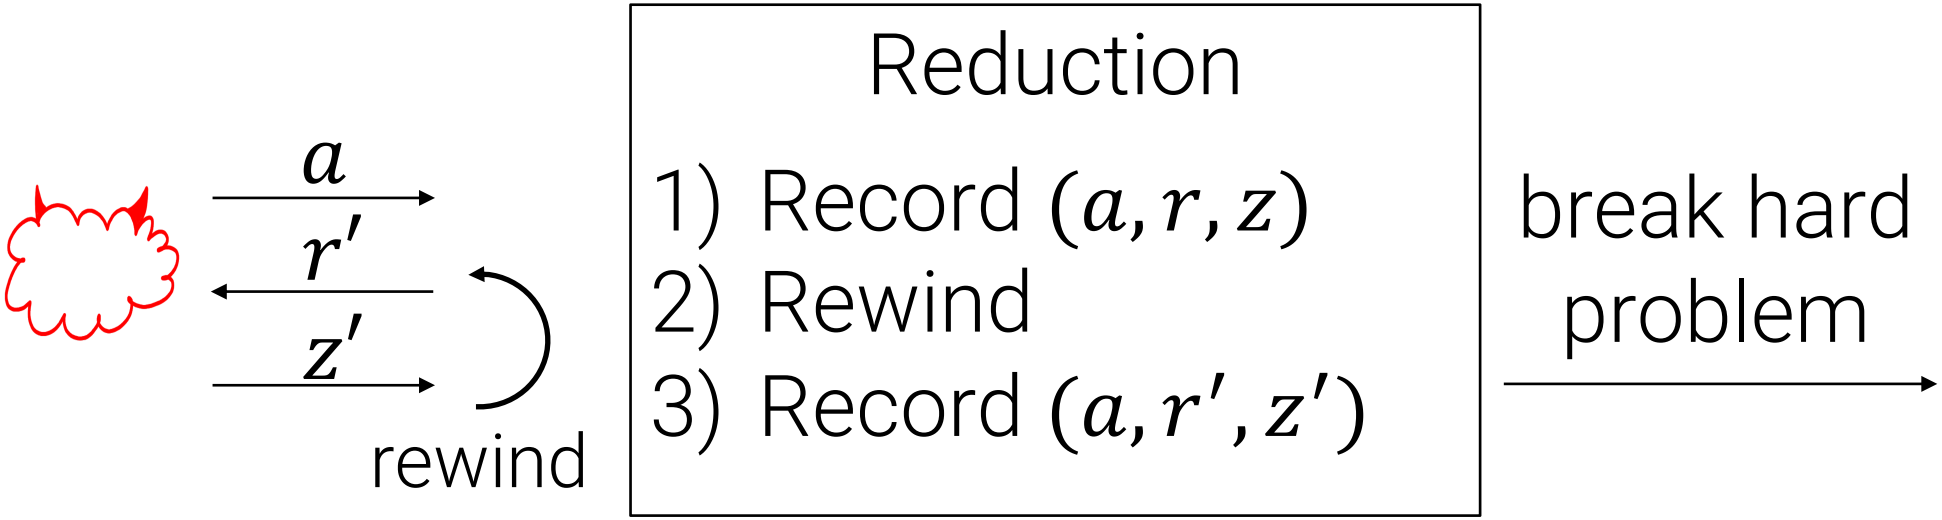

Problem: unclear how to rewind a quantum adversary since measuring its response may disturb its state.

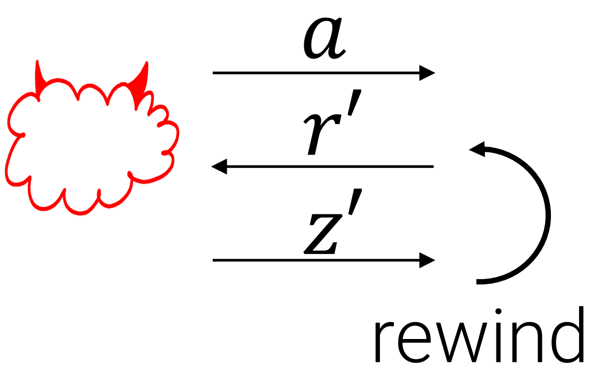

Reduction

1) Record (a, r, z)
2) Rewind
3) Record (a, r', z')

break hard problem

Problem: unclear how to rewind a quantum adversary since measuring its response may disturb its state.

An adversary that detects this disturbance could stop giving valid responses!

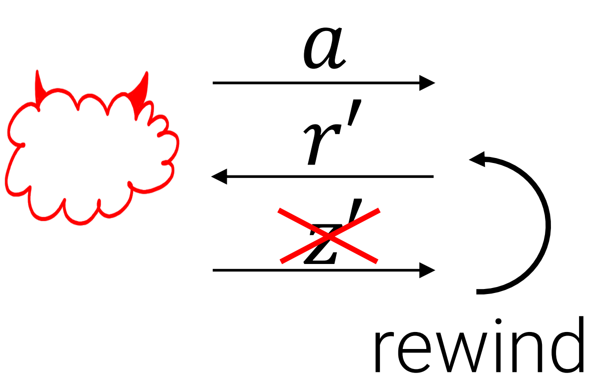

Reduction

1) Record (a, r, z)
2) Rewind
3) ~~Record (a, r', z')~~

~~break hard problem~~

For this talk, the goal of rewinding is to record the adversary's responses to multiple challenges.

For this talk, the goal of rewinding is to record the adversary’s responses to multiple challenges.**

**Disclaimer:
This is how rewinding is commonly used to prove *soundness*, but it doesn't capture applications such as zero knowledge.

For this talk, the goal of rewinding is to record the adversary's responses to multiple challenges.

We'll focus on Kilian's succinct argument protocol, a central result that captures the difficulty of rewinding.

Succinct Arguments for NP [Kilian92]

Claim: $x \in \mathrm{3SAT}$

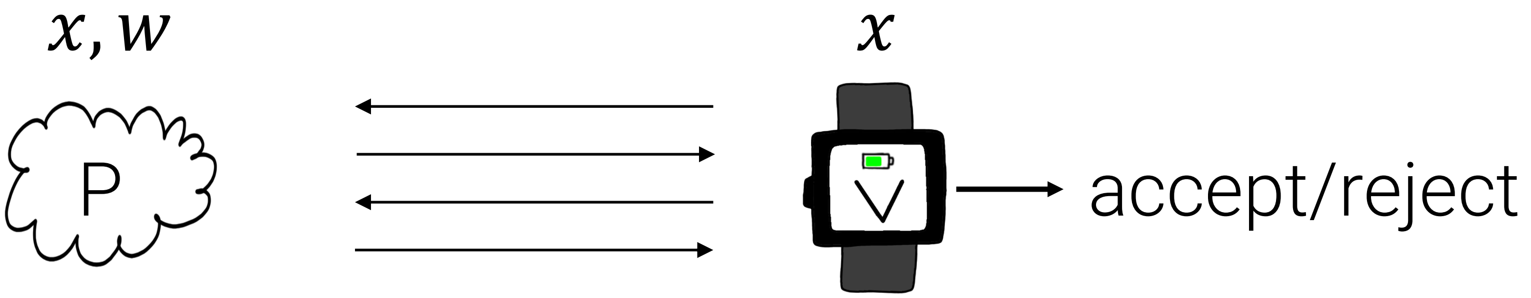

Succinct Arguments for NP [Kilian92]

Claim: $x \in \mathrm{3SAT}$

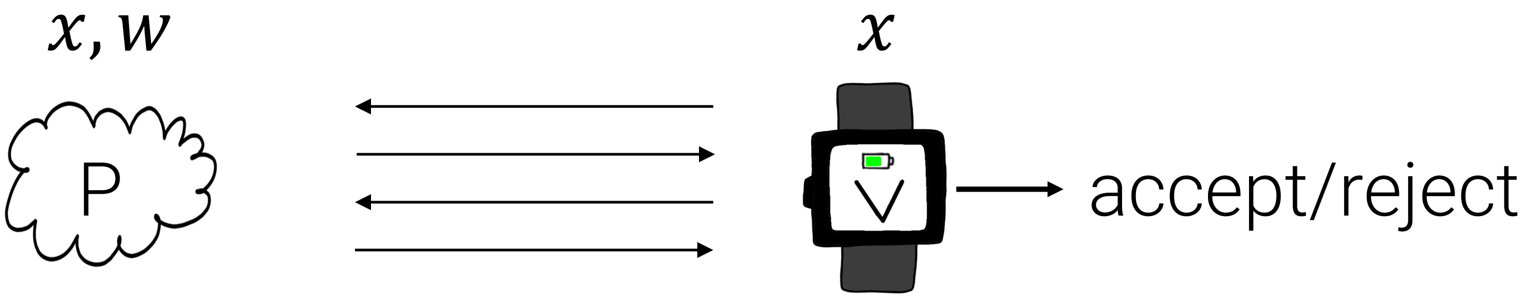

"Succinct" = communication + verifier efficiency is $\mathrm{poly}(\lambda, \log(|x| + |w|))$

Succinct Arguments for NP [Kilian92]

Claim: $x \in \mathrm{3SAT}$

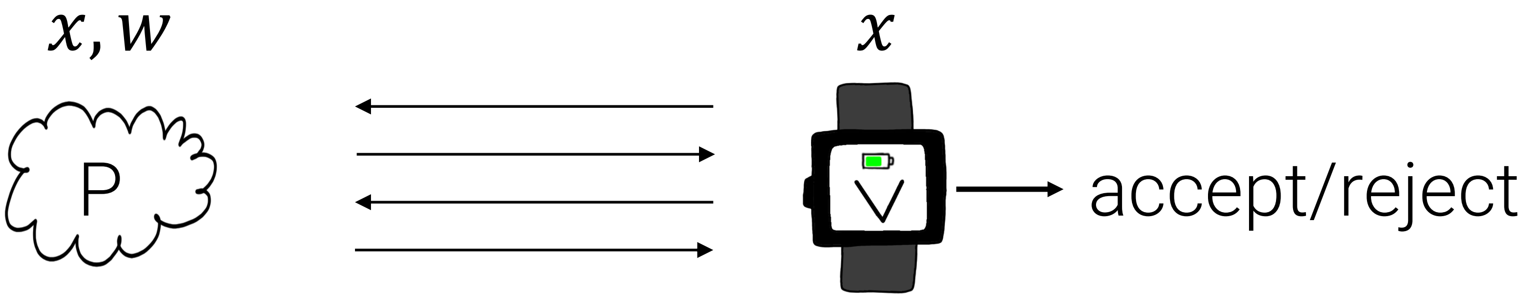

"Succinct" = communication + verifier efficiency is $\mathrm{poly}(\lambda, \log(|x| + |w|))$

"Argument" = sound against *efficient* cheating 

Succinct Arguments for NP [Kilian92]

Claim: $x \in \mathsf{3SAT}$

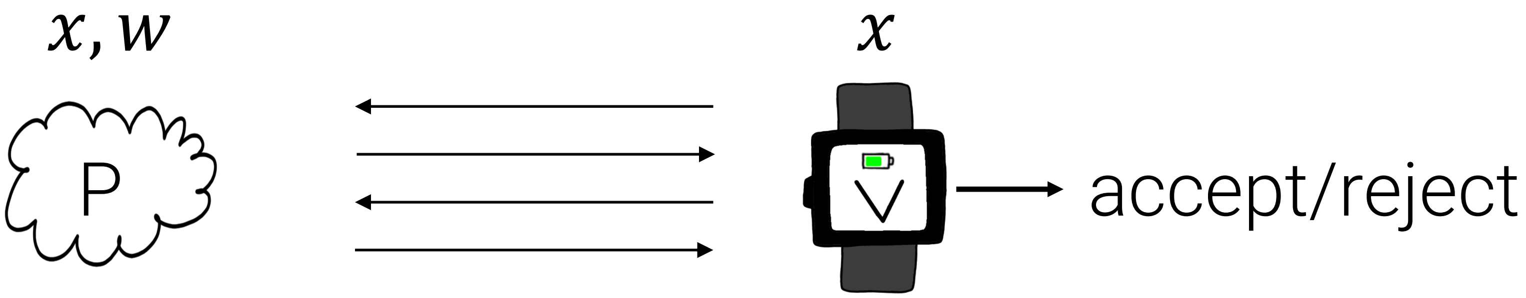

[Kilian92] constructs a 4-message succinct argument for NP from collision-resistant hash functions (CRHFs).

Succinct Arguments for NP [Kilian92]

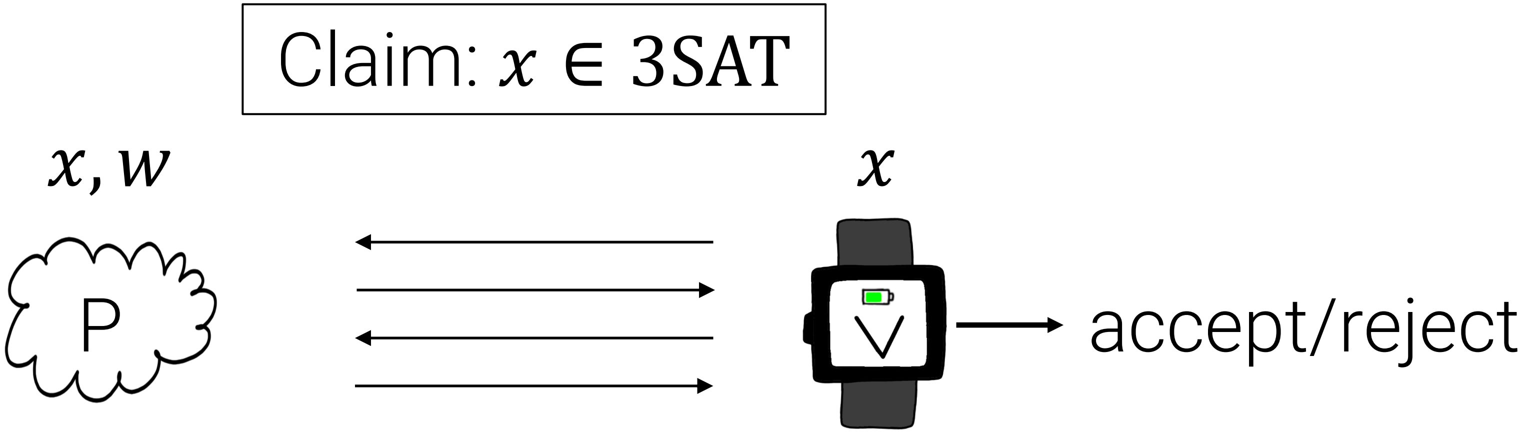

[Kilian92] constructs a 4-message succinct argument for NP from collision-resistant hash functions (CRHFs).

In other words, under a mild computational assumption, any NP statement can be verified $\mathbf{poly}(\lambda, \mathbf{log}(|x| + |w|))$ time!

Succinct Arguments for NP [Kilian92]

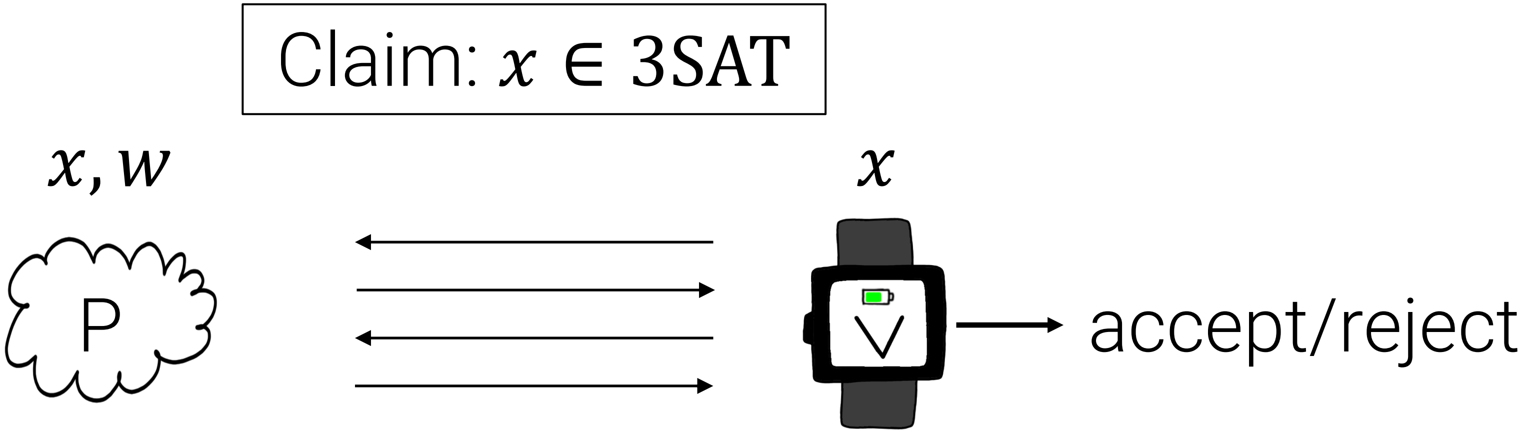

[Kilian92] constructs a 4-message succinct argument for NP from collision-resistant hash functions (CRHFs).

Many applications: universal arguments [BG01], zero knowledge [Barak01], SNARGs [Micali94, BCS16], …

However, post-quantum soundness of Kilian’s protocol remained an open question.

However, post-quantum soundness of Kilian's protocol remained an open question.

- Known reductions for Kilian rewind the attacker to get an *arbitrary polynomial number* of accepting transcripts.

However, post-quantum soundness of Kilian's protocol remained an open question.

- Known reductions for Kilian rewind the attacker to get an *arbitrary polynomial number* of accepting transcripts.
- Existing quantum rewinding techniques [U12,DFMS19] are fundamentally stuck at a *far smaller (constant)* number of rewinds.

However, post-quantum soundness of Kilian's protocol remained an open question.

- Known reductions for Kilian rewind the attacker to get an *arbitrary polynomial number* of accepting transcripts.
- Existing quantum rewinding techniques [U12,DFMS19] are fundamentally stuck at a *far smaller (constant)* number of rewinds.

In this work, we resolve this problem.

This Work

We give a general technique to rewind any quantum attacker as many times as desired.

This Work

We give a general technique to rewind any quantum attacker as many times as desired.

Consequences:

- Kilian is post-quantum sound if the CRHF is quantum-binding*.

This Work

We give a general technique to rewind any quantum attacker as many times as desired.

Consequences:

- Kilian is post-quantum sound if the CRHF is quantum-binding*.

* The CRHF must be *collapsing* — the standard definition of binding for quantum adversaries [Unruh16]. These exist assuming the quantum hardness of Learning with Errors (LWE).

This Work

We give a general technique to rewind any quantum attacker as many times as desired.

Consequences:

- Kilian is post-quantum sound if the CRHF is quantum-binding*.
- Many other protocols, e.g., [GMW86] 3-coloring, [Blum86] Hamiltonicity have optimal post-quantum soundness.

* The CRHF must be *collapsing* — the standard definition of binding for quantum adversaries [Unruh16]. These exist assuming the quantum hardness of Learning with Errors (LWE).

Recall Kilian’s protocol

Kilian's protocol

Compile a ***probabilistically checkable proof**** (PCP) into an interactive argument system using cryptography.

*[BFLS91,FGLSS91,AS92,ALMSS92]

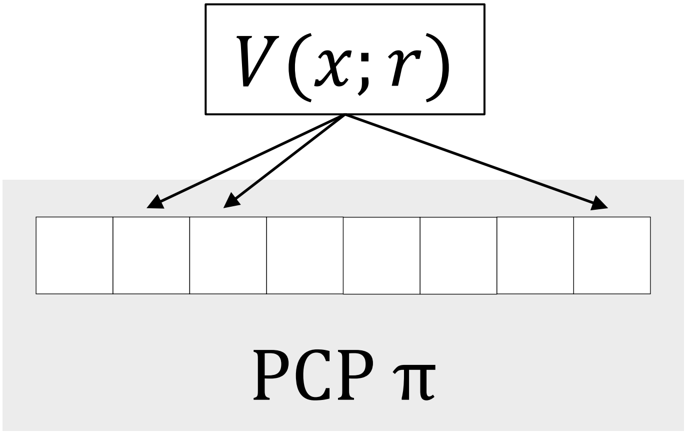

Kilian's protocol

Compile a ***probabilistically checkable proof**** (PCP) into an interactive argument system using cryptography.

*[BFLS91,FGLSS91,AS92,ALMSS92]

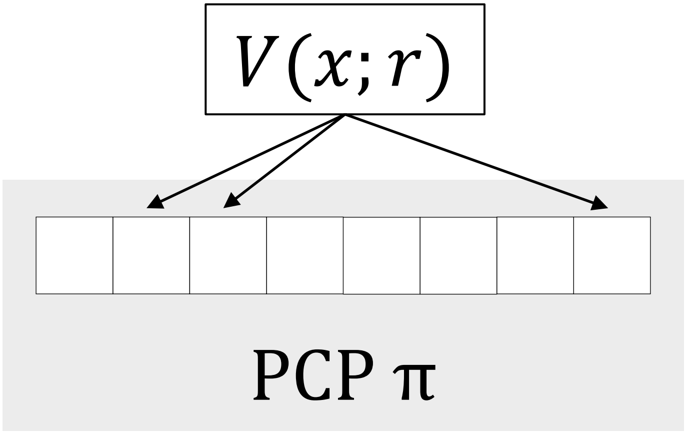

Kilian's protocol

Compile a ***probabilistically checkable proof**** (PCP) into an interactive argument system using cryptography.

*[BFLS91,FGLSS91,AS92,ALMSS92]

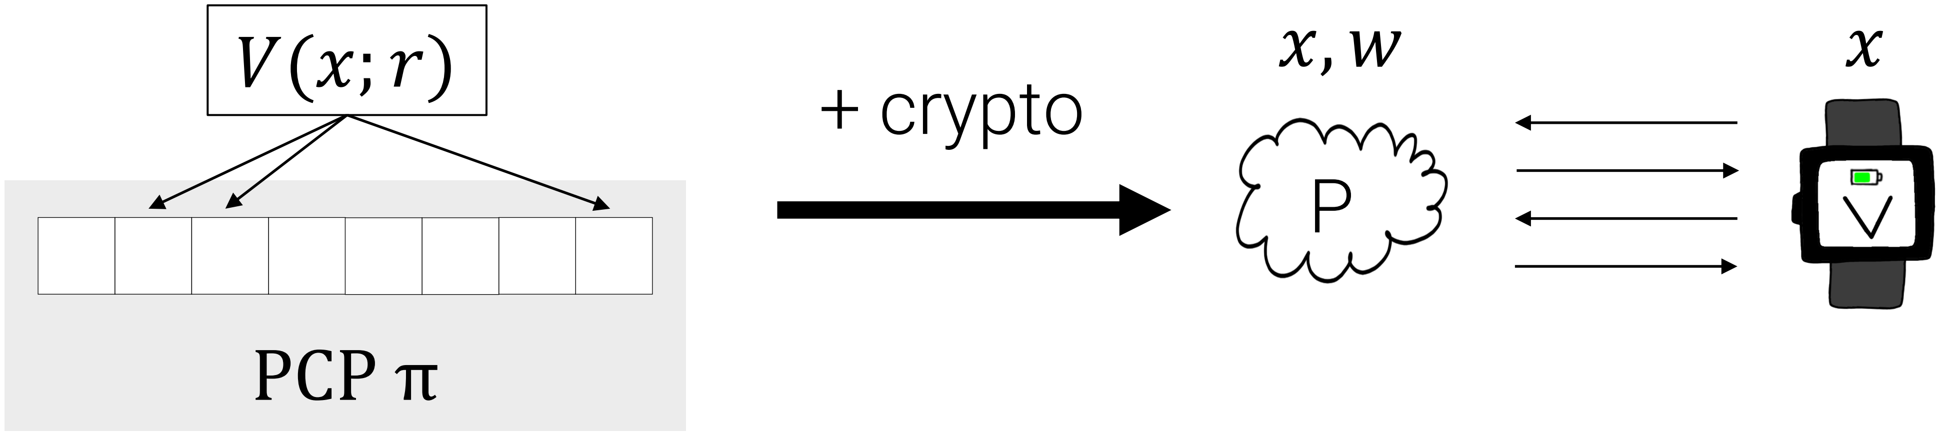

Kilian’s protocol

x, w $\qquad\qquad$ x

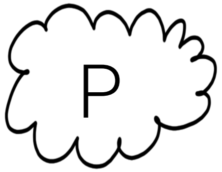

Encode w as PCP π

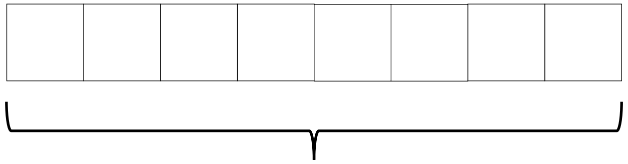

PCP π

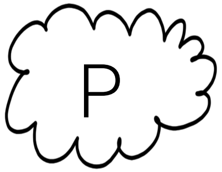

sends short commitment to PCP π.

Kilian's protocol

x, w $\qquad$ x

CRHF h

Encode w as PCP π

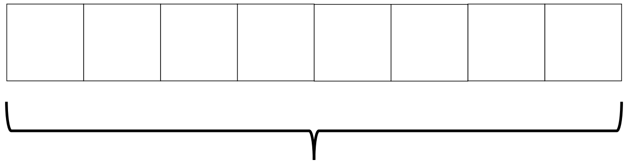

PCP π

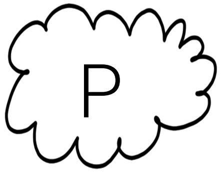

sends short commitment to PCP π.

Kilian's protocol

Encode w as PCP π

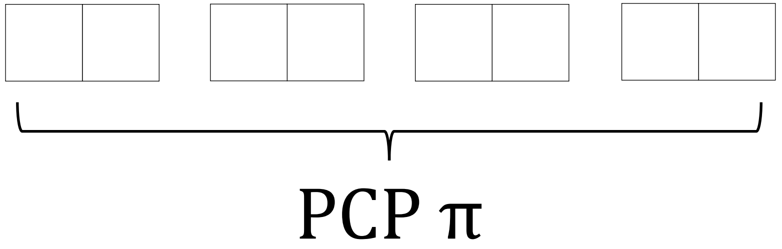

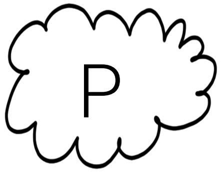

sends short commitment to PCP π.

Kilian's protocol

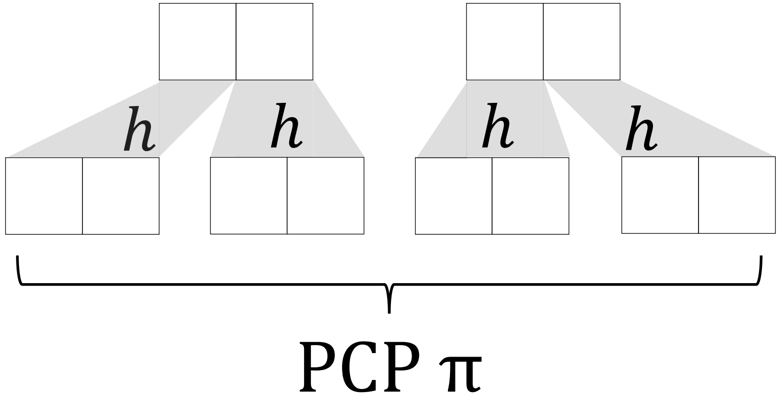

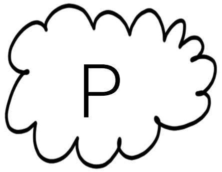

sends short commitment to PCP π.

Kilian's protocol

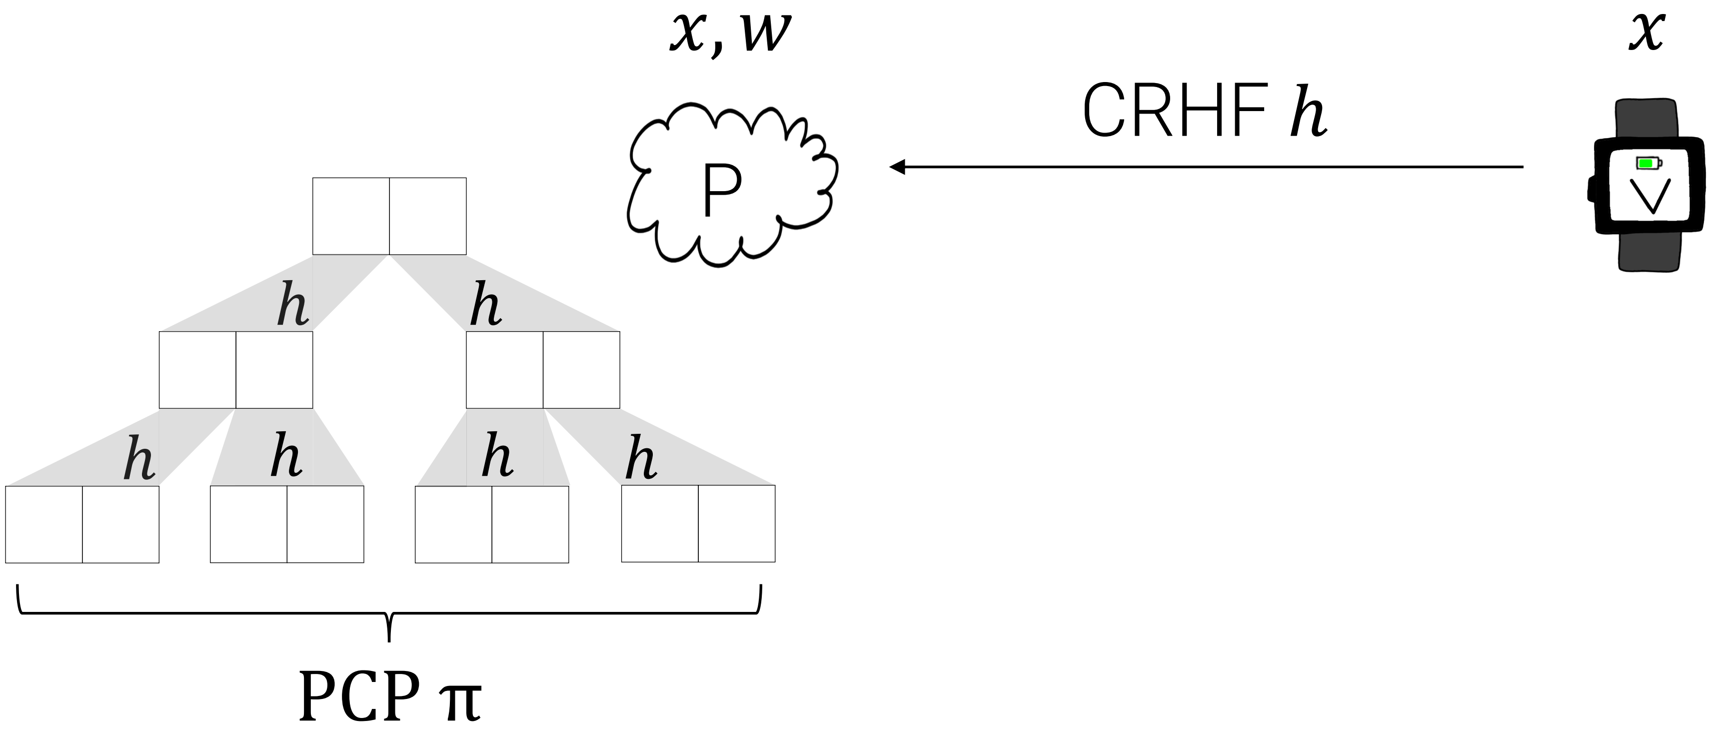

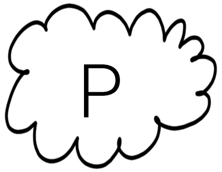

sends short commitment to PCP π.

Kilian’s protocol

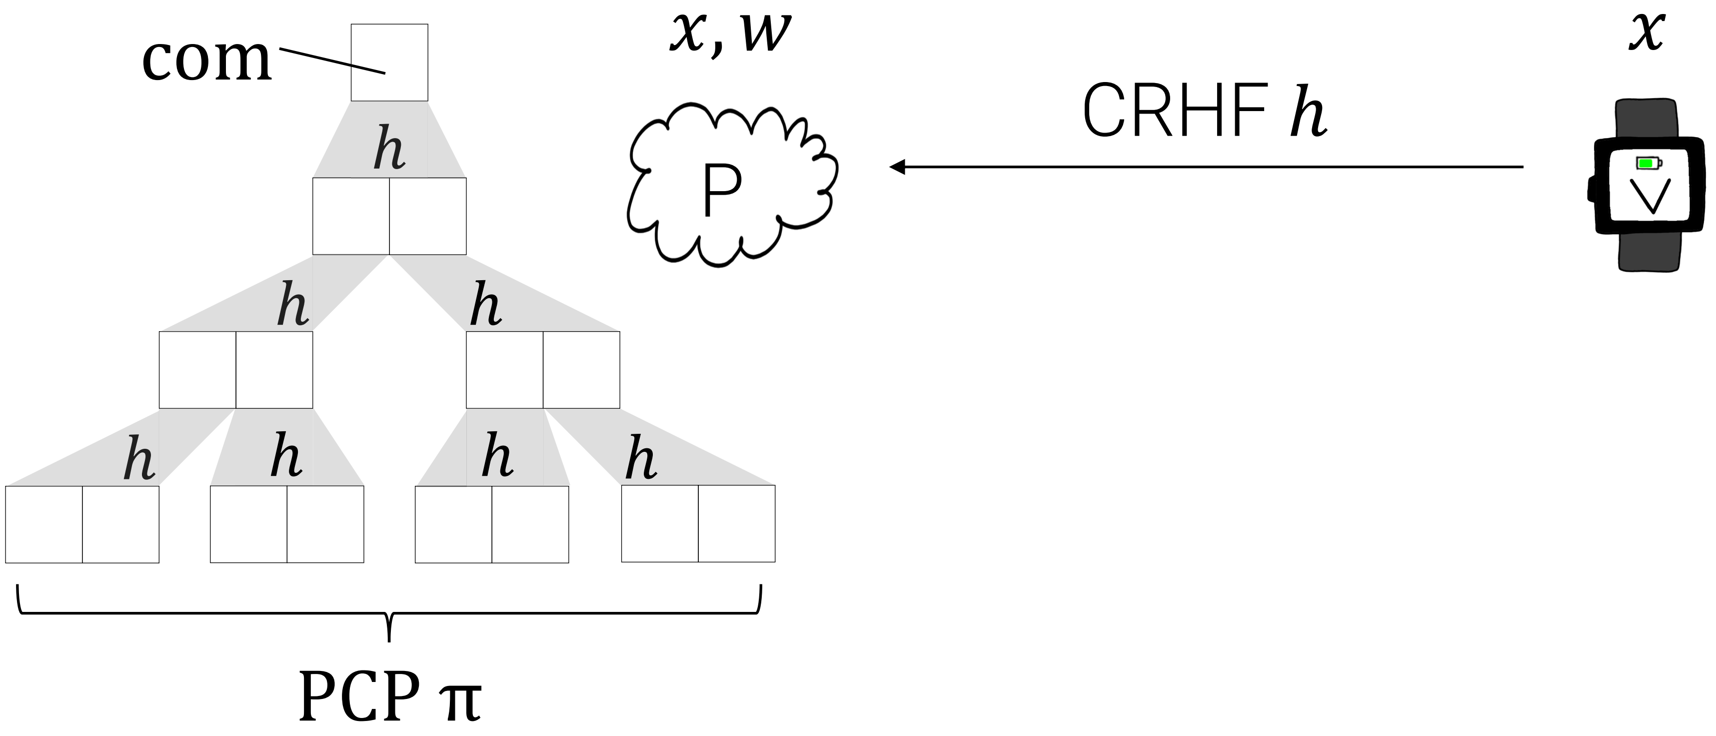

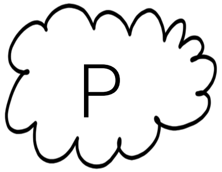

sends short commitment to PCP π.

Kilian's protocol

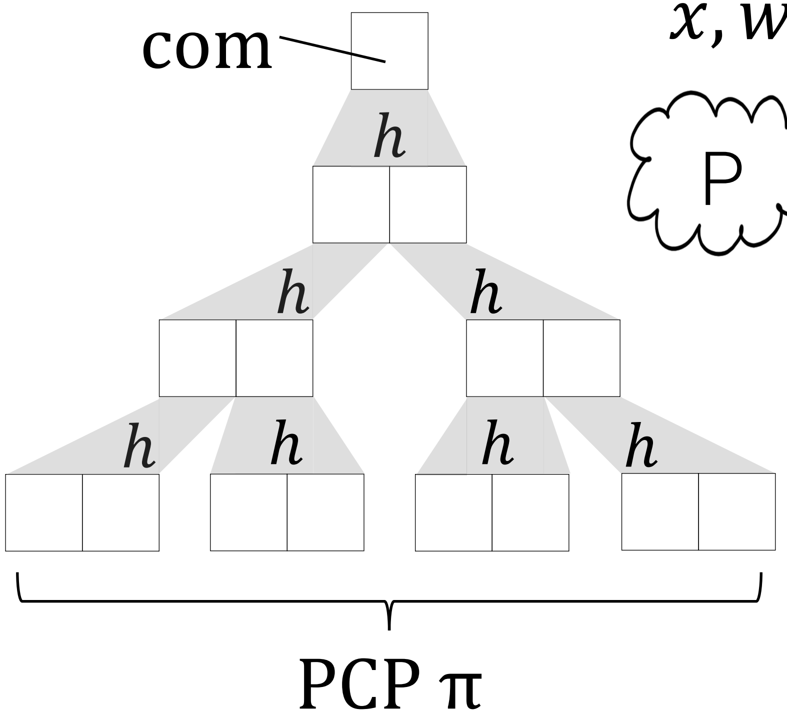

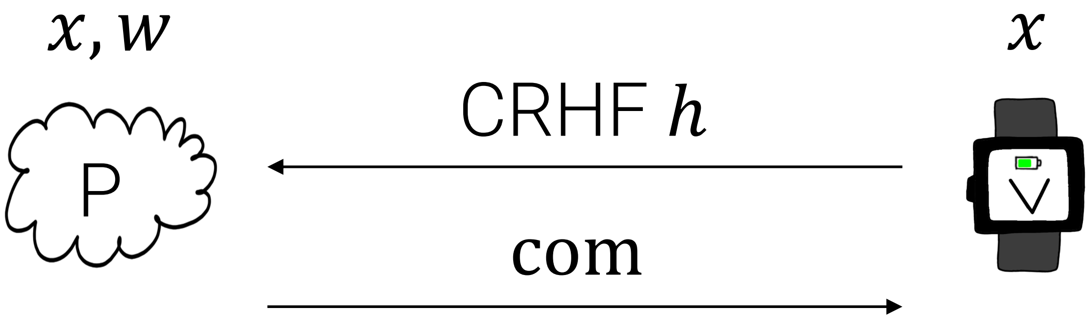

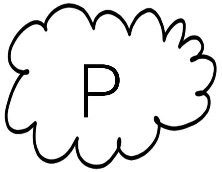

sends short commitment to PCP π.

Kilian's protocol

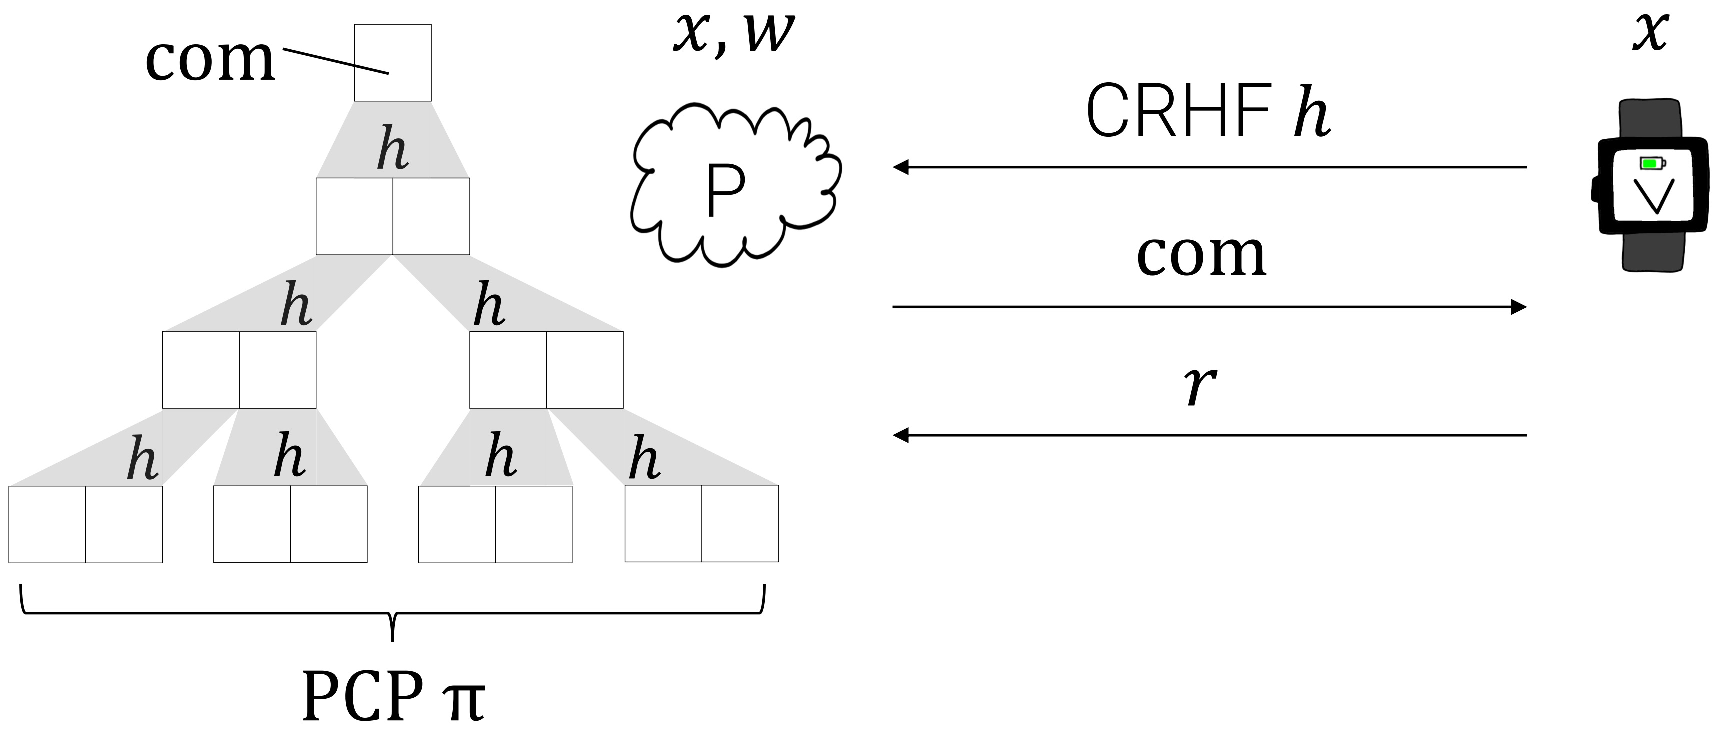

 samples PCP verifier coins $r \leftarrow R$.

Kilian's protocol

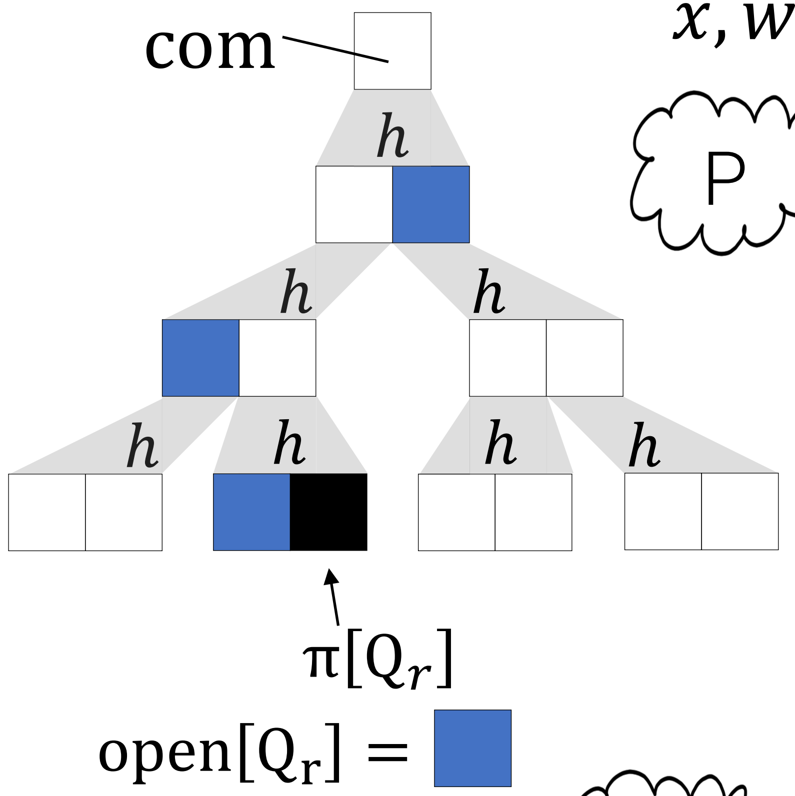

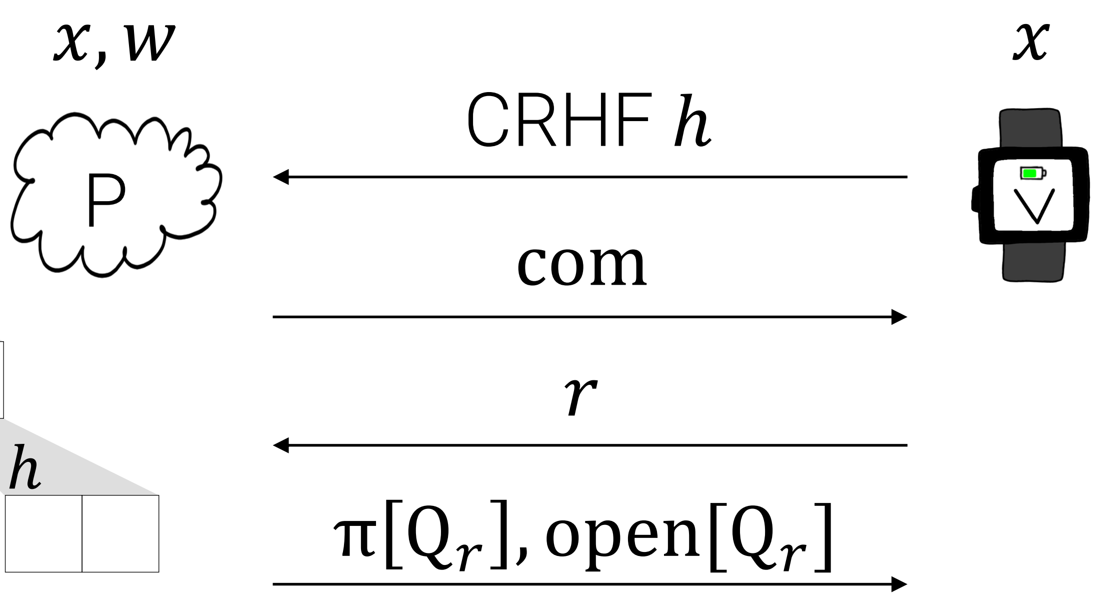

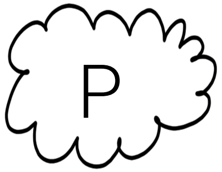

sends $\pi[Q_r]$ + opening proofs

Q_r = indices PCP verifier checks on random coins r

Kilian's protocol

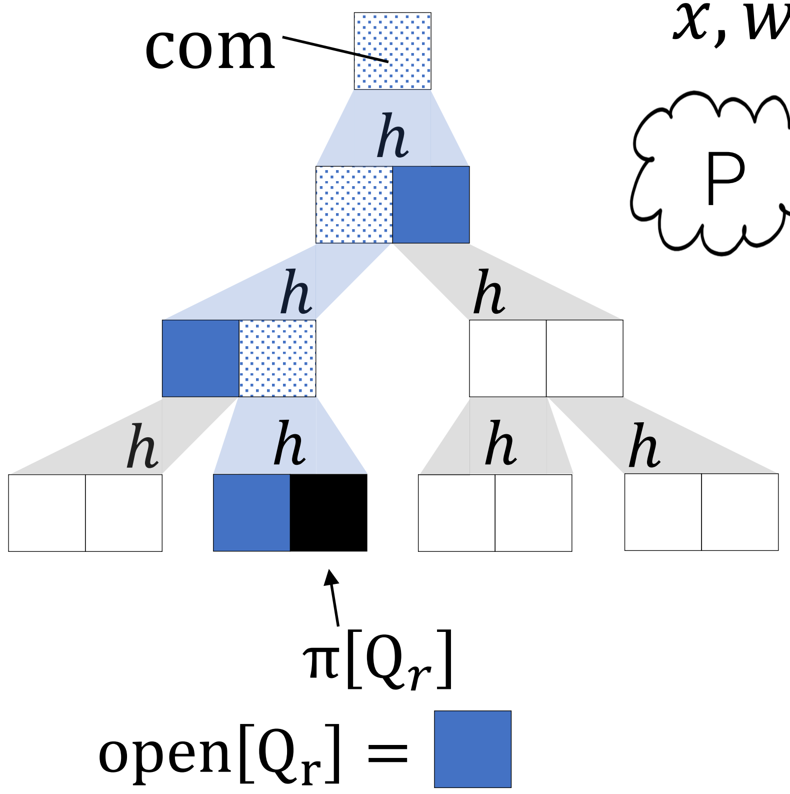

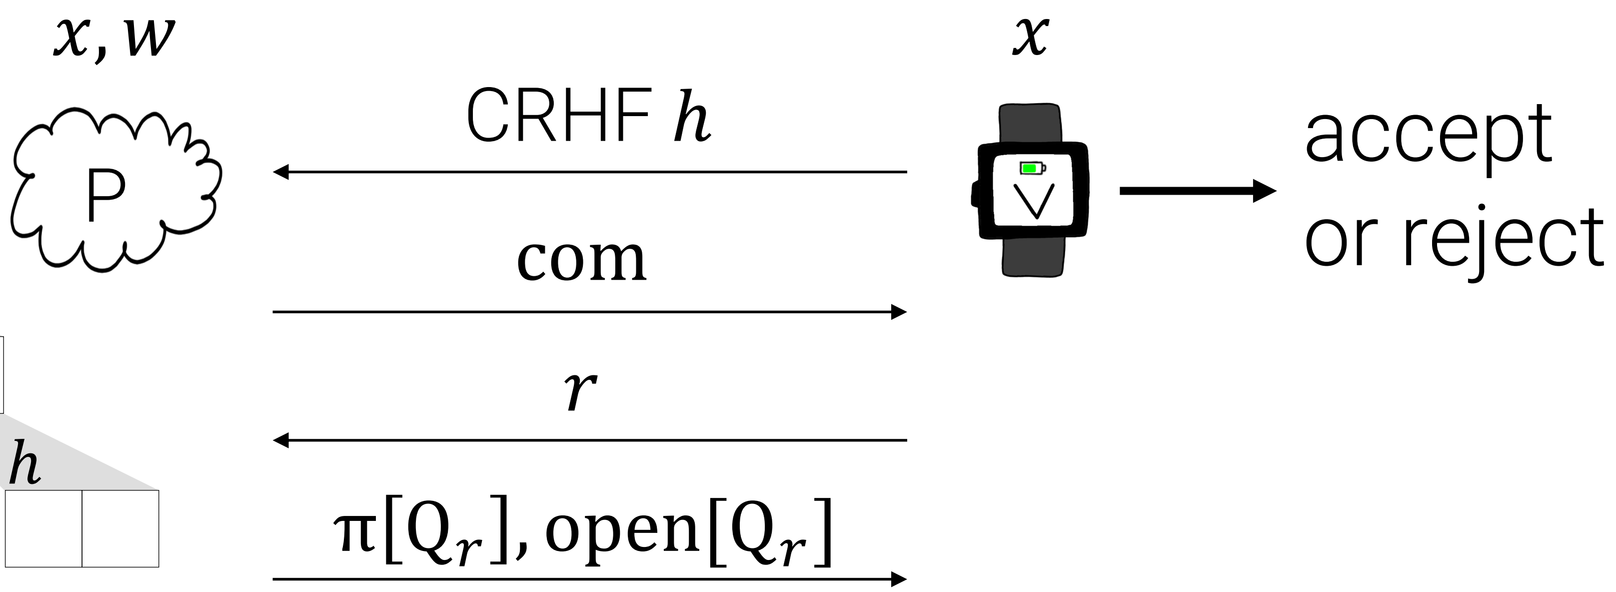

accepts if openings valid
+ PCP verifier accepts

Classical Security

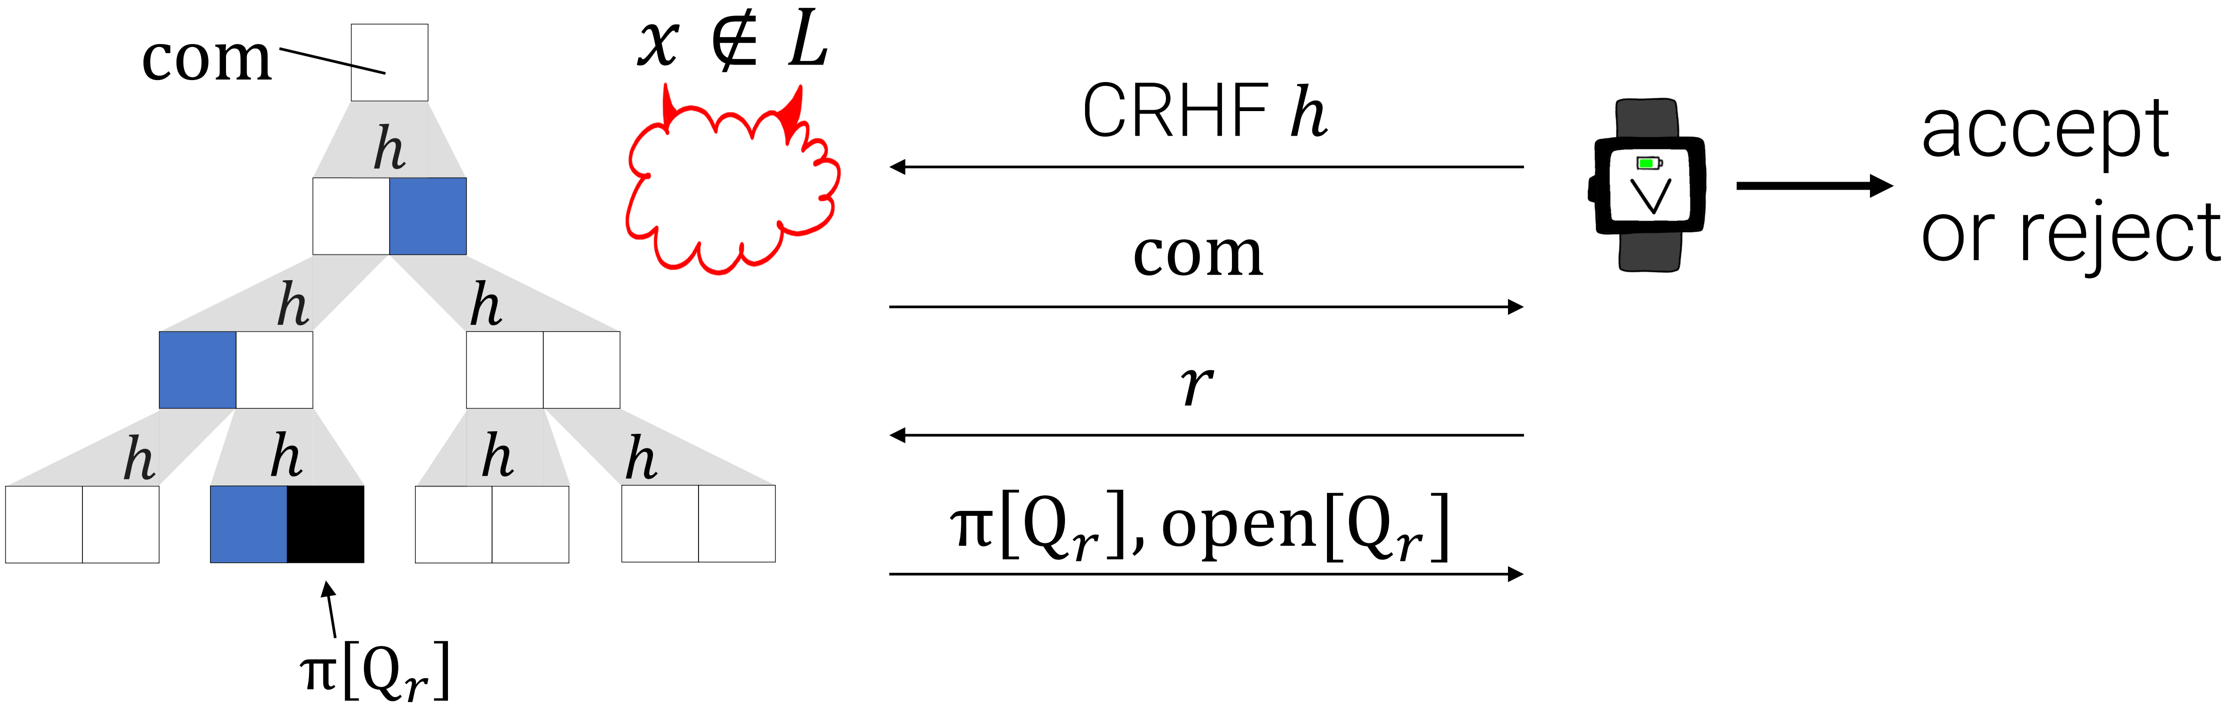

Intuition: want to show that the CRHF forces  to respond consistently with some PCP string π.

Classical Security

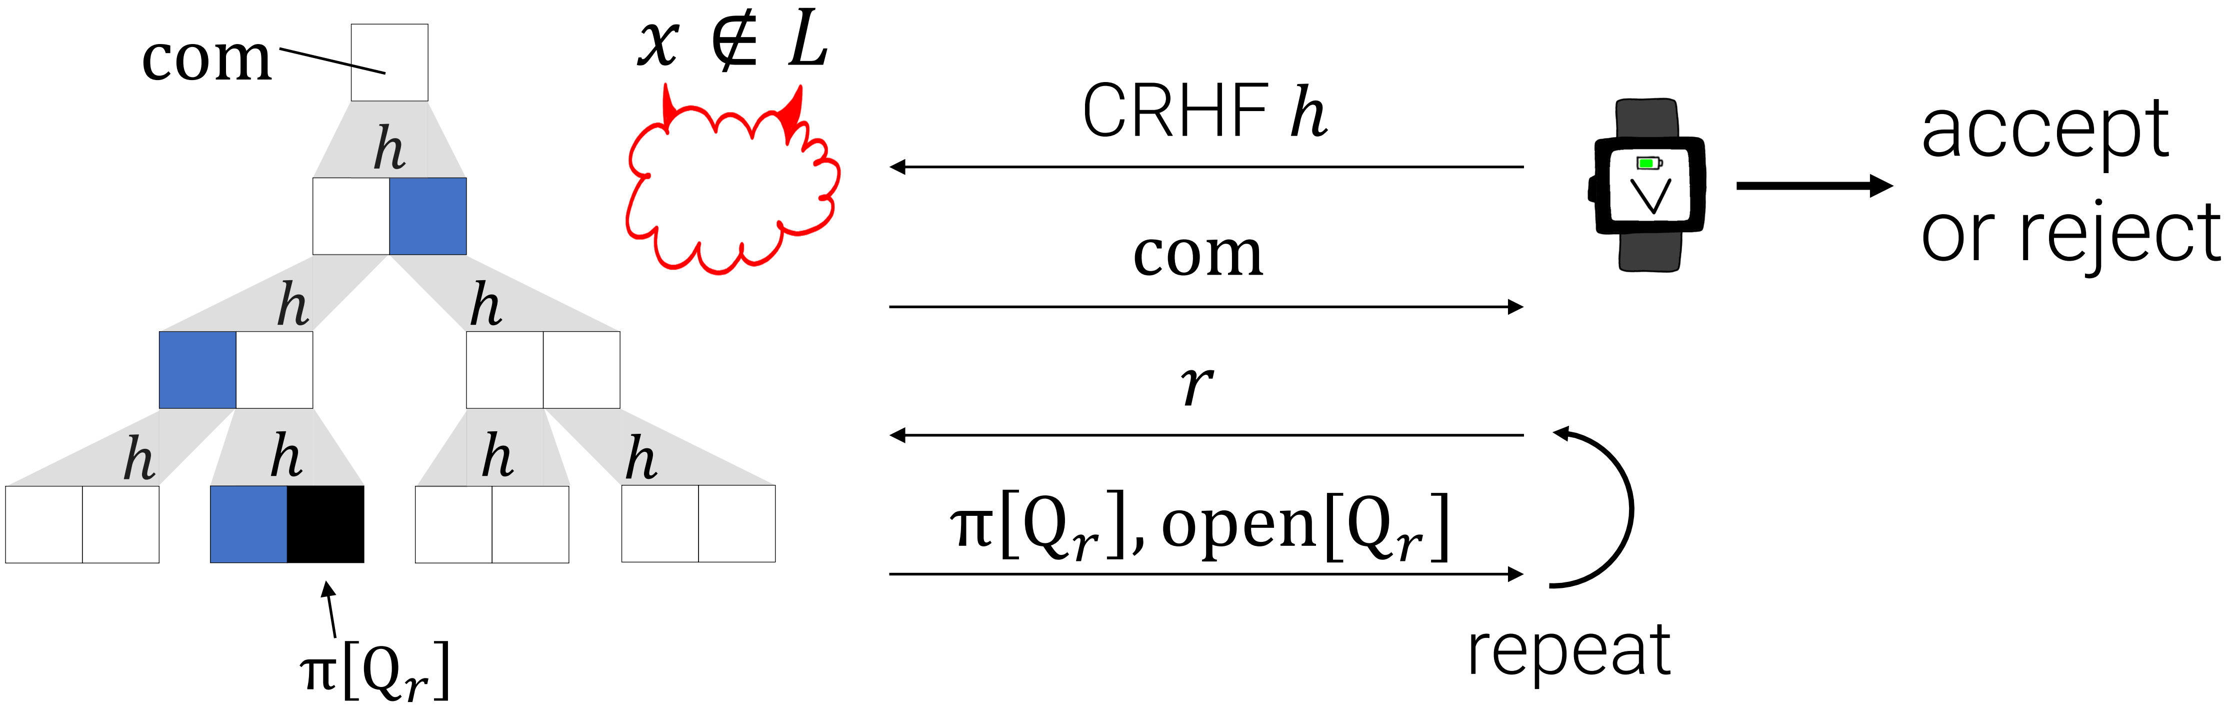

Intuition: want to show that the CRHF forces [cloud figure] to respond consistently with some PCP string π.

Formalize by ***rewinding*** last two messages many times.

Classical Security

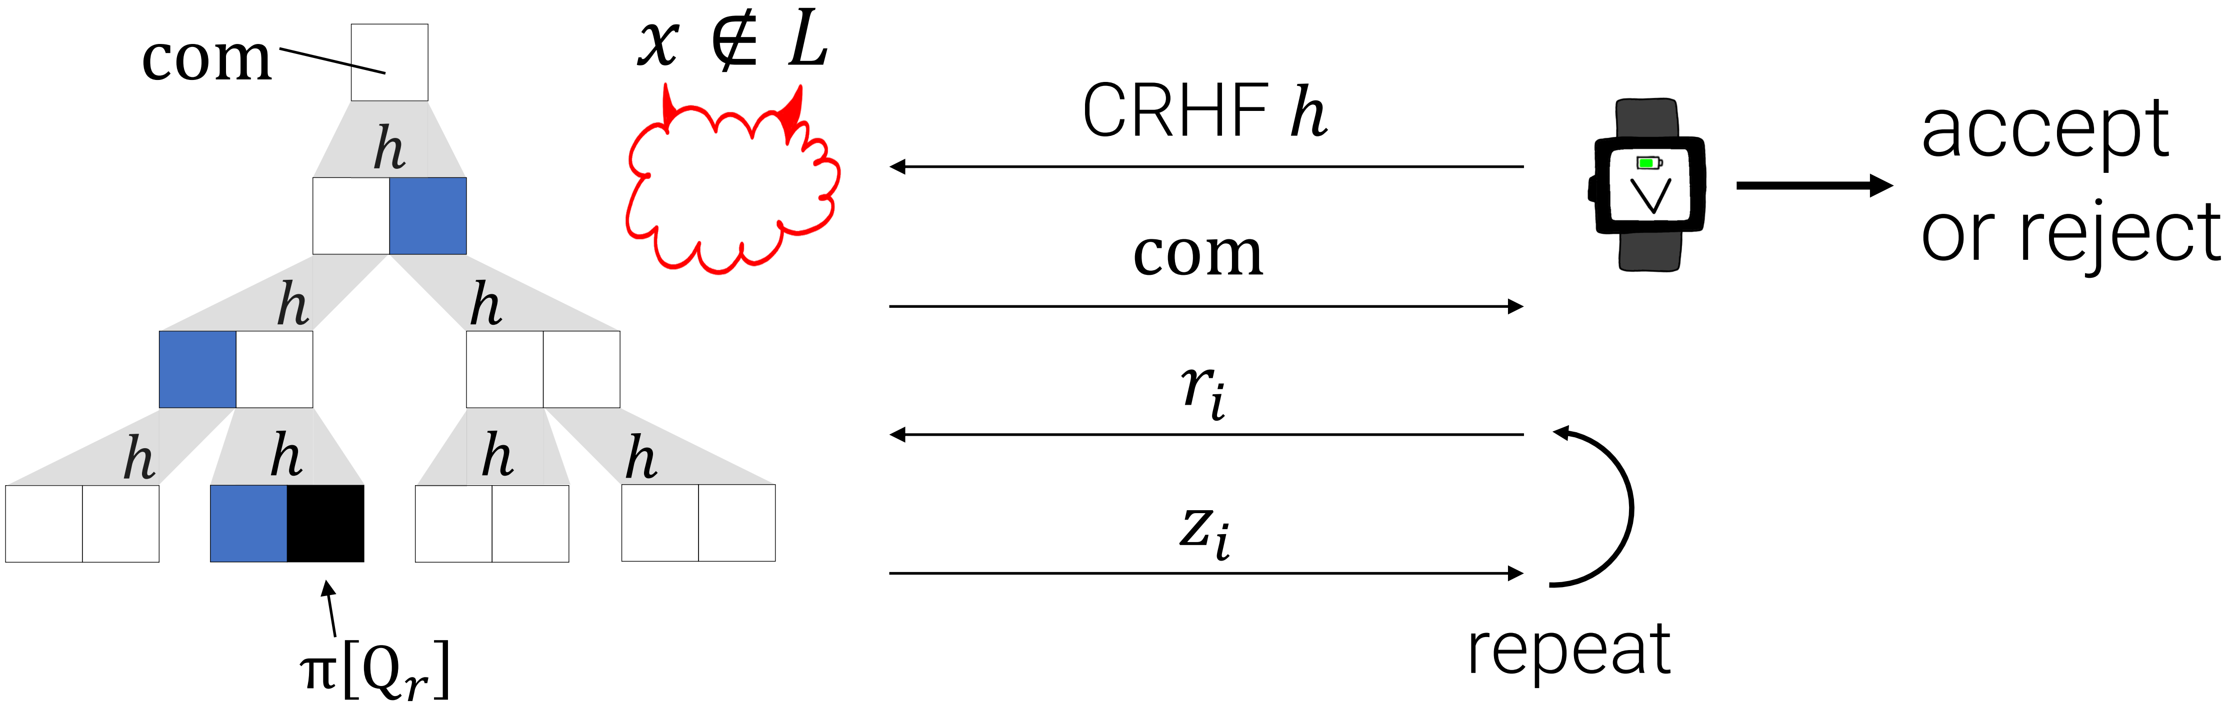

Reduction's goal: record *many* accepting transcripts (r_i, z_i)

Classical Security

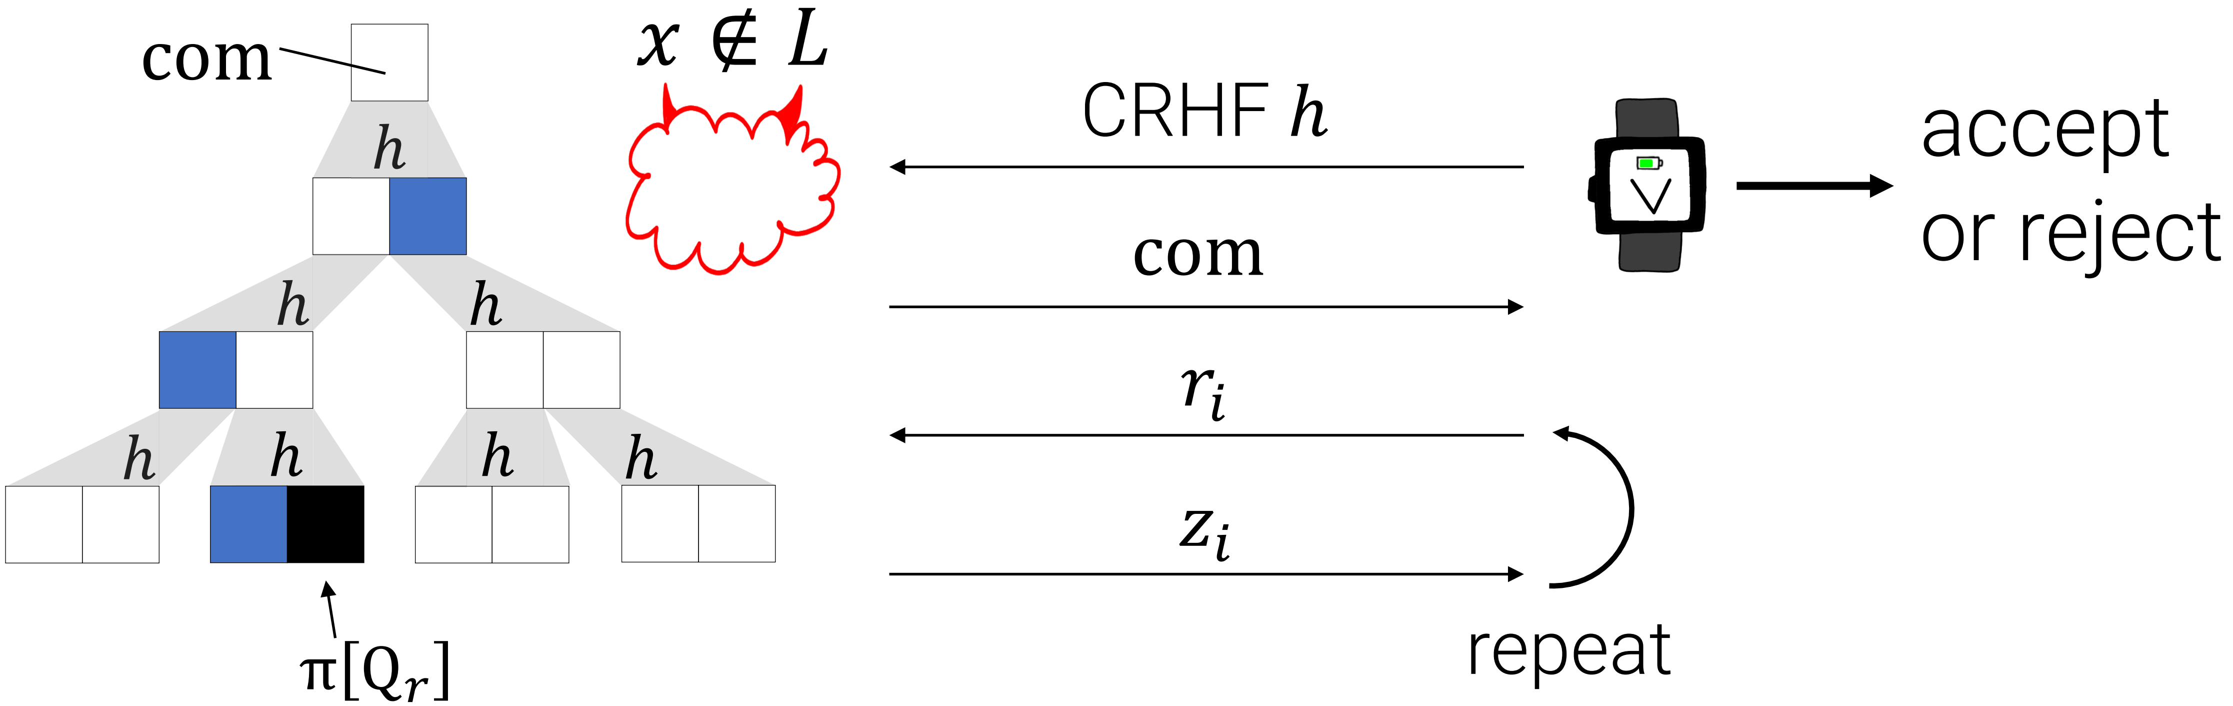

Reduction's goal: record *many* accepting transcripts (r_i, z_i)

Eventually finds impossible π OR collision.

Pr[PCP verifier accepts π] > PCP soundness error

Classical Security

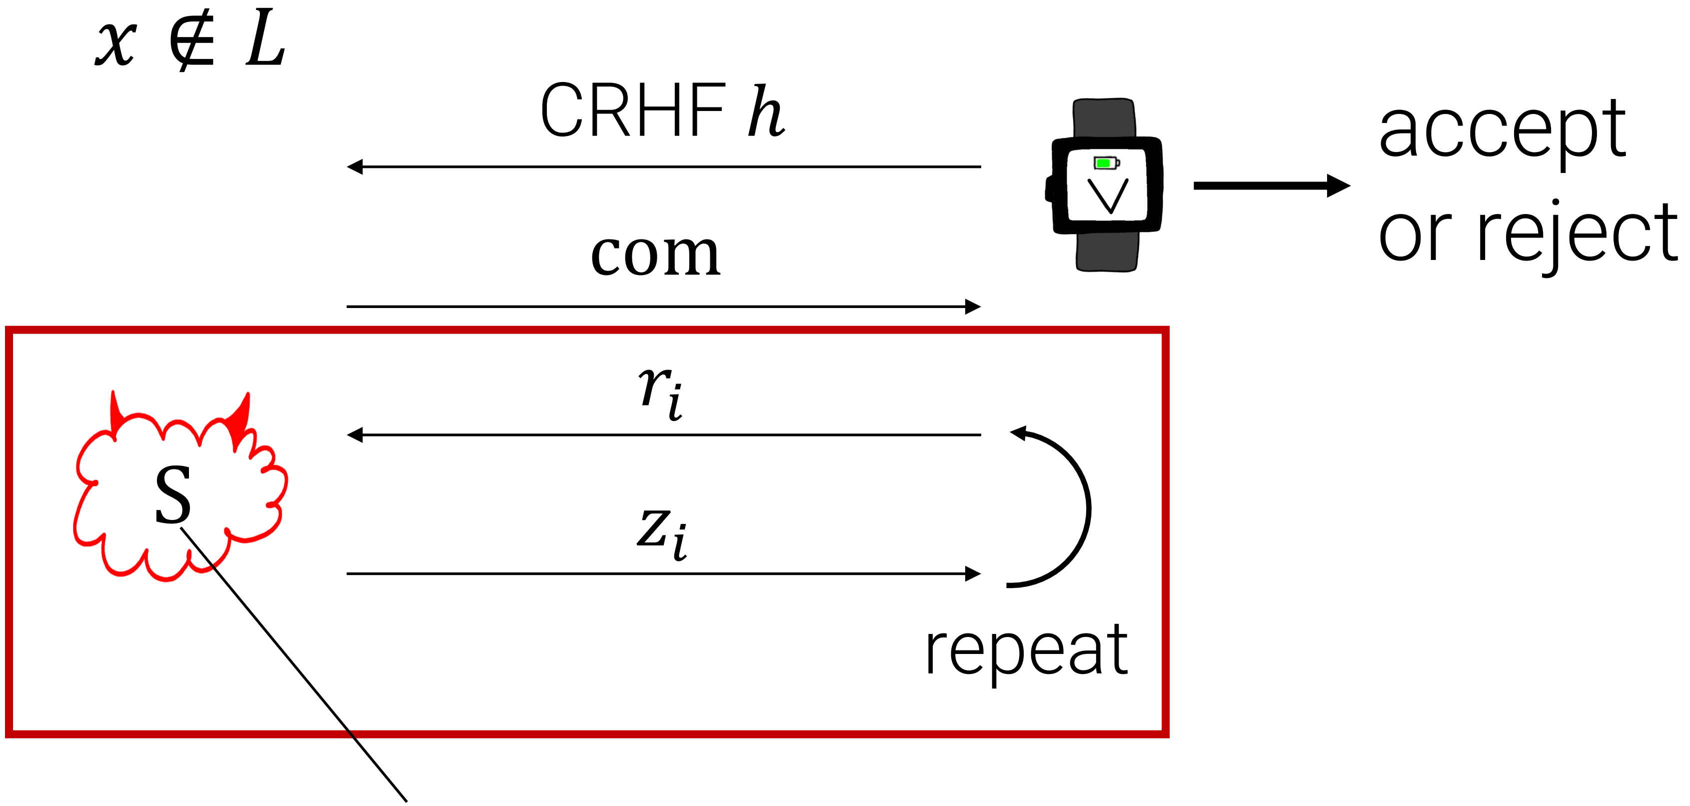

S = internal state before last two messages

rest of talk: consider "challenge-response" game

The Challenge-Response Game

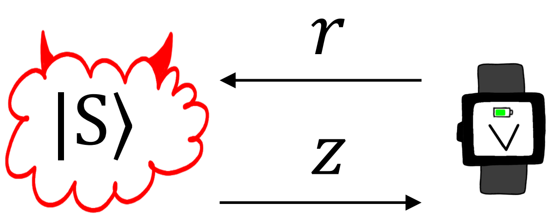

1) sample $r \leftarrow R$.
2) win if $V(r, z) = 1$.

Define ***success probability*** of $|S\rangle \coloneqq \Pr_{r \leftarrow R}[$ 

$\text{wins}]$

The Challenge-Response Game

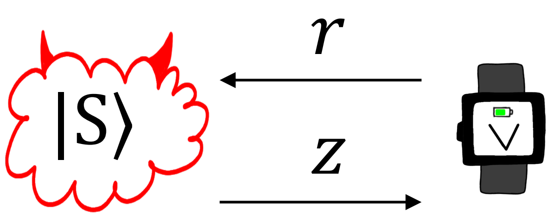

1) sample $r \leftarrow R$.
2) win if $V(r, z) = 1$.

Define *success probability* of $|S\rangle :=$ $\Pr_{r \leftarrow R}[$  $\text{wins}]$

Goal: Given $|S\rangle$ with success probability $1/\text{poly}(\lambda)$, output many accepting transcripts (r_i, z_i)

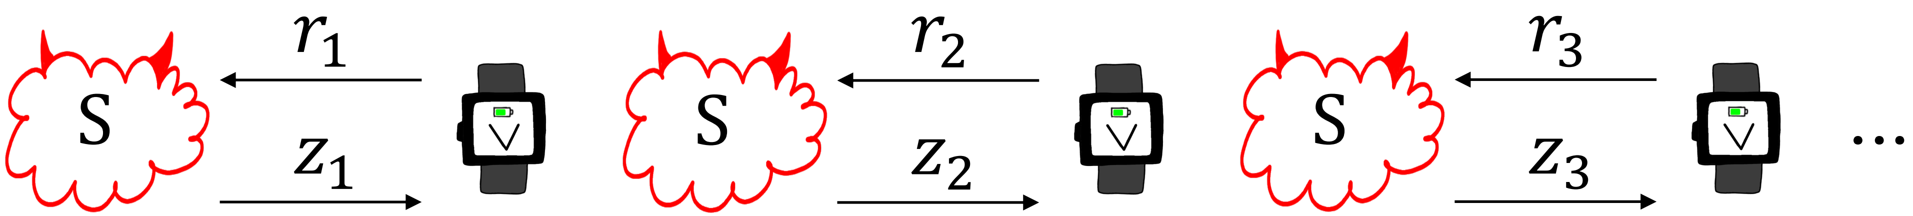

When $|S\rangle$ is classical, can run many trials by resetting the prover's state.

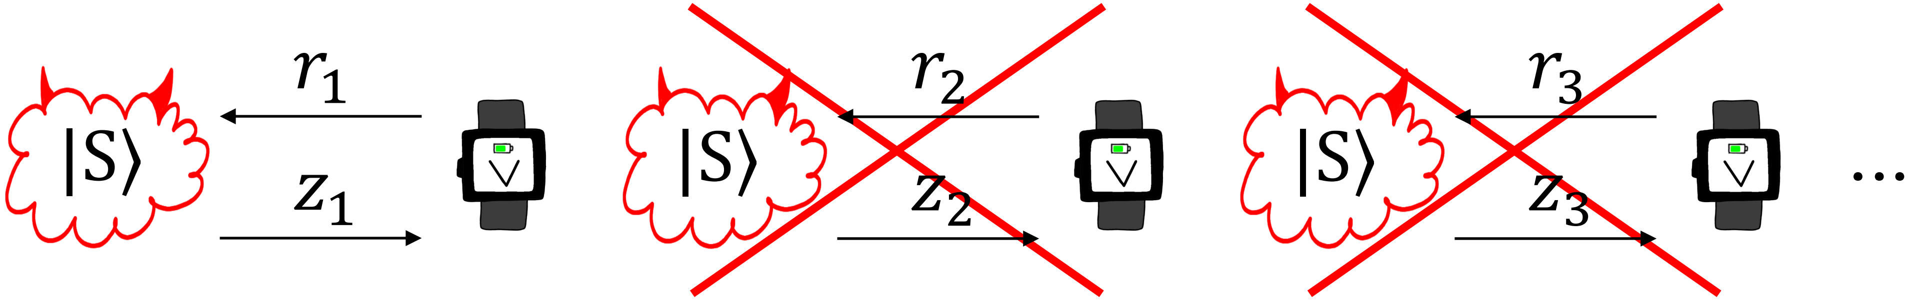

If $|S\rangle$ is quantum, we can't reset the state since a single trial requires measuring z, which disturbs $|S\rangle$.

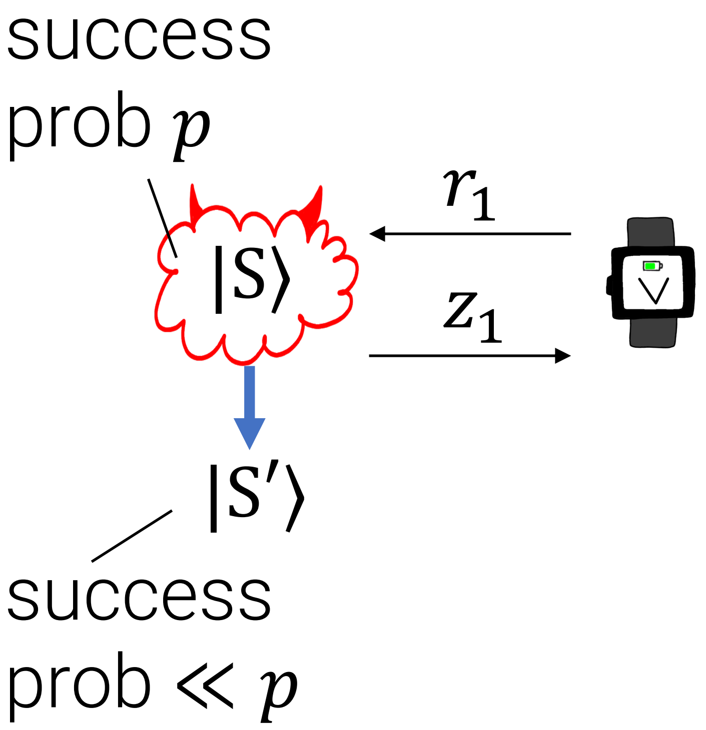

If $|S\rangle$ is quantum, we can't reset the state since a single trial requires measuring z, which disturbs $|S\rangle$.

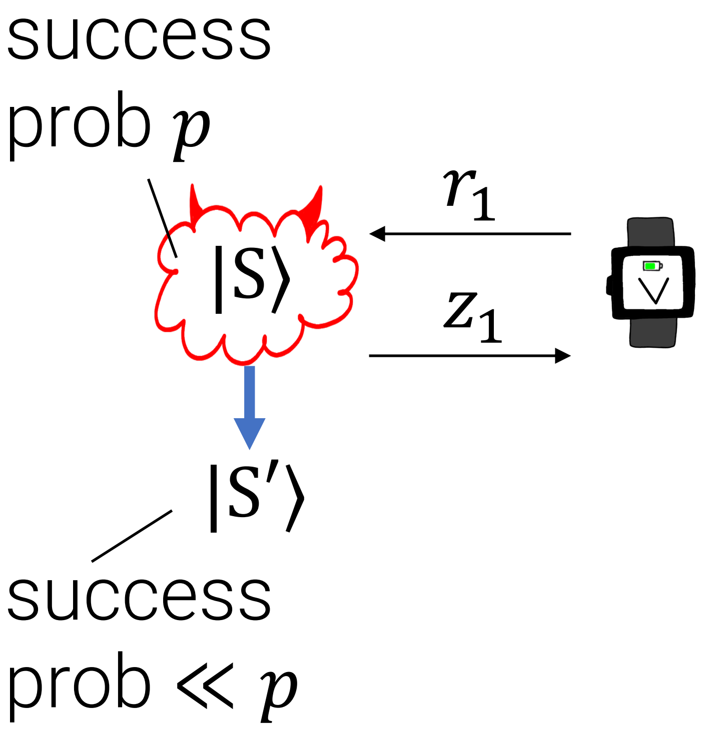

Problem: $|S'\rangle$ might not be a successful adversary!

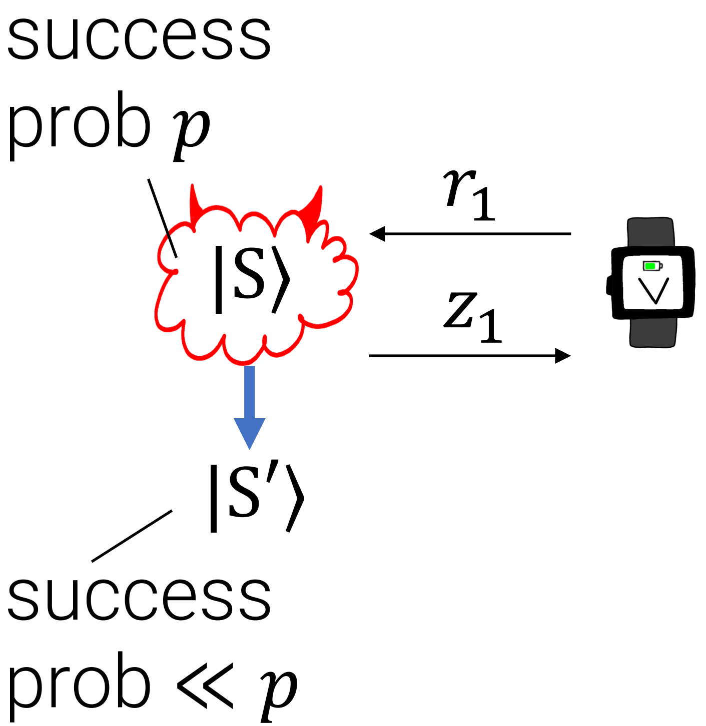

Problem: $|S'\rangle$ might not be a successful adversary!

This work: we devise a "repair" procedure to restore the original success probability.

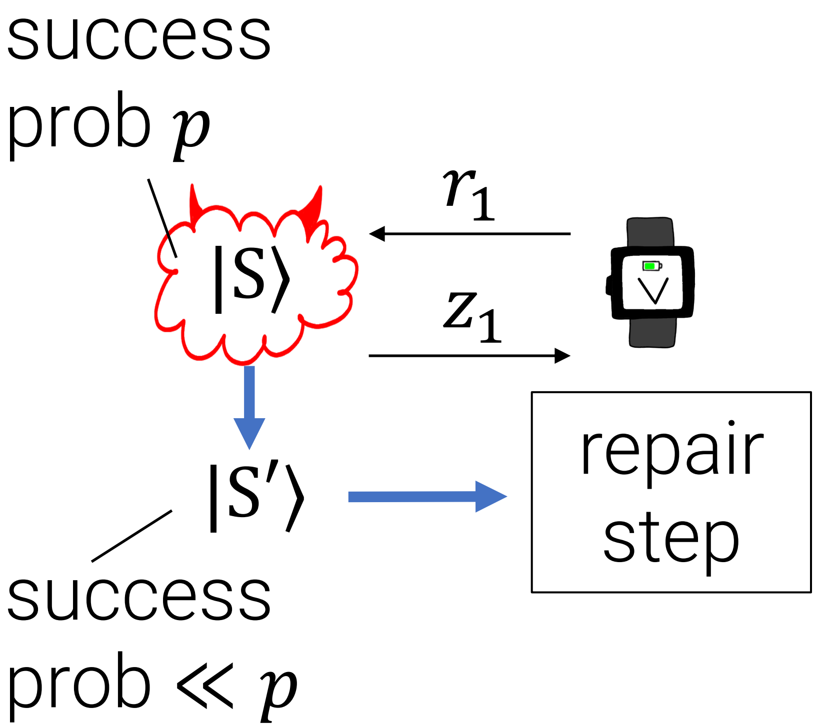

Problem: $|S'\rangle$ might not be a successful adversary!

This work: we devise a "repair" procedure to restore the original success probability.

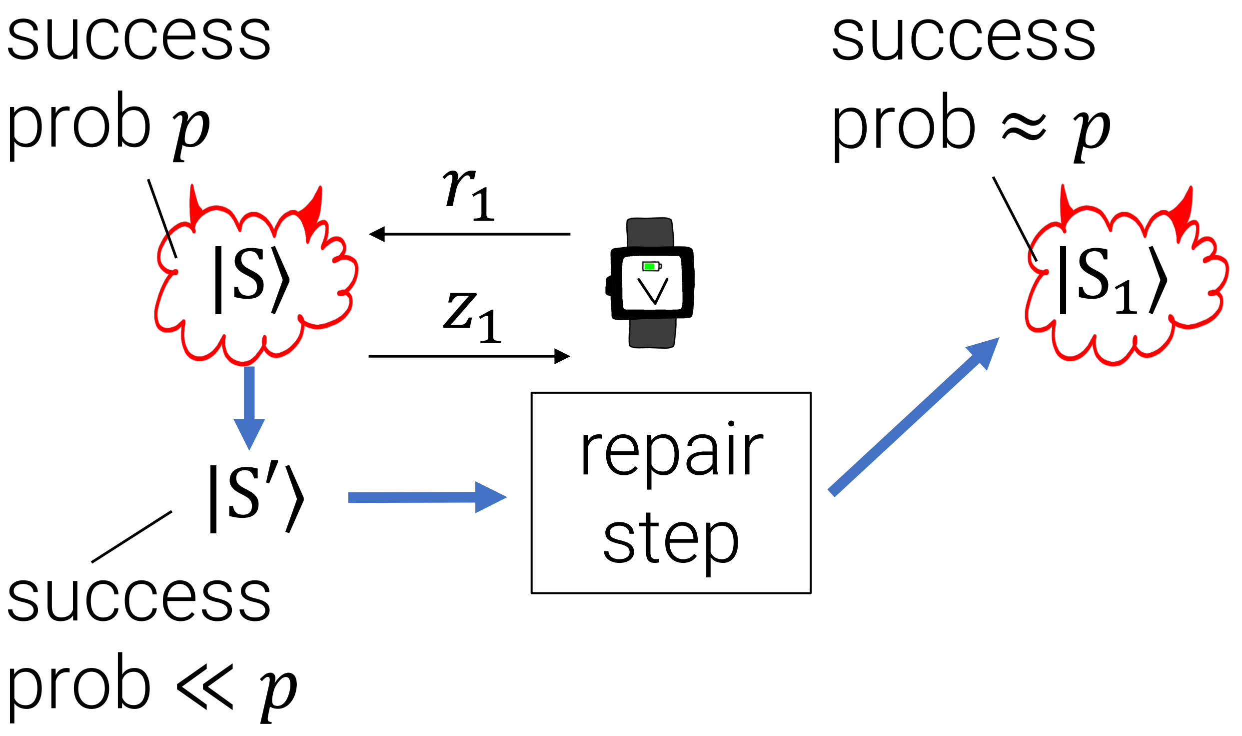

Problem: $|S'\rangle$ might not be a successful adversary!

This work: we devise a "repair" procedure to restore the original success probability.

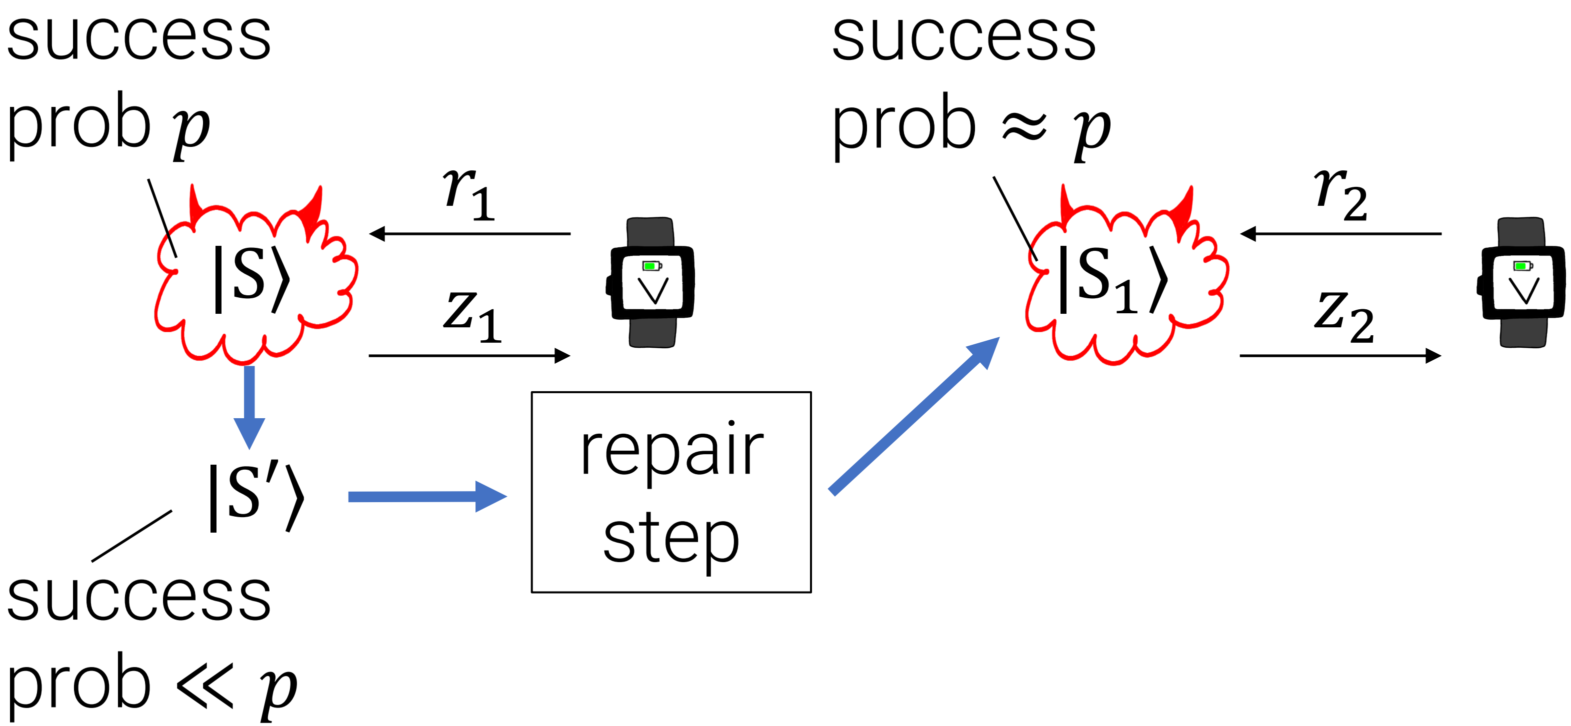

Problem: $|S'\rangle$ might not be a successful adversary!

This work: we devise a "repair" procedure to restore the original success probability.

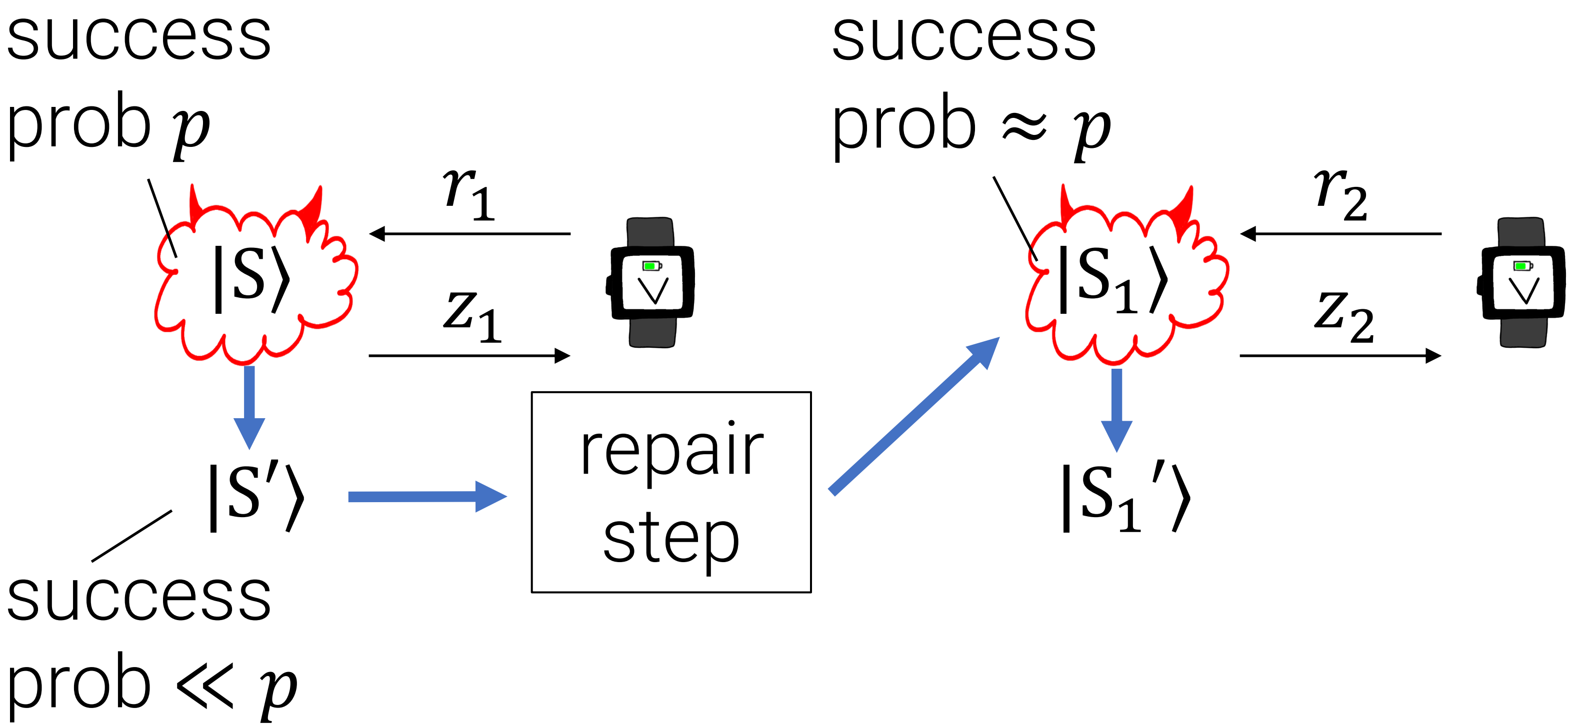

Problem: $|S'\rangle$ might not be a successful adversary!

This work: we devise a “repair” procedure to restore the original success probability.

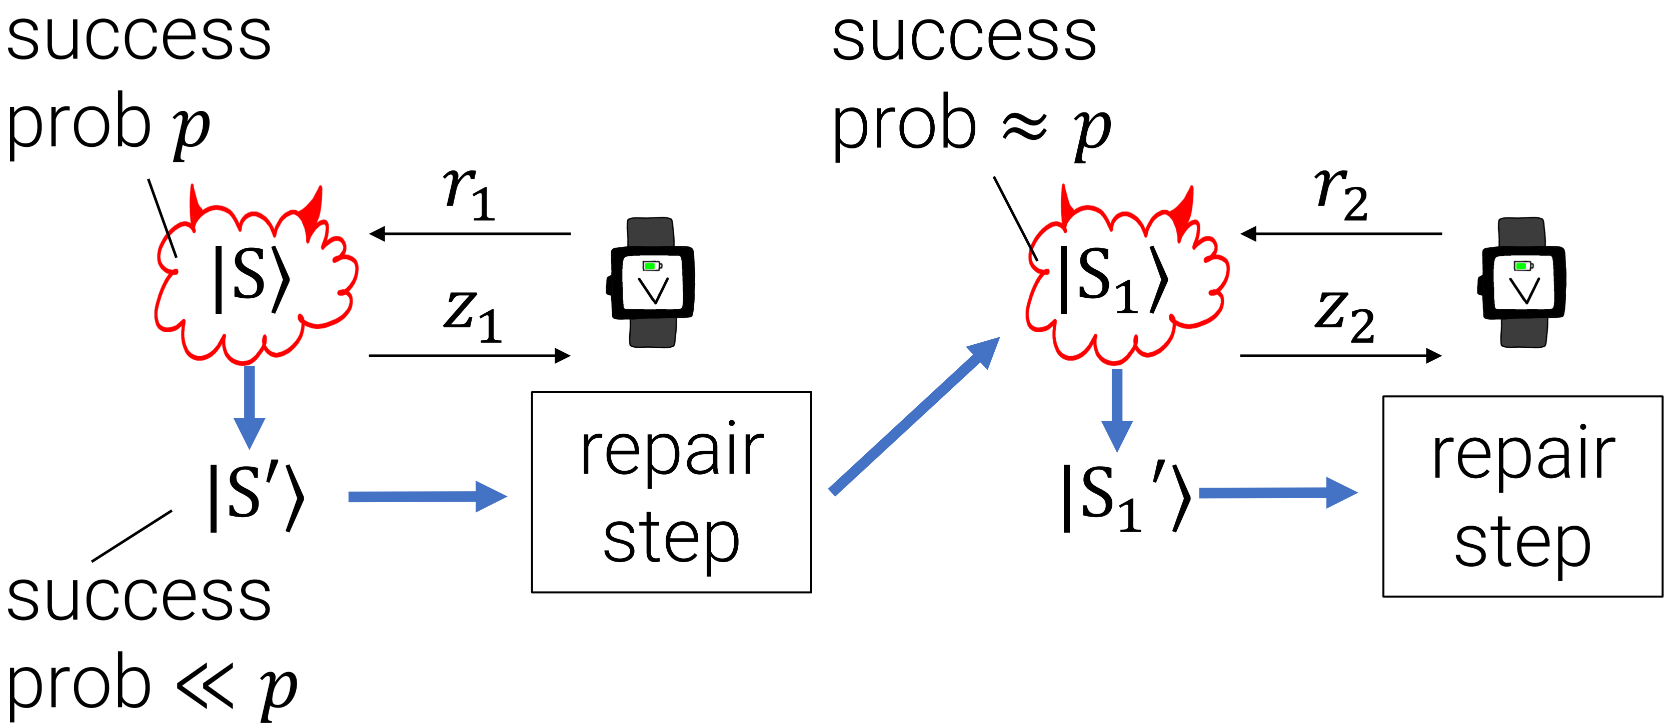

Problem: $|S'\rangle$ might not be a successful adversary!

This work: we devise a "repair" procedure to restore the original success probability.

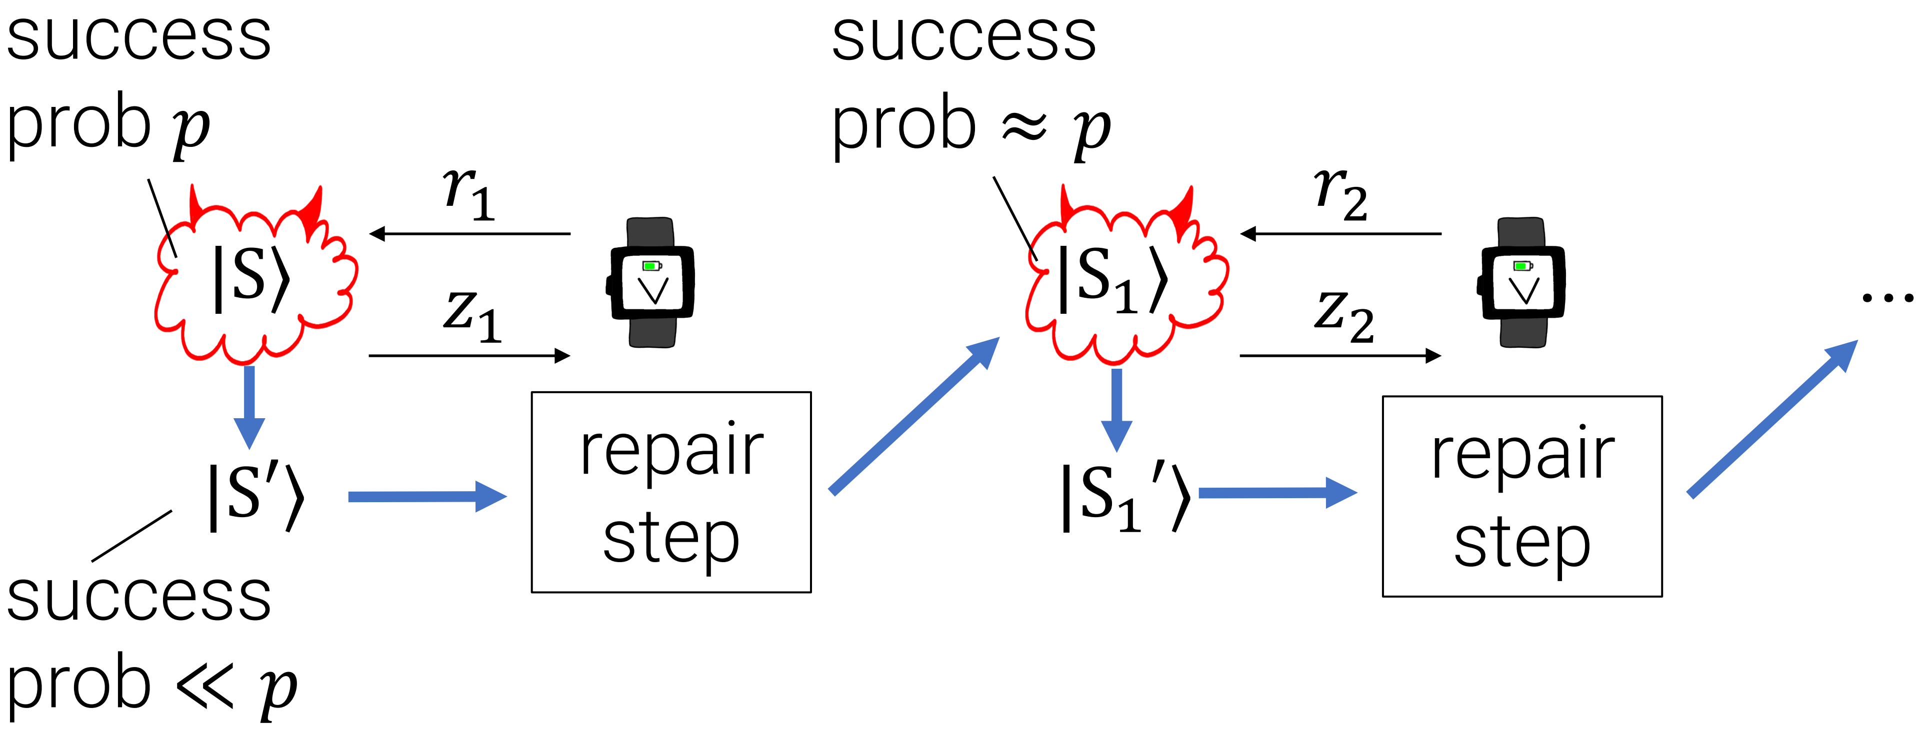

Problem: $|S'\rangle$ might not be a successful adversary!

This work: we devise a "repair" procedure to restore the original success probability.

First, we'll need to recall a technique of [Unruh12] to reduce *measuring the prover's response* to *measuring the verifier's decision*.

Recording the Verifier's Decision [Unruh12]

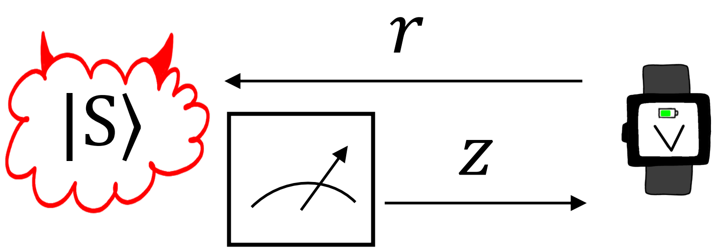

Naïve Measurement:

Measure $\sum|z\rangle$ right away.

Recording the Verifier's Decision [Unruh12]

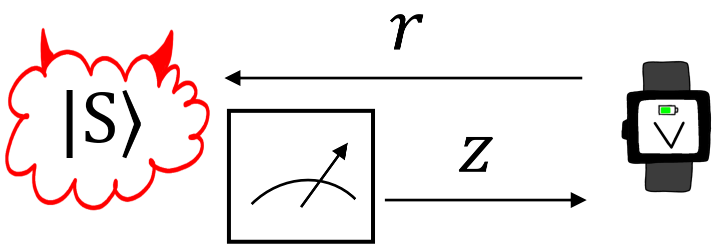

Naïve Measurement:

Measure $\sum|z\rangle$ right away.

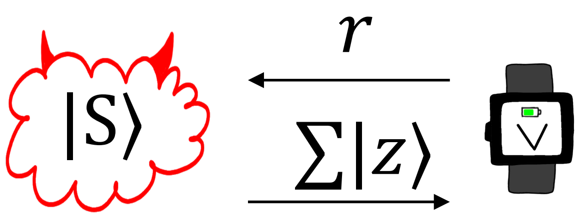

"Lazy" Measurement:

(1) Compute + measure $V(r,z)$.

(2) Measure z if $V(r,z) = 1$.

Recording the Verifier's Decision [Unruh12]

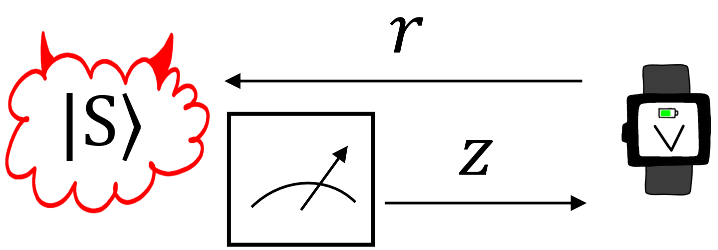

Naïve Measurement:

Measure $\sum|z\rangle$ right away.

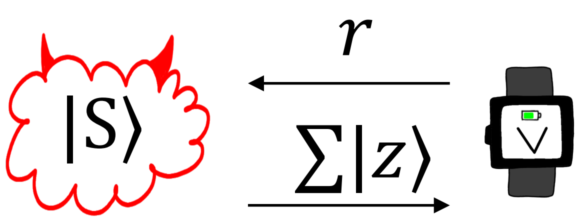

"Lazy" Measurement:

(1) Compute + measure $V(r,z)$.
(2) Measure z if $V(r,z) = 1$.

[U12]: For protocols with ***unique responses***, measurement in step (2) causes *no disturbance*!

Recording the Verifier's Decision [Unruh12]

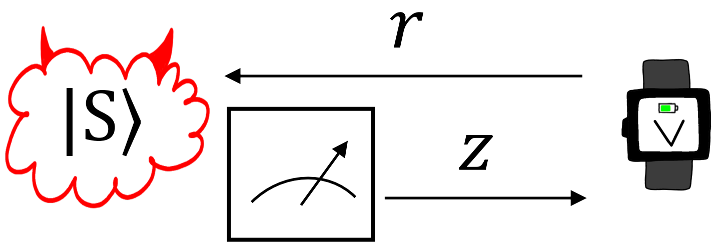

Naïve Measurement:

Measure $\sum|z\rangle$ right away.

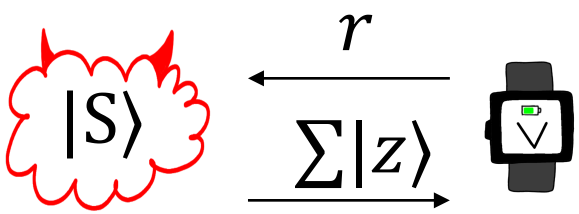

"Lazy" Measurement:

(1) Compute + measure $V(r, z)$.
(2) Measure z if $V(r, z) = 1$.

[U12]: For protocols with ***unique responses***, measurement in step (2) causes *no disturbance*!

- Kilian's protocol doesn't have this property.

Recording the Verifier's Decision [Unruh12]

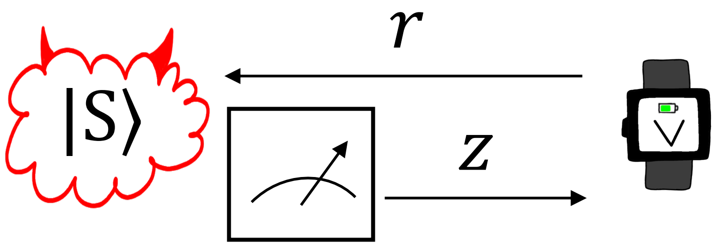

Naïve Measurement:

Measure $\sum|z\rangle$ right away.

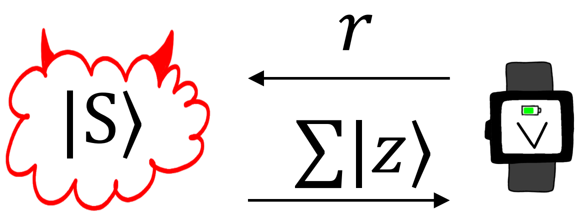

"Lazy" Measurement:

(1) Compute + measure $V(r, z)$.
(2) Measure z if $V(r, z) = 1$.

[U12]: For protocols with ***unique responses***, measurement in step (2) causes *no disturbance*!

- Kilian's protocol doesn't have this property.
- However, if the CRHF h is *quantum-binding* (collapsing [U16]), then step (2) is *computationally undetectable*.

Recording the Verifier's Decision [Unruh12]

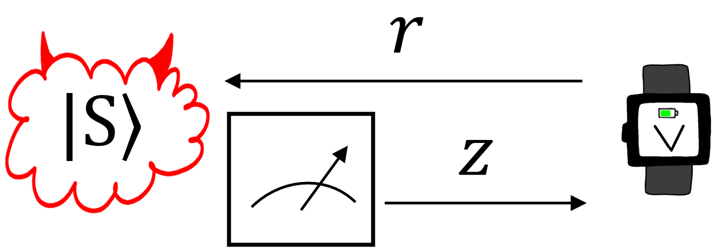

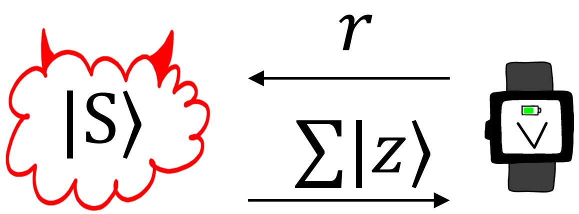

Naïve Measurement:

Measure $\sum|z\rangle$ right away.

"Lazy" Measurement:

(1) Compute + measure $V(r,z)$.

(2) Measure z if $V(r,z) = 1$.

It therefore *suffices* to only perform step (1) and simply try to make the verifier accept on many random challenges.

Recording the Verifier's Decision [Unruh12]

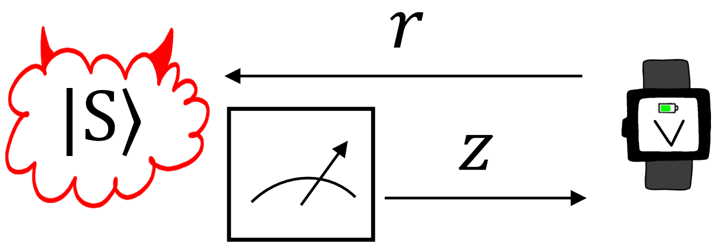

Naïve Measurement:

Measure $\sum|z\rangle$ right away.

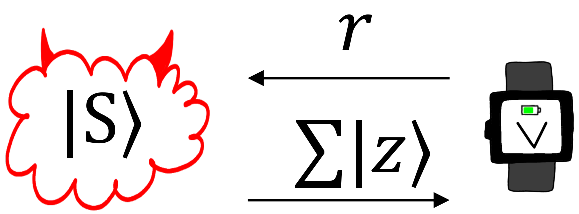

"Lazy" Measurement:

(1) Compute + measure $V(r,z)$.
(2) Measure z if $V(r,z) = 1$.

It therefore *suffices* to only perform step (1) and simply try to make the verifier accept on many random challenges.

This will imply a full reduction that performs step (1) and (2), since (2) is computationally undetectable.

Takeaway: can just measure the verifier’s decision, so we only have to “repair” one-bit disturbance.

With this in mind, let’s turn to state repair.

State Repair Intuition: Alternating Projections

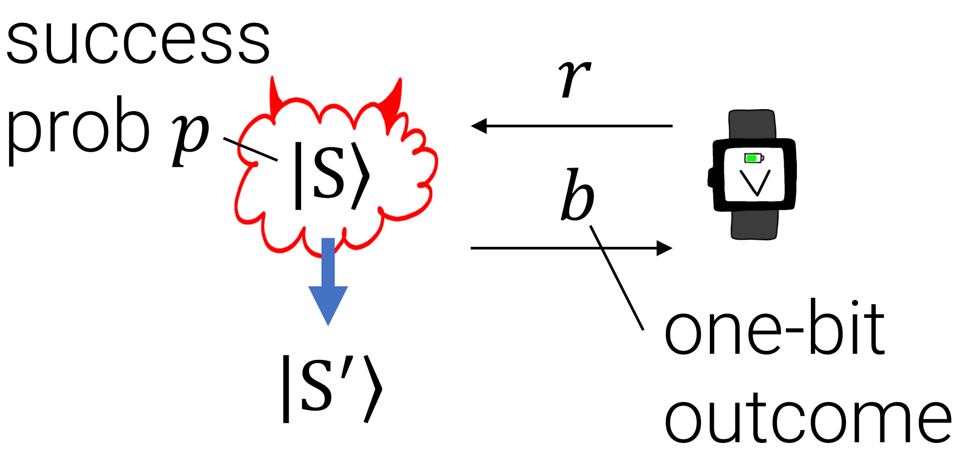

State Repair Intuition: Alternating Projections

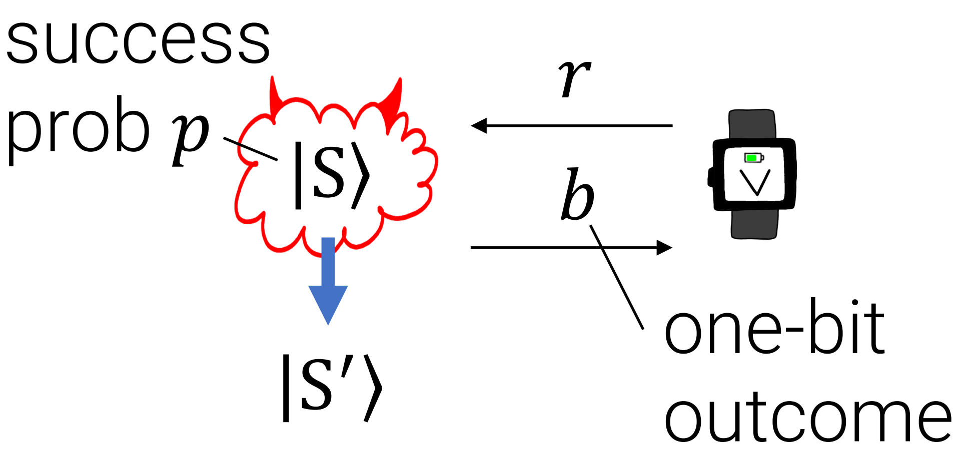

State Repair (High-Level Idea)

1) Identify a "good subspace" Π_p where $|S\rangle \in \Pi_p$, and moreover *all states* in Π_p have success prob $\geq p$.

State Repair Intuition: Alternating Projections

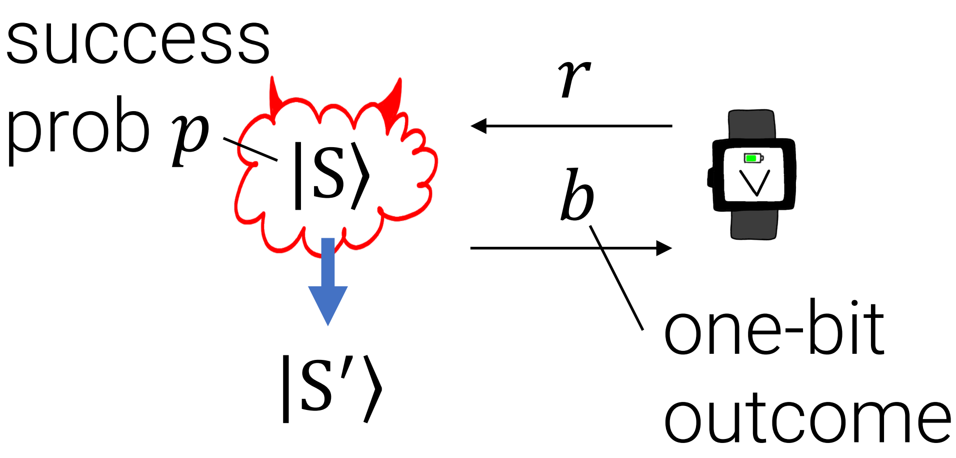

State Repair (High-Level Idea)

1) Identify a "good subspace" Π_p where $|S\rangle \in \Pi_p$, and moreover *all states* in Π_p have success prob $\geq p$.
2) Alternate two projective measurements until the state is in Π_p:

State Repair Intuition: Alternating Projections

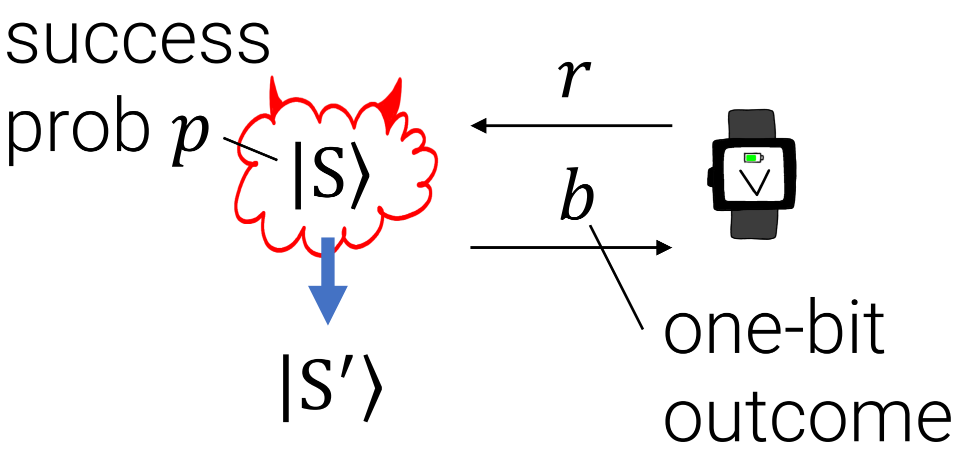

State Repair (High-Level Idea)

1) Identify a "good subspace" Π_p where $|S\rangle \in \Pi_p$, and moreover *all states* in Π_p have success prob $\geq p$.
2) Alternate two projective measurements until the state is in Π_p:
   - the binary measurement $(\Pi_r, \mathbb{I} - \Pi_r)$ that disturbed $|S\rangle$

State Repair Intuition: Alternating Projections

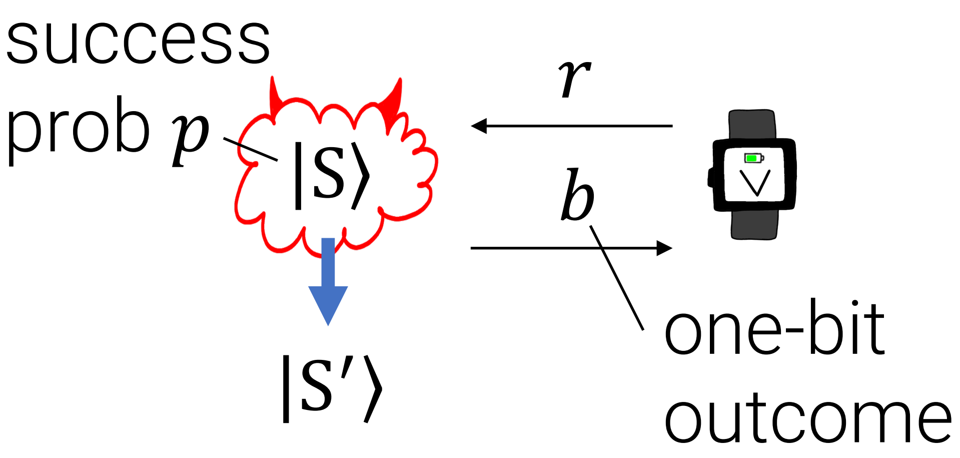

State Repair (High-Level Idea)

1) Identify a "good subspace" Π_p where $|S\rangle \in \Pi_p$, and moreover *all states* in Π_p have success prob $\geq p$.
2) Alternate two projective measurements until the state is in Π_p:
   - the binary measurement $(\Pi_r, \mathbb{I} - \Pi_r)$ that disturbed $|S\rangle$
   - the binary measurement $(\Pi_p, \mathbb{I} - \Pi_p)$

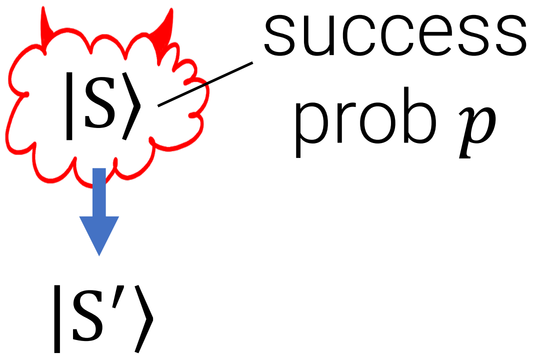

State Repair (High-Level Idea)

1) Identify a "good subspace" Π_p where $|S\rangle \in \Pi_p$, and moreover *all states* in Π_p have success prob $\geq p$.
2) Alternate two projective measurements until the state is in Π_p:
   - the binary measurement $(\Pi_r, \mathbb{I} - \Pi_r)$ that disturbed $|S\rangle$
   - the binary measurement $(\Pi_p, \mathbb{I} - \Pi_p)$

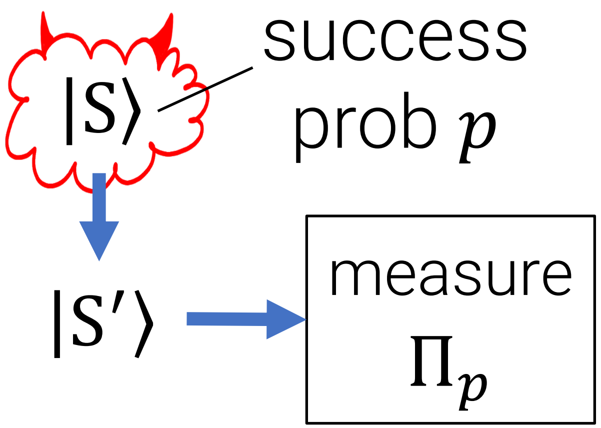

State Repair (High-Level Idea)

1) Identify a "good subspace" Π_p where $|S\rangle \in \Pi_p$, and moreover *all states* in Π_p have success prob $\geq p$.
2) Alternate two projective measurements until the state is in Π_p:
   - the binary measurement $(\Pi_r, \mathbb{I} - \Pi_r)$ that disturbed $|S\rangle$
   - the binary measurement $(\Pi_p, \mathbb{I} - \Pi_p)$

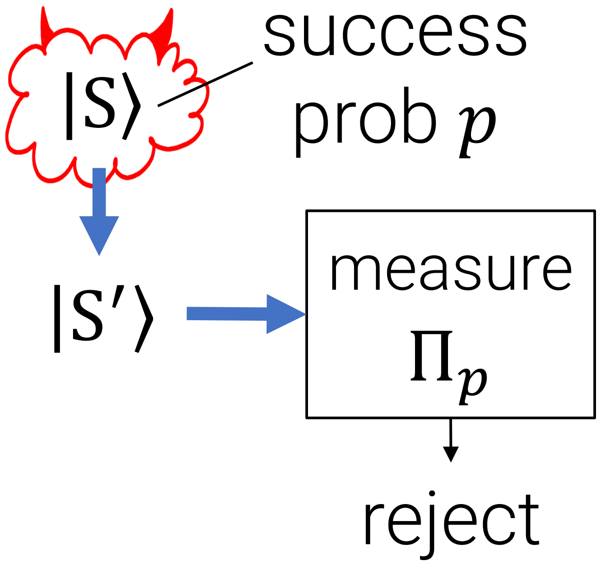

State Repair (High-Level Idea)

1) Identify a “good subspace” Π_p where $|S\rangle \in \Pi_p$, and moreover *all states* in Π_p have success prob $\geq p$.
2) Alternate two projective measurements until the state is in Π_p:
   - the binary measurement $(\Pi_r, \mathbb{I} - \Pi_r)$ that disturbed $|S\rangle$
   - the binary measurement $(\Pi_p, \mathbb{I} - \Pi_p)$

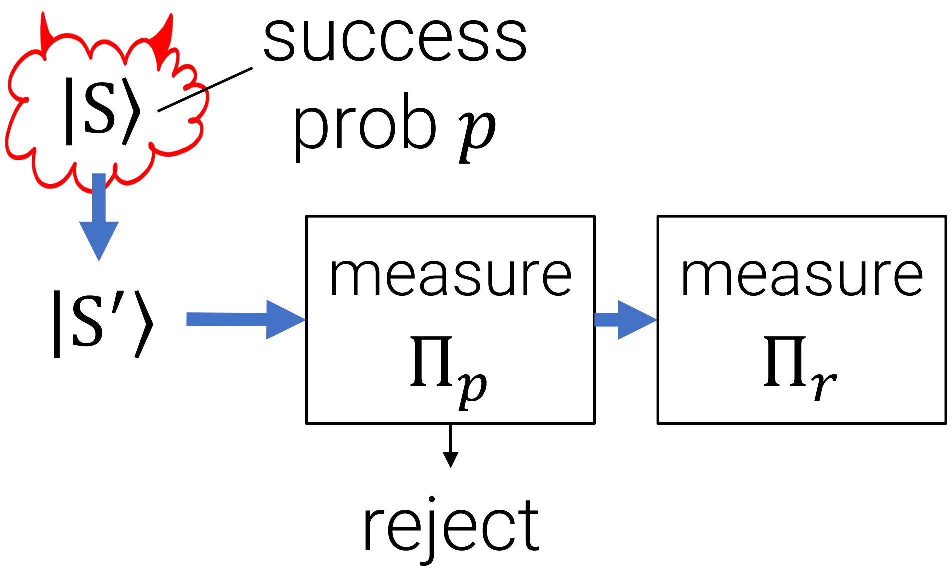

State Repair (High-Level Idea)

1) Identify a "good subspace" Π_p where $|S\rangle \in \Pi_p$, and moreover *all states* in Π_p have success prob $\geq p$.
2) Alternate two projective measurements until the state is in Π_p:
   - the binary measurement $(\Pi_r, \mathbb{I} - \Pi_r)$ that disturbed $|S\rangle$
   - the binary measurement $(\Pi_p, \mathbb{I} - \Pi_p)$

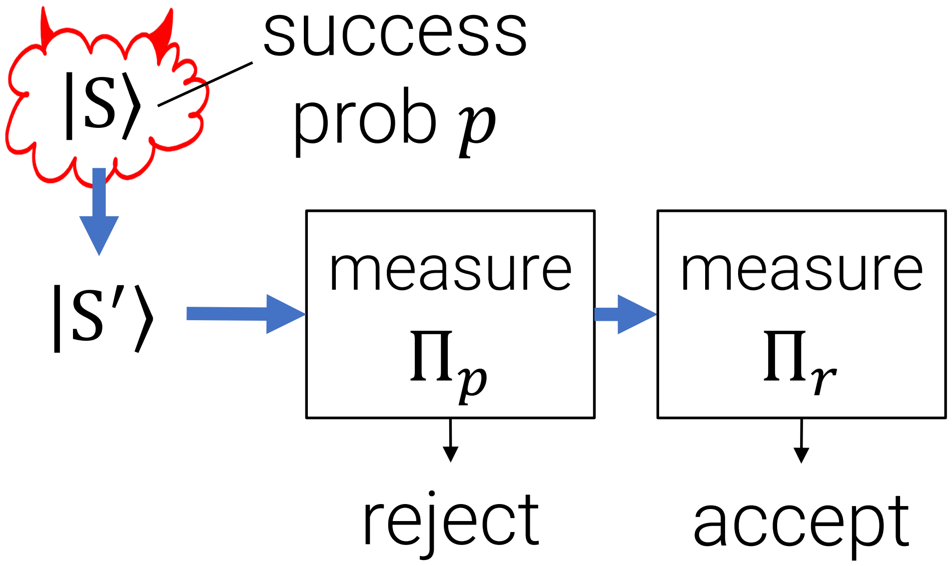

State Repair (High-Level Idea)

1) Identify a “good subspace” Π_p where $|S\rangle \in \Pi_p$, and moreover *all states* in Π_p have success prob $\geq p$.
2) Alternate two projective measurements until the state is in Π_p:
   - the binary measurement $(\Pi_r, \mathbb{I} - \Pi_r)$ that disturbed $|S\rangle$
   - the binary measurement $(\Pi_p, \mathbb{I} - \Pi_p)$

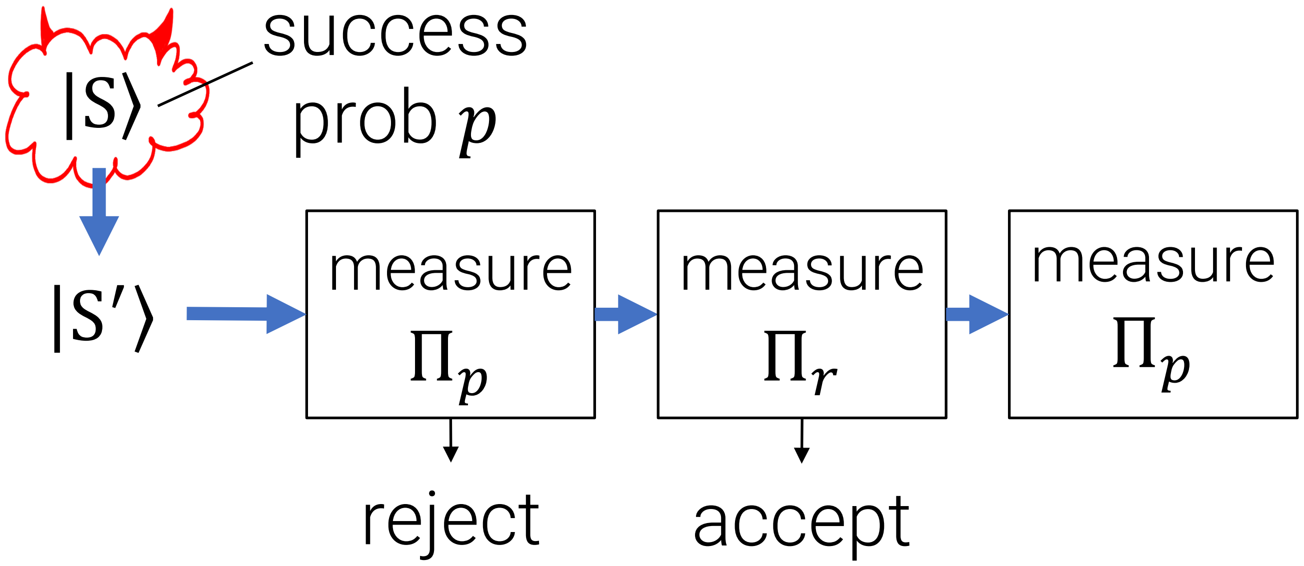

State Repair (High-Level Idea)

1) Identify a "good subspace" Π_p where $|S\rangle \in \Pi_p$, and moreover *all states* in Π_p have success prob $\geq p$.
2) Alternate two projective measurements until the state is in Π_p:
   - the binary measurement $(\Pi_r, \mathbb{I} - \Pi_r)$ that disturbed $|S\rangle$
   - the binary measurement $(\Pi_p, \mathbb{I} - \Pi_p)$

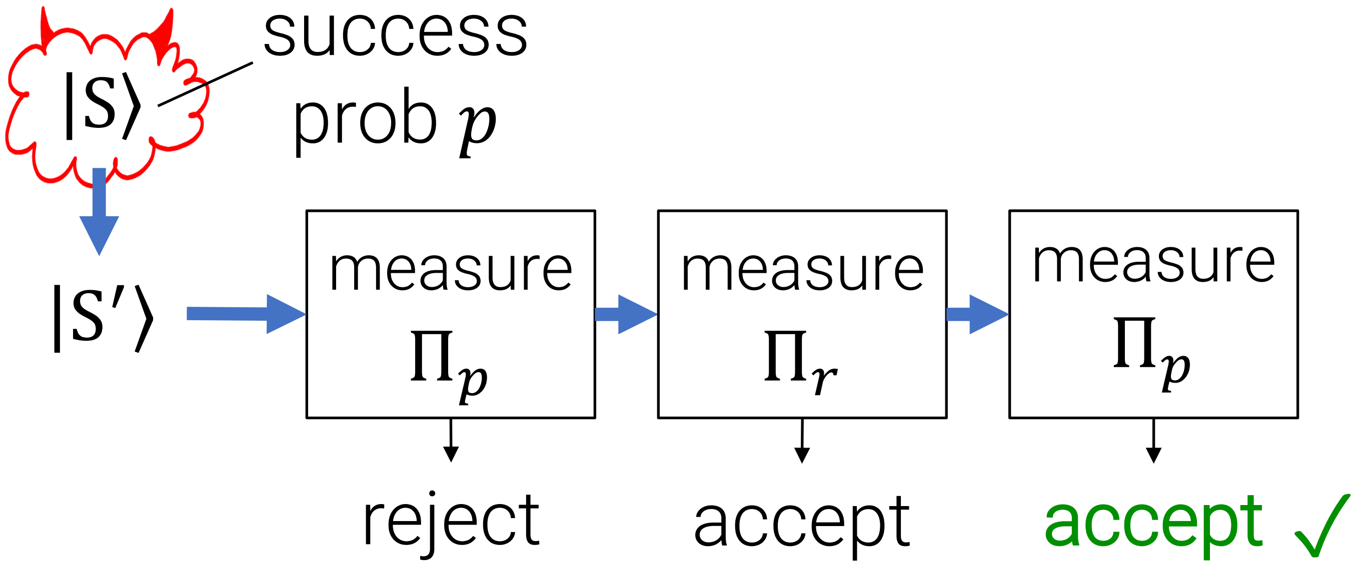

State Repair (High-Level Idea)

1) Identify a "good subspace" Π_p where $|S\rangle \in \Pi_p$, and moreover *all states* in Π_p have success prob $\geq p$.
2) Alternate two projective measurements until the state is in Π_p:
   - the binary measurement $(\Pi_r, \mathbb{I} - \Pi_r)$ that disturbed $|S\rangle$
   - the binary measurement $(\Pi_p, \mathbb{I} - \Pi_p)$

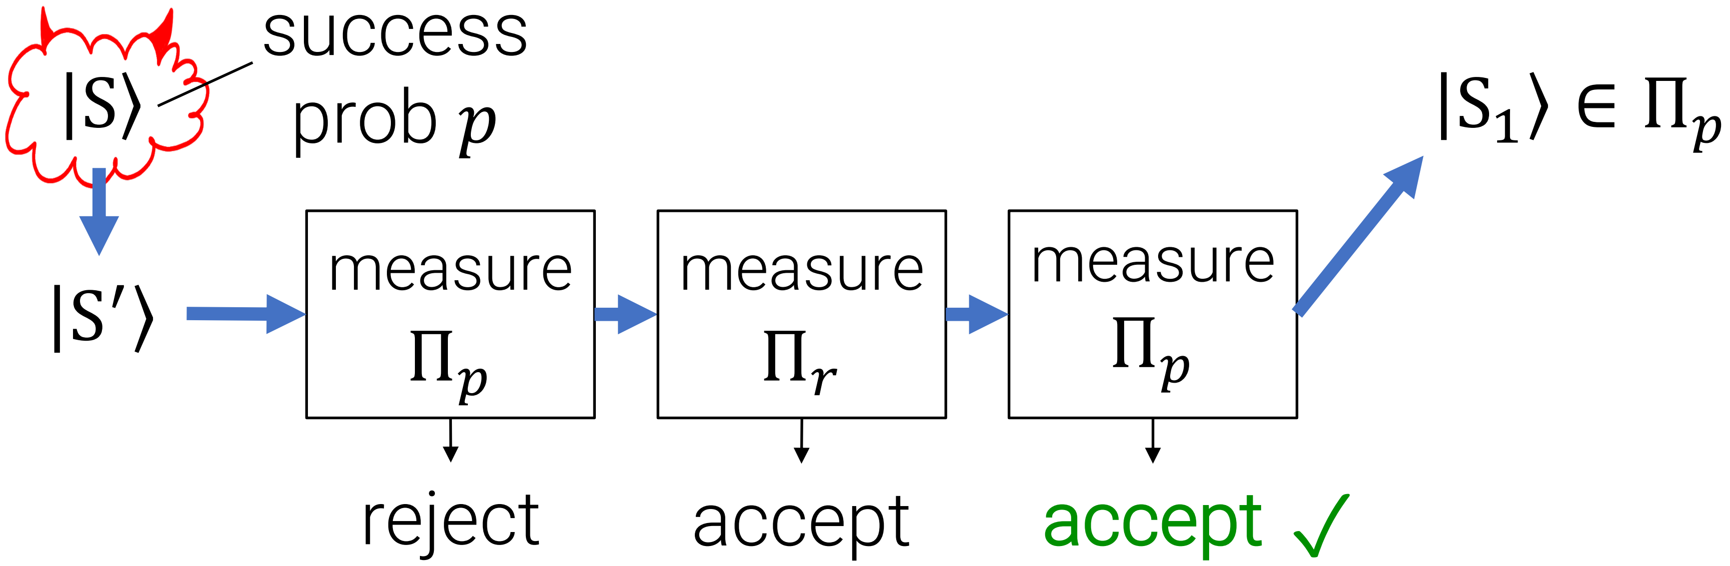

State Repair (High-Level Idea)

1) Identify a “good subspace” Π_p where $|S\rangle \in \Pi_p$, and moreover *all states* in Π_p have success prob $\geq p$.
2) Alternate two projective measurements until the state is in Π_p:
   - the binary measurement $(\Pi_r, \mathbb{I} - \Pi_r)$ that disturbed $|S\rangle$
   - the binary measurement $(\Pi_p, \mathbb{I} - \Pi_p)$

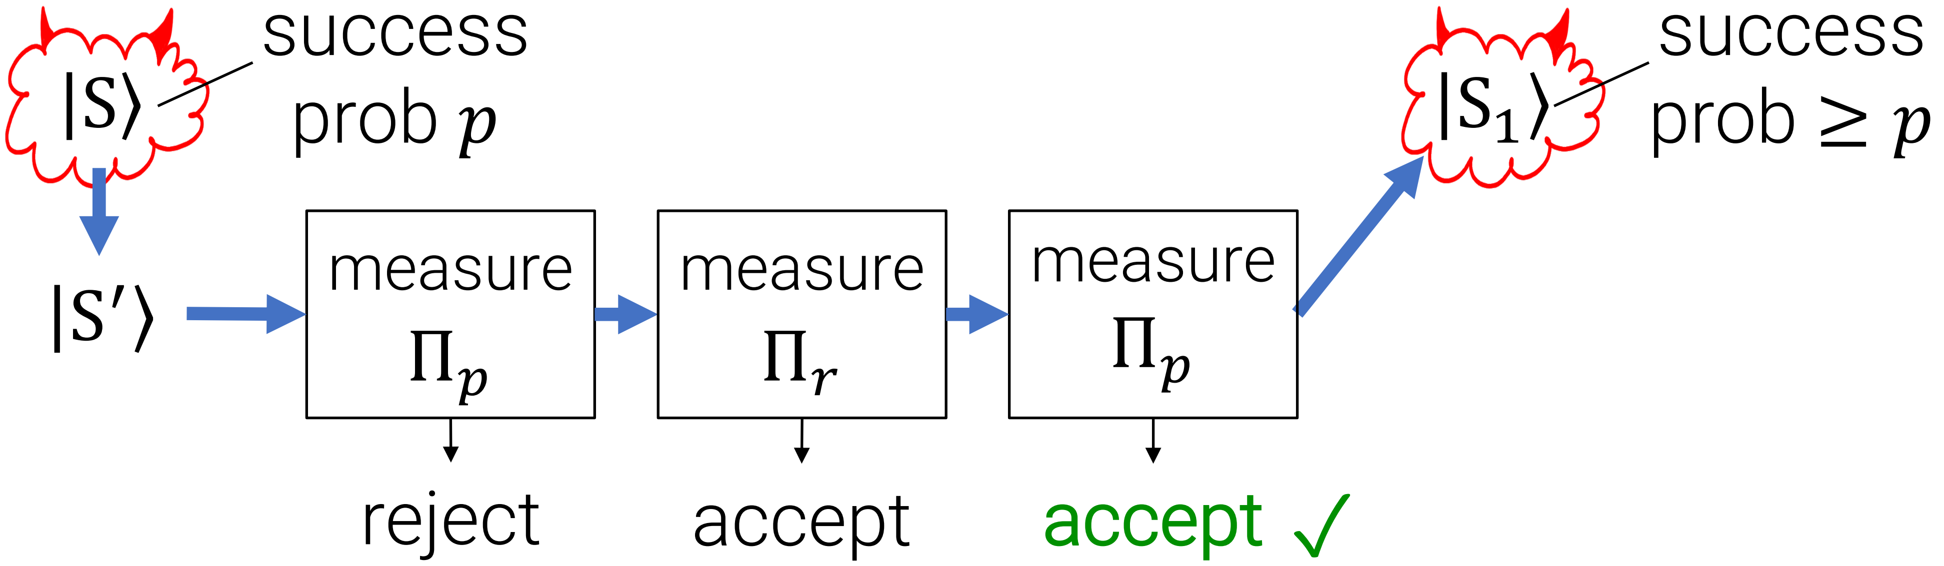

State Repair (High-Level Idea)

1) Identify a "good subspace" Π_p where $|S\rangle \in \Pi_p$, and moreover *all states* in Π_p have success prob $\geq p$.
2) Alternate two projective measurements until the state is in Π_p:
   - the binary measurement $(\Pi_r, \mathbb{I} - \Pi_r)$ that disturbed $|S\rangle$
   - the binary measurement $(\Pi_p, \mathbb{I} - \Pi_p)$

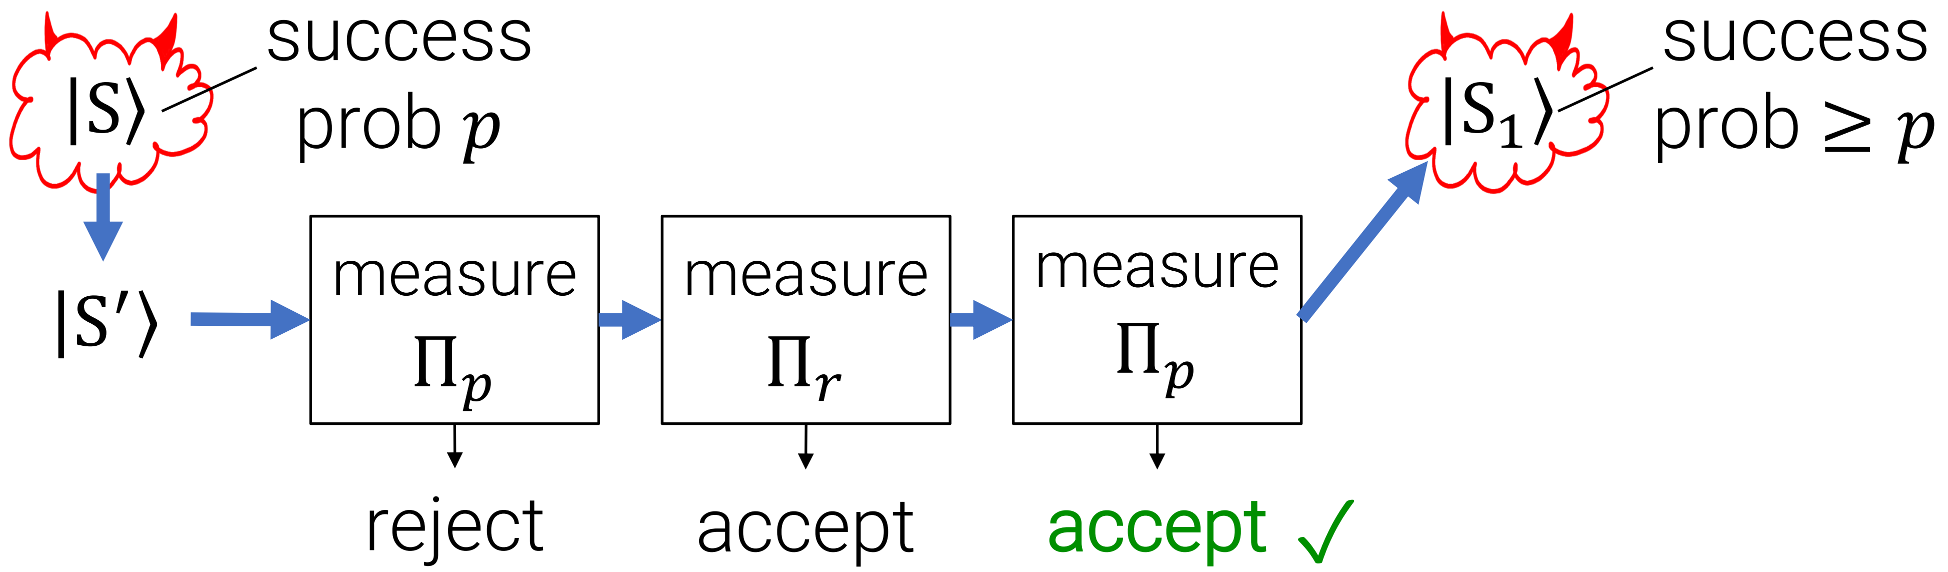

Missing Details

1) How do we know this process terminates?

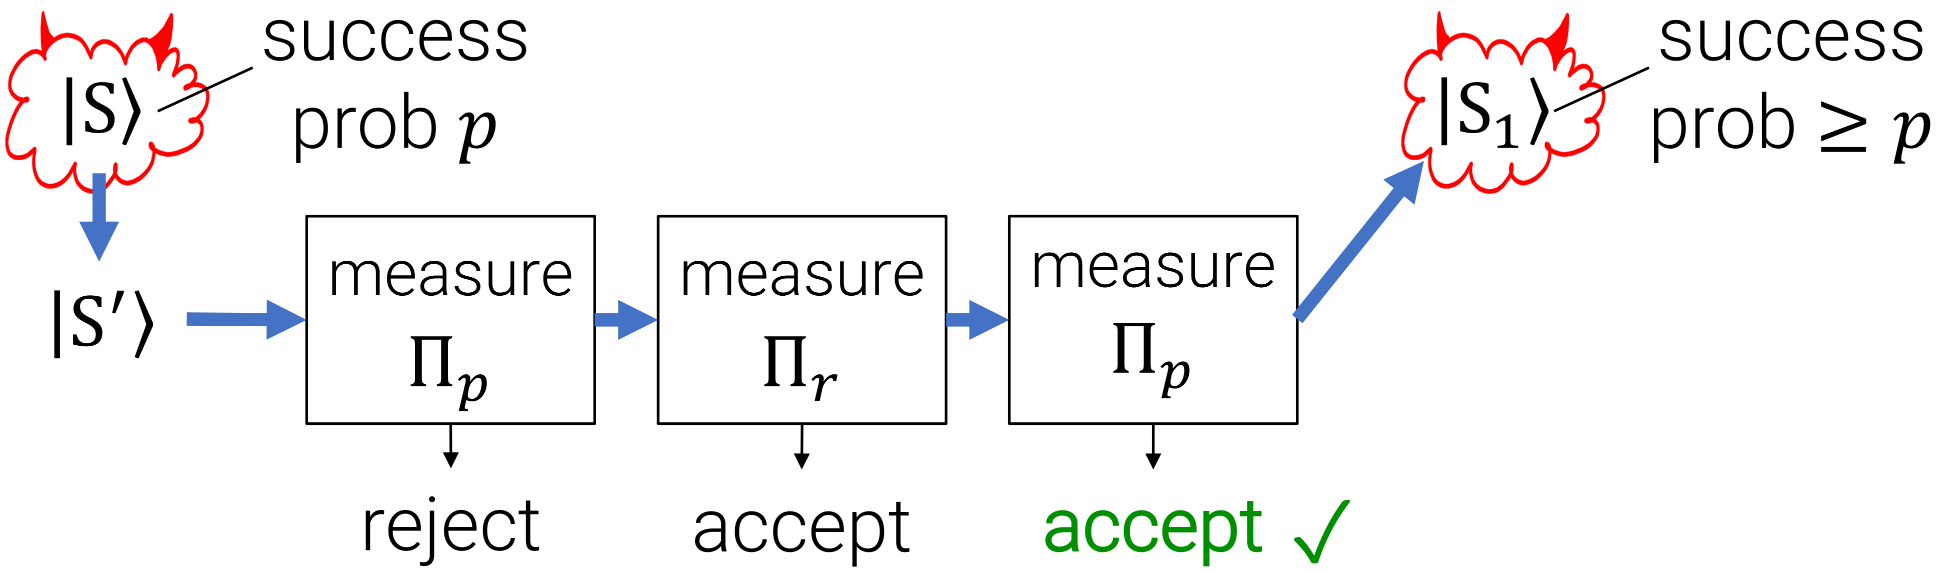

Missing Details

1) How do we know this process terminates?

2) How do we define Π_p? (In particular, we need to be able to measure Π_p efficiently.)

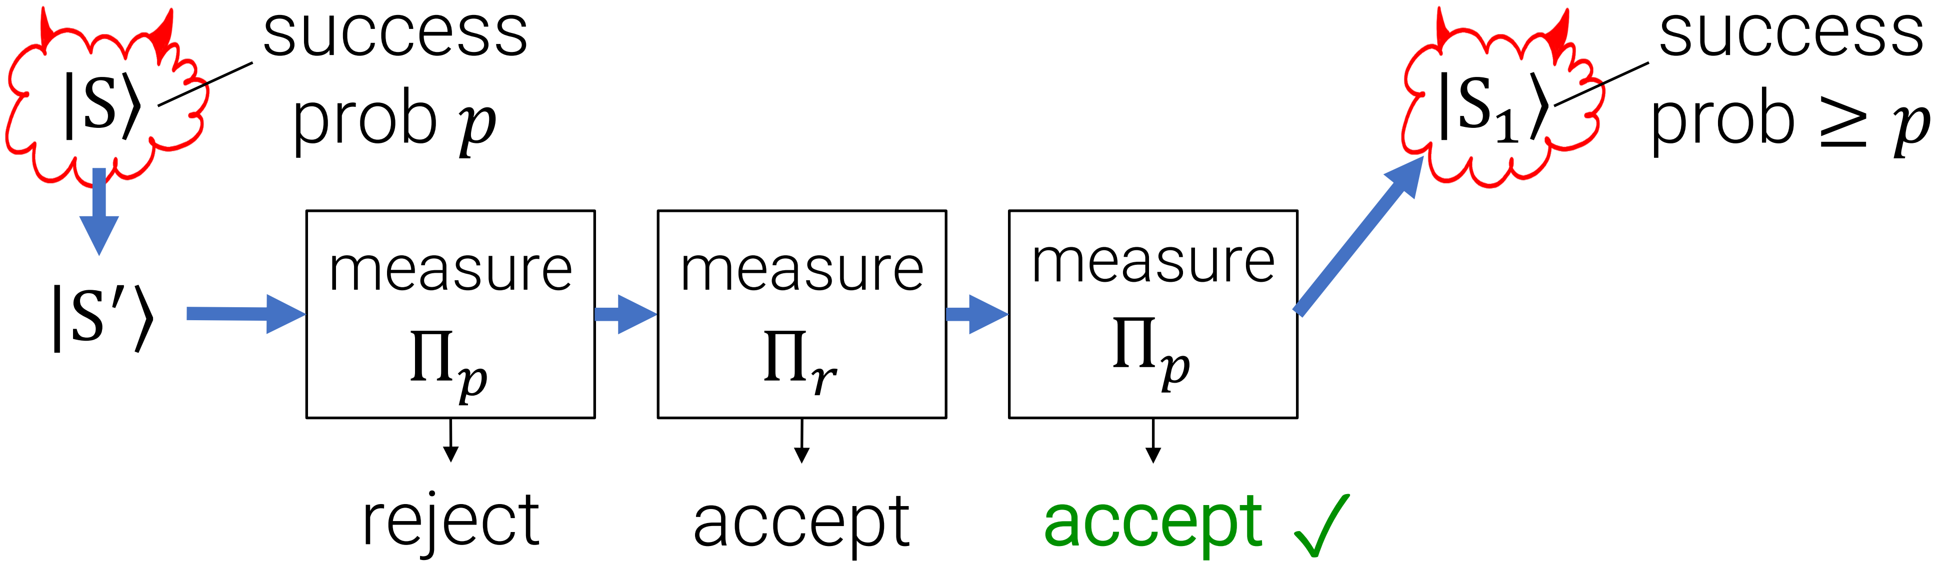

1) How do we know this process terminates?

To see why alternating Π_p and Π_r measurements eventually produces a state in Π_p, consider the 2-dim case:

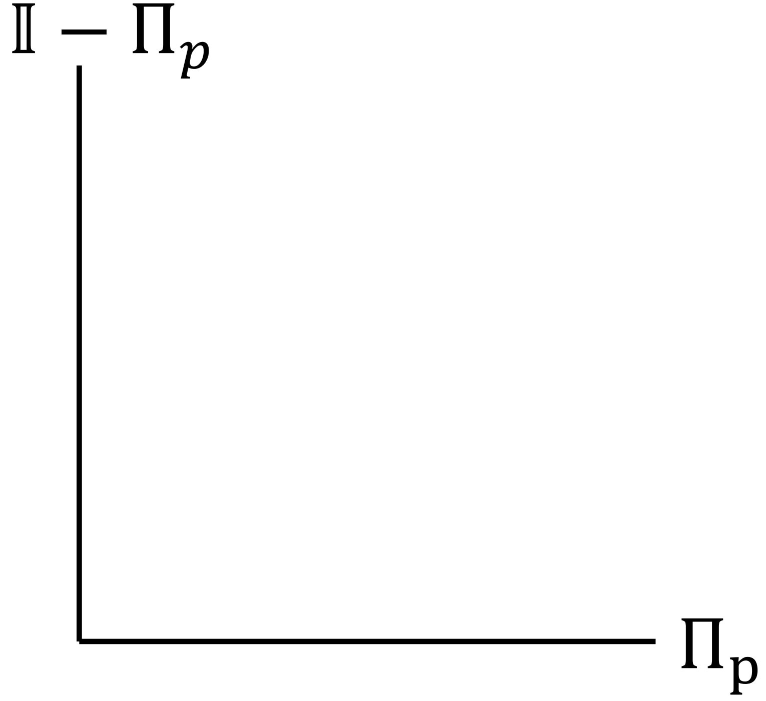

To see why alternating Π_p and Π_r measurements eventually produces a state in Π_p, consider the 2-dim case:

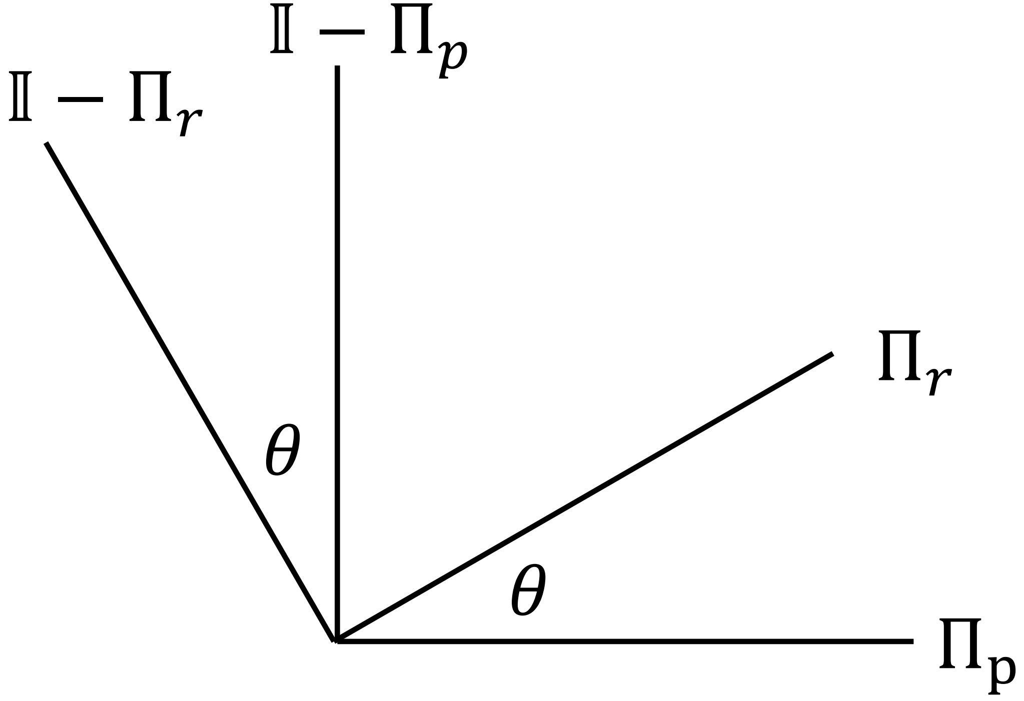

To see why alternating Π_p and Π_r measurements eventually produces a state in Π_p, consider the 2-dim case:

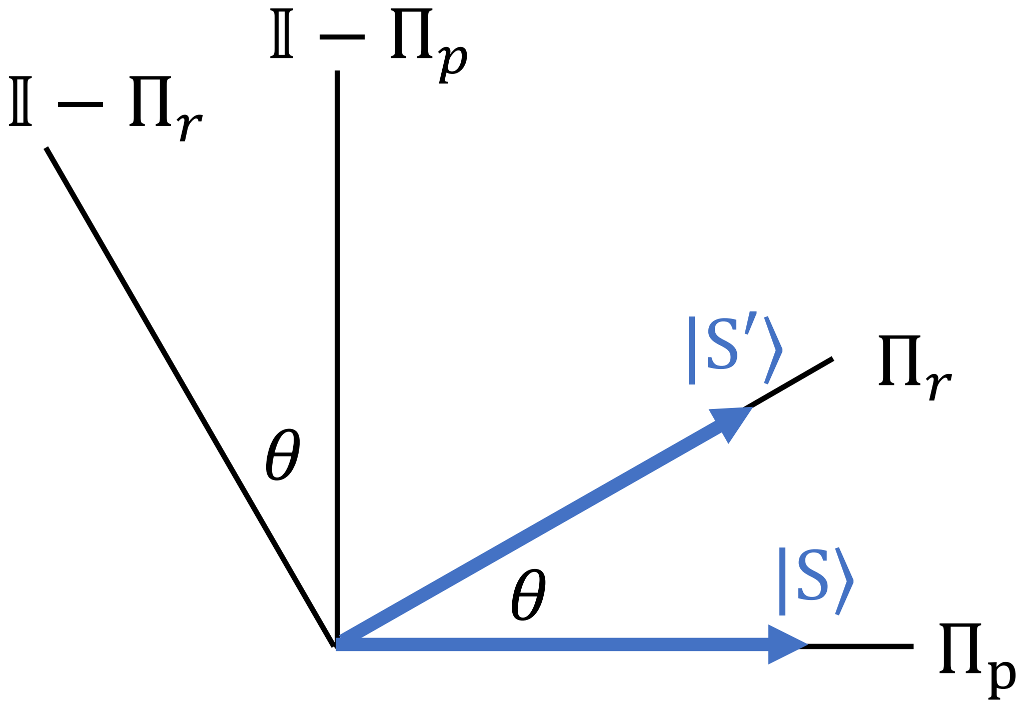

To see why alternating Π_p and Π_r measurements eventually produces a state in Π_p, consider the 2-dim case:

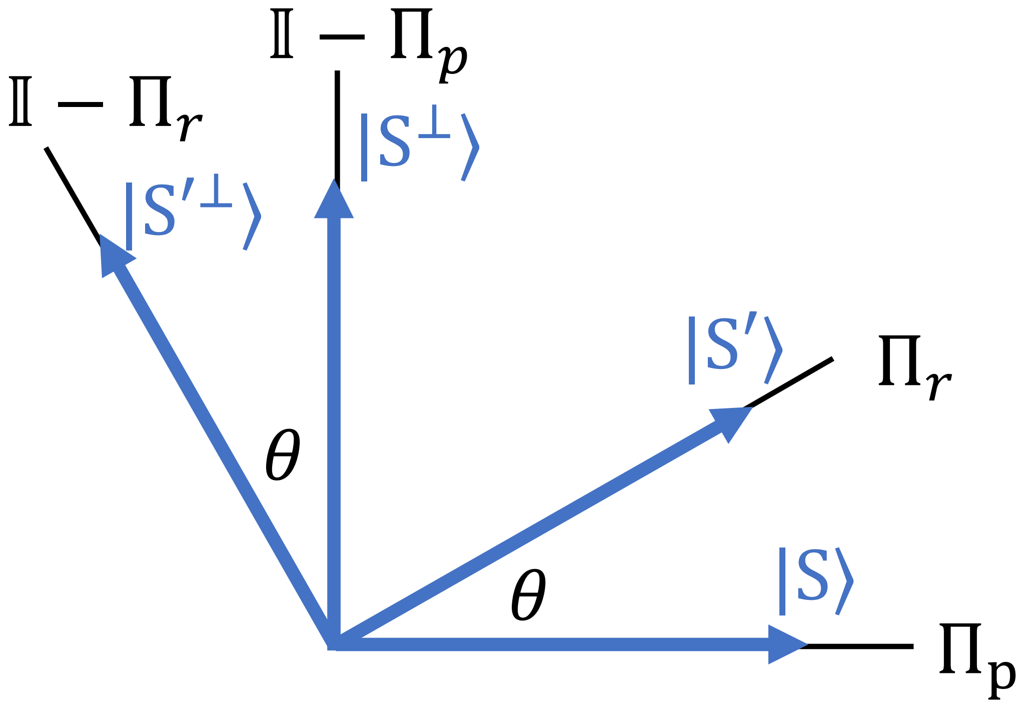

To see why alternating Π_p and Π_r measurements eventually produces a state in Π_p, consider the 2-dim case:

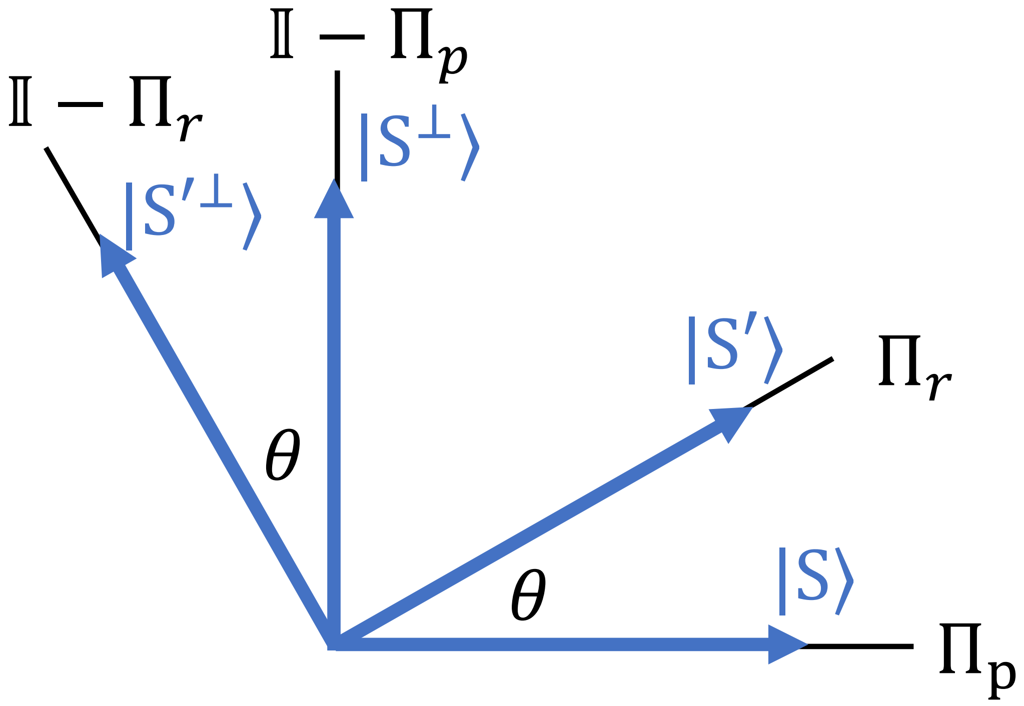

- State “jumps” between the 4 labeled states.

To see why alternating Π_p and Π_r measurements eventually produces a state in Π_p, consider the 2-dim case:

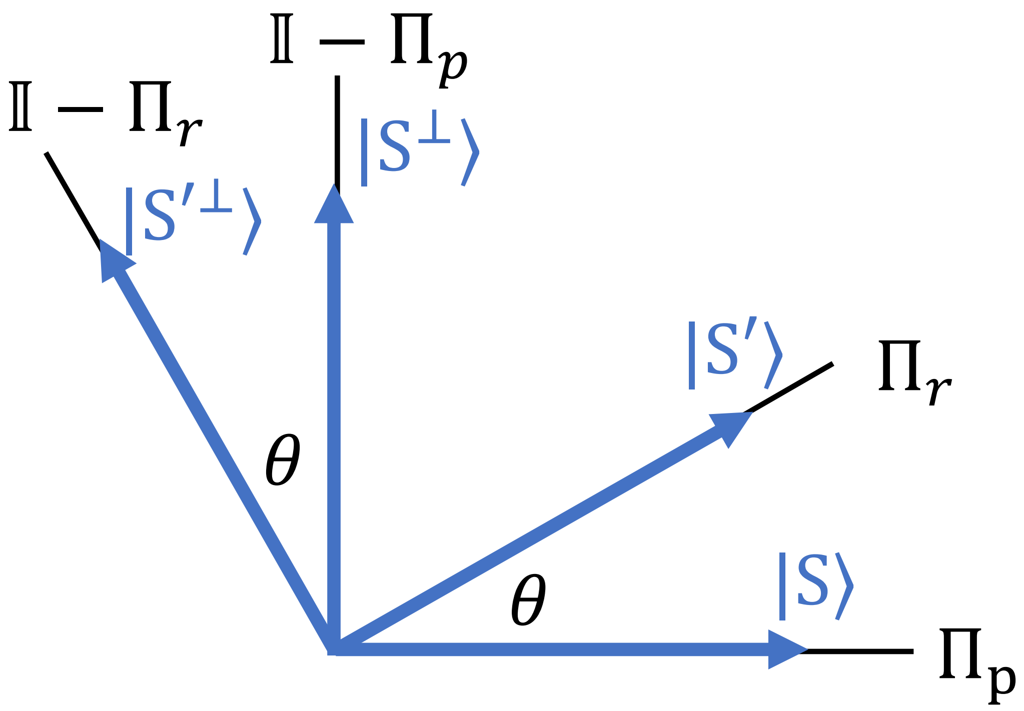

- State "jumps" between the 4 labeled states.
- If we start at $|S\rangle \in \Pi_p$, we return to Π_p in expected $O(1)$ steps for any θ.

To see why alternating Π_p and Π_r measurements eventually produces a state in Π_p, consider the 2-dim case:

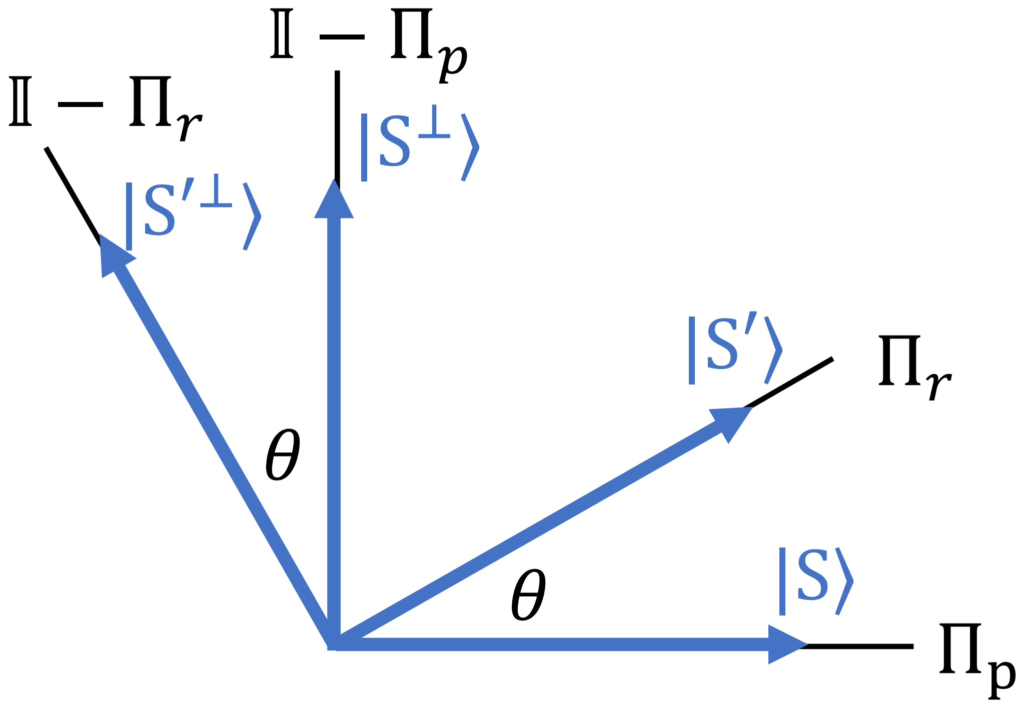

- State "jumps" between the 4 labeled states.
- If we start at $|S\rangle \in \Pi_p$, we return to Π_p in expected $O(1)$ steps for any θ.
- Jordan's lemma extends this to higher dimensions.

To see why alternating Π_p and Π_r measurements eventually produces a state in Π_p, consider the 2-dim case:

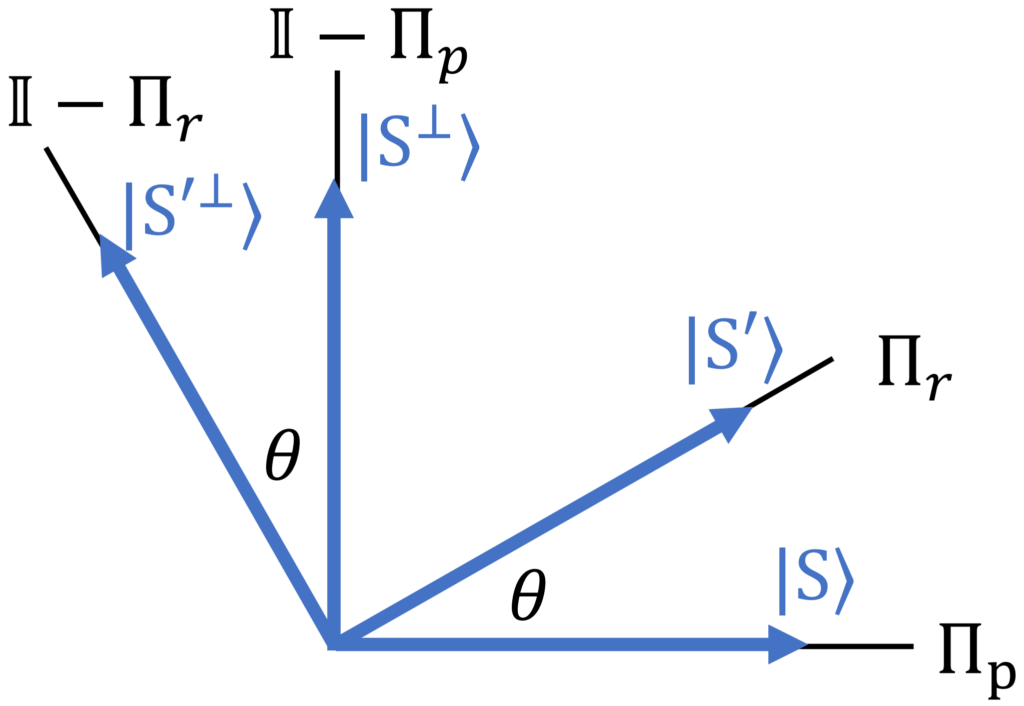

- State "jumps" between the 4 labeled states.
- If we start at $|S\rangle \in \Pi_p$, we return to Π_p in expected $O(1)$ steps for any θ.
- Jordan's lemma extends this to higher dimensions.

Note: this works for any two binary projective measurements.

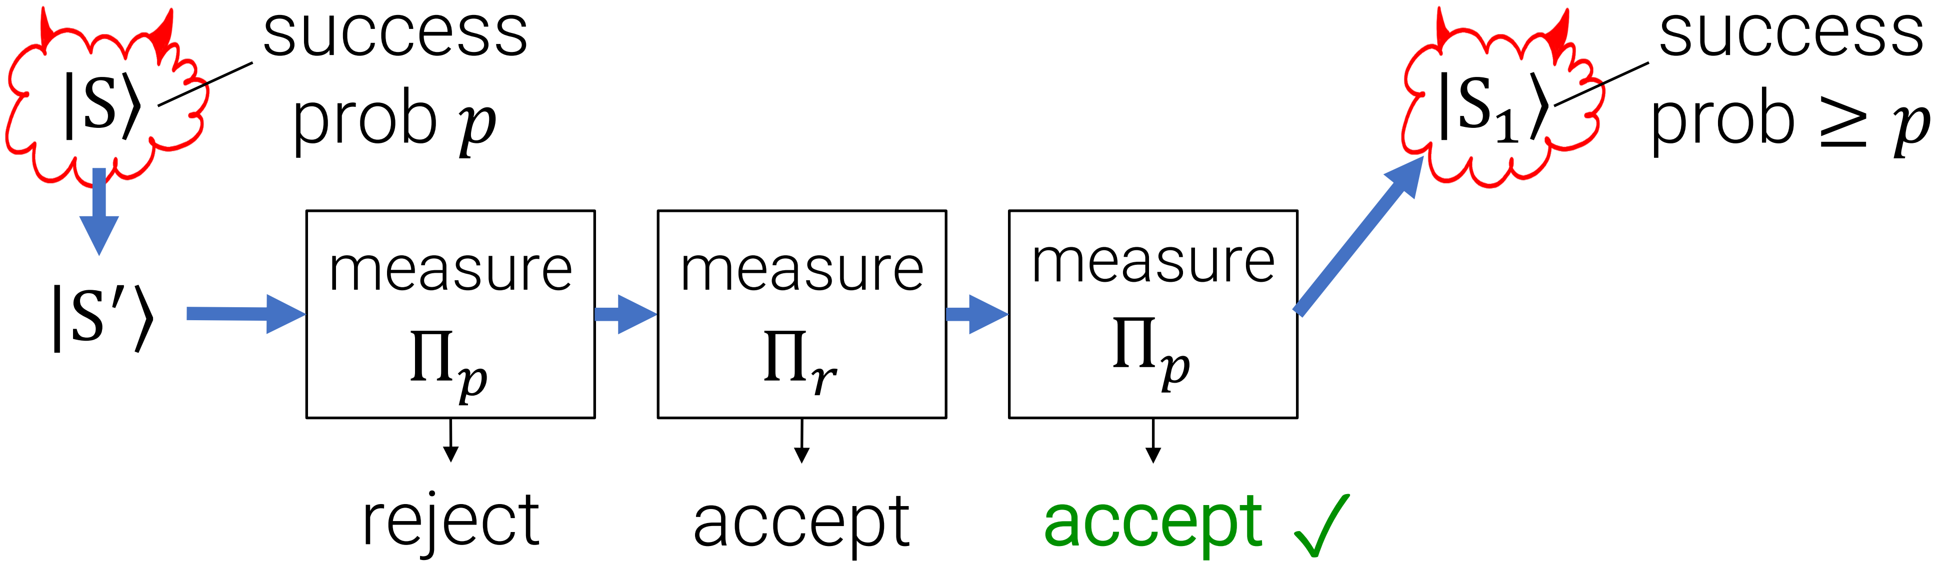

Missing Details

~~1) How do we know this process terminates?~~

2) How do we define Π_p? (In particular, we need to be able to measure Π_p efficiently.)

As currently specified, a projection Π_p onto states with success prob $\geq p$ is unlikely to be efficient.

2) How do we define Π_p? (In particular, we need to be able to measure Π_p efficiently.)

As currently specified, a projection Π_p onto states with success prob $\geq p$ is unlikely to be efficient.

However, we can achieve a relaxed version of this guarantee using a technique of [MW05].

2) How do we define Π_p? (In particular, we need to be able to measure Π_p efficiently.)

Recap: The Marriott-Watrous Procedure

Given a binary-output quantum circuit C and an input $|S\rangle$, [MW05] gives a procedure to *estimate* $\Pr[\, C(|S\rangle) \to 1]$ to any precision.

([MW05] use this procedure for QMA amplification)

Recap: The Marriott-Watrous Procedure

Given a binary-output quantum circuit C and an input $|S\rangle$, [MW05] gives a procedure to *estimate* $\Pr[\, C(|S\rangle) \to 1]$ to any precision.

We'll use [MW05] to estimate *success probability*.

Recap: The Marriott-Watrous Procedure

Given a binary-output quantum circuit C and an input $|S\rangle$, [MW05] gives a procedure to ***estimate*** $\Pr[\, C(|S\rangle) \rightarrow 1]$ to any precision.

We'll use [MW05] to estimate ***success probability***.

$C(|S\rangle)$:

1) Prepare the superposition of challenges $\sum_{r \in R} |r\rangle$.
2) Compute (in superposition) the response of adversary $|S\rangle$.
3) Output $V(r, z)$.

success
prob p_0 —— $|S\rangle$

For this talk, we'll only need to know two things about the MW estimator.

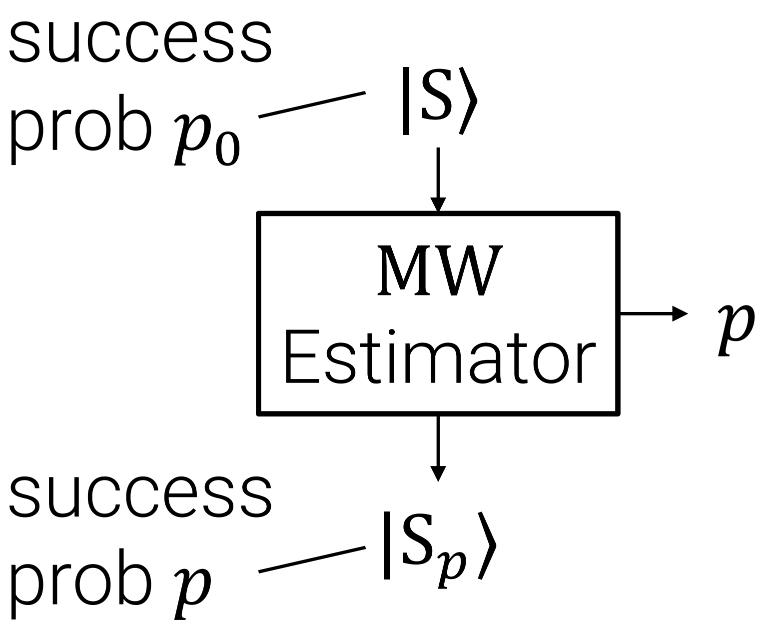

For this talk, we'll only need to know two things about the MW estimator.

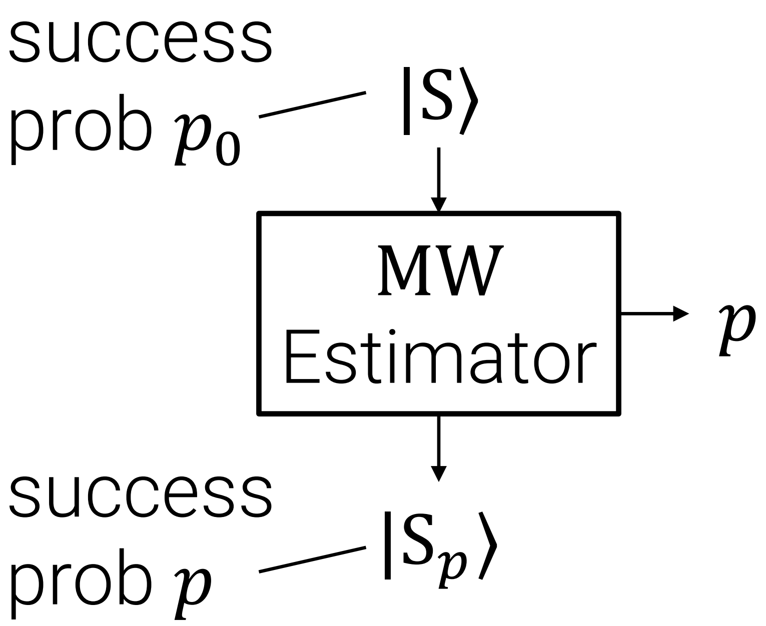

For this talk, we'll only need to know two things about the MW estimator.

Key Properties

1) $\mathbb{E}[p] = p_0$

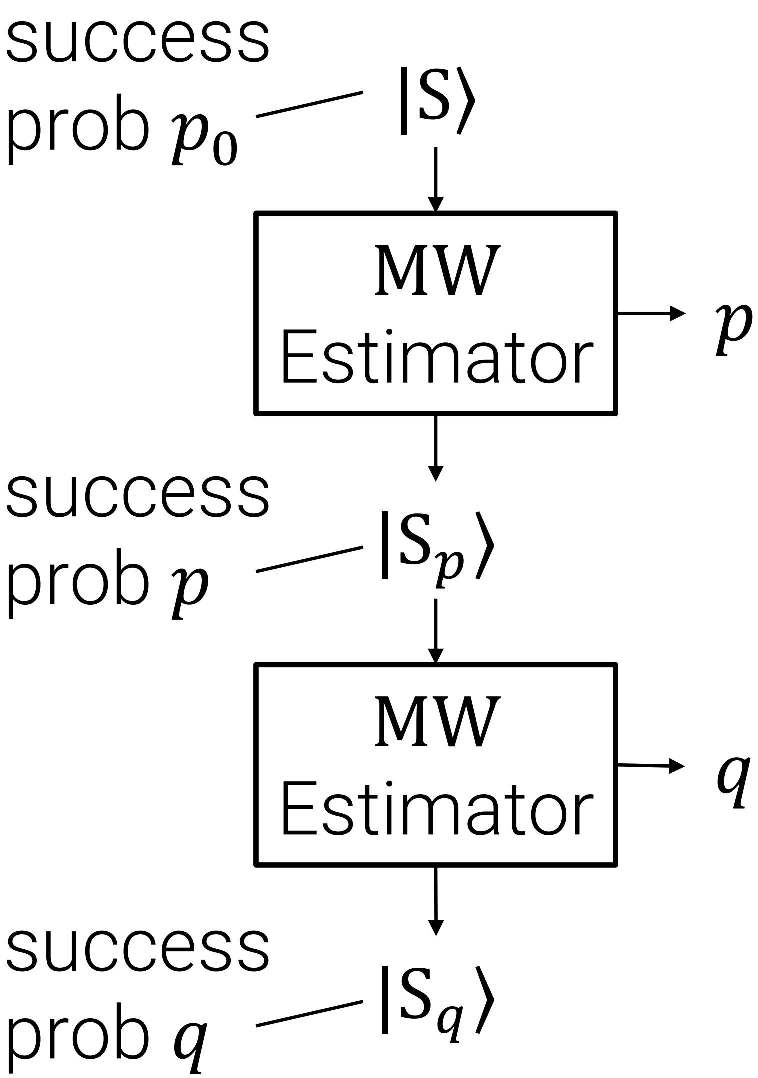

For this talk, we'll only need to know two things about the MW estimator.

Key Properties

1) $\mathbb{E}[p] = p_0$

2) If we apply MW twice, the two outcomes p, q are close with high probability.

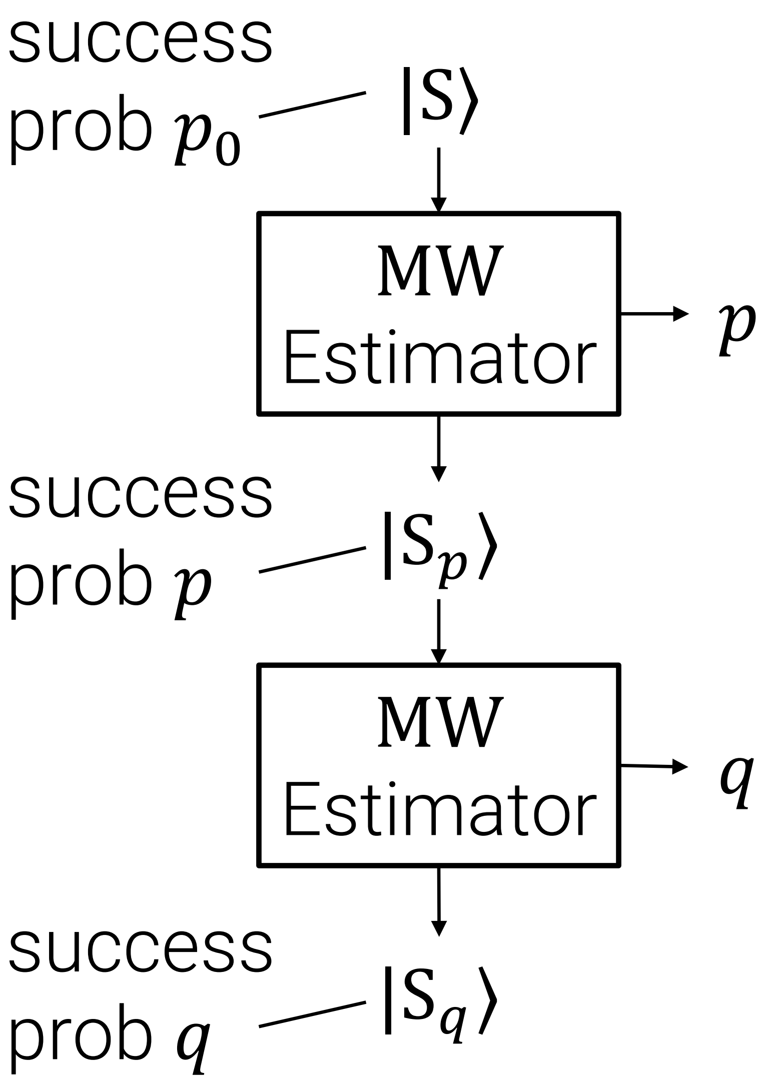

For this talk, we'll only need to know two things about the MW estimator.

Key Properties

1) $\mathbb{E}[p] = p_0$

2) If we apply MW twice, the two outcomes p, q are close with high probability. Formally, MW achieves

$$\Pr[|p - q| \leq \varepsilon] \geq 1 - \delta$$

with $poly(\frac{1}{\varepsilon}, \log\left(\frac{1}{\delta}\right))$ runtime.

For this talk, we'll only need to know two things about the $\mathbf{MW}$ estimator.

As in [Zha20], we call this "(ε, δ)-almost-projective."

Key Properties

1) $\mathbb{E}[p] = p_0$

2) If we apply $\mathbf{MW}$ twice, the two outcomes p, q are close with high probability. Formally, $\mathbf{MW}$ achieves

$$\Pr[|p - q| \leq \varepsilon] \geq 1 - \delta$$

with $poly(\frac{1}{\varepsilon}, \log\left(\frac{1}{\delta}\right))$ runtime.

Let's see how [MW05] fits into our approach.

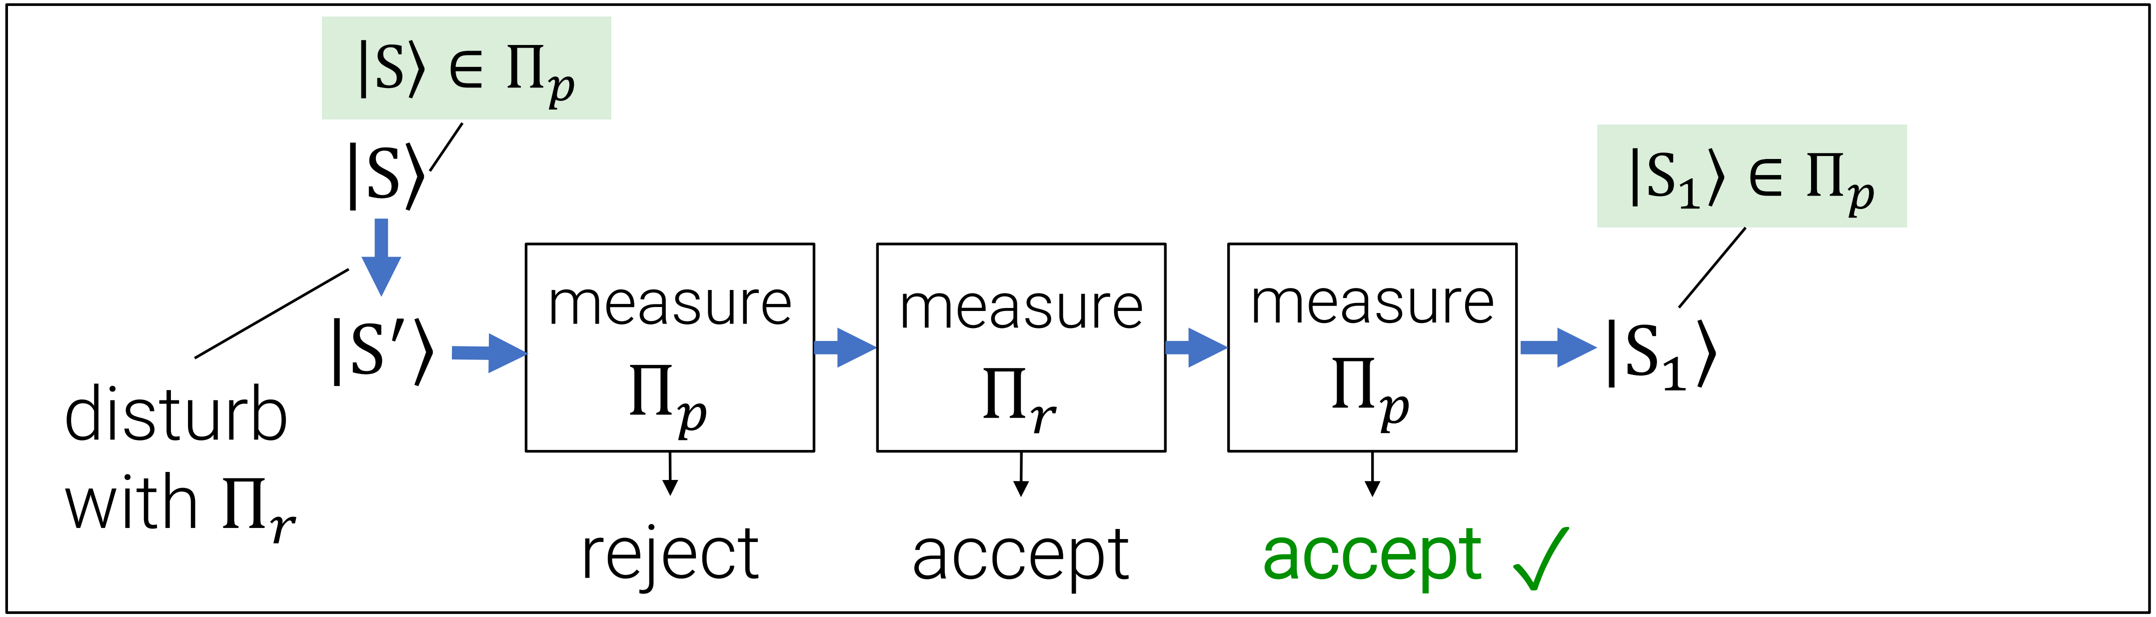

Recall: our high-level approach assumed we could measure Π_p, a projection onto states with success prob $\geq p$.

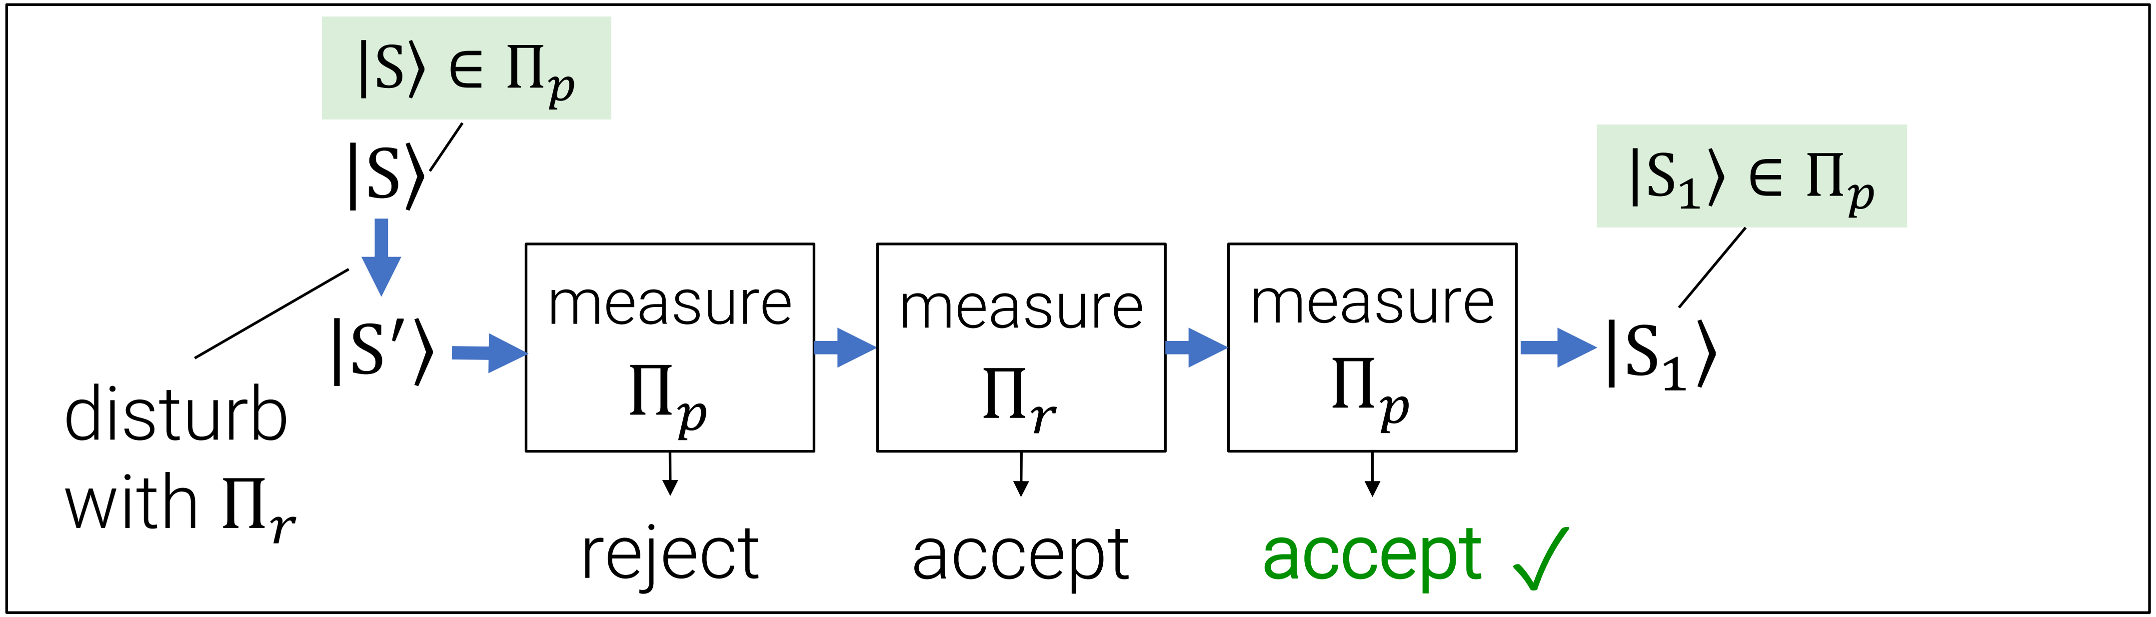

In reality, the closest thing we have is a binary measurement MW_p that runs MW and accepts if the output is $\geq p$.

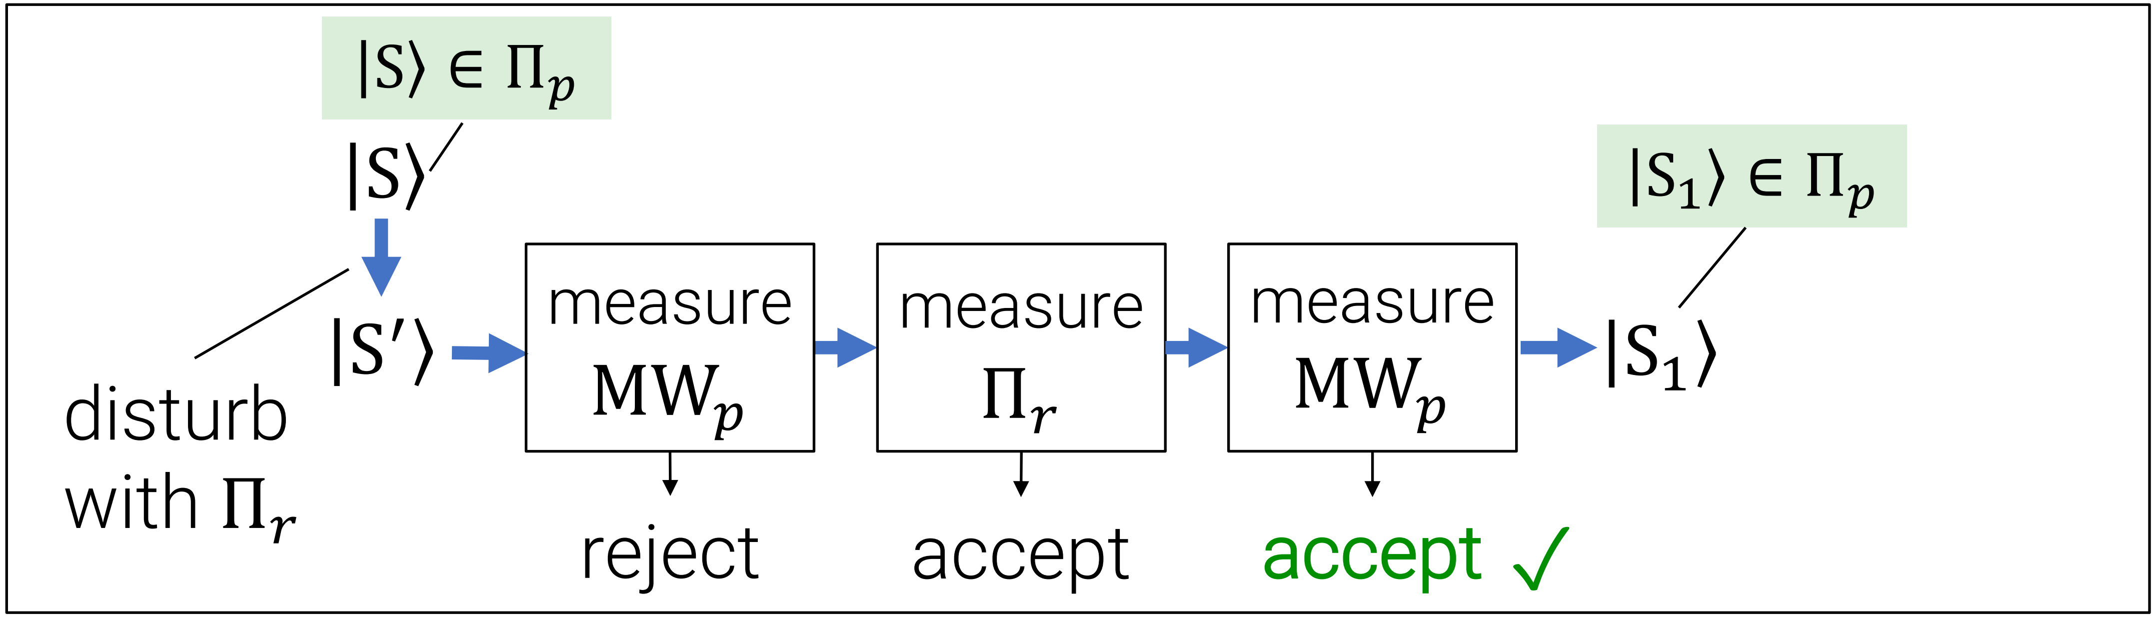

In reality, the closest thing we have is a binary measurement $\mathbf{MW}_p$ that runs $\mathbf{MW}$ and accepts if the output is $\geq p$.

We can easily swap out the Π_p measurements for $\mathbf{MW}_p$, but we also need to update the invariant that we want a state $\in \Pi_p$.

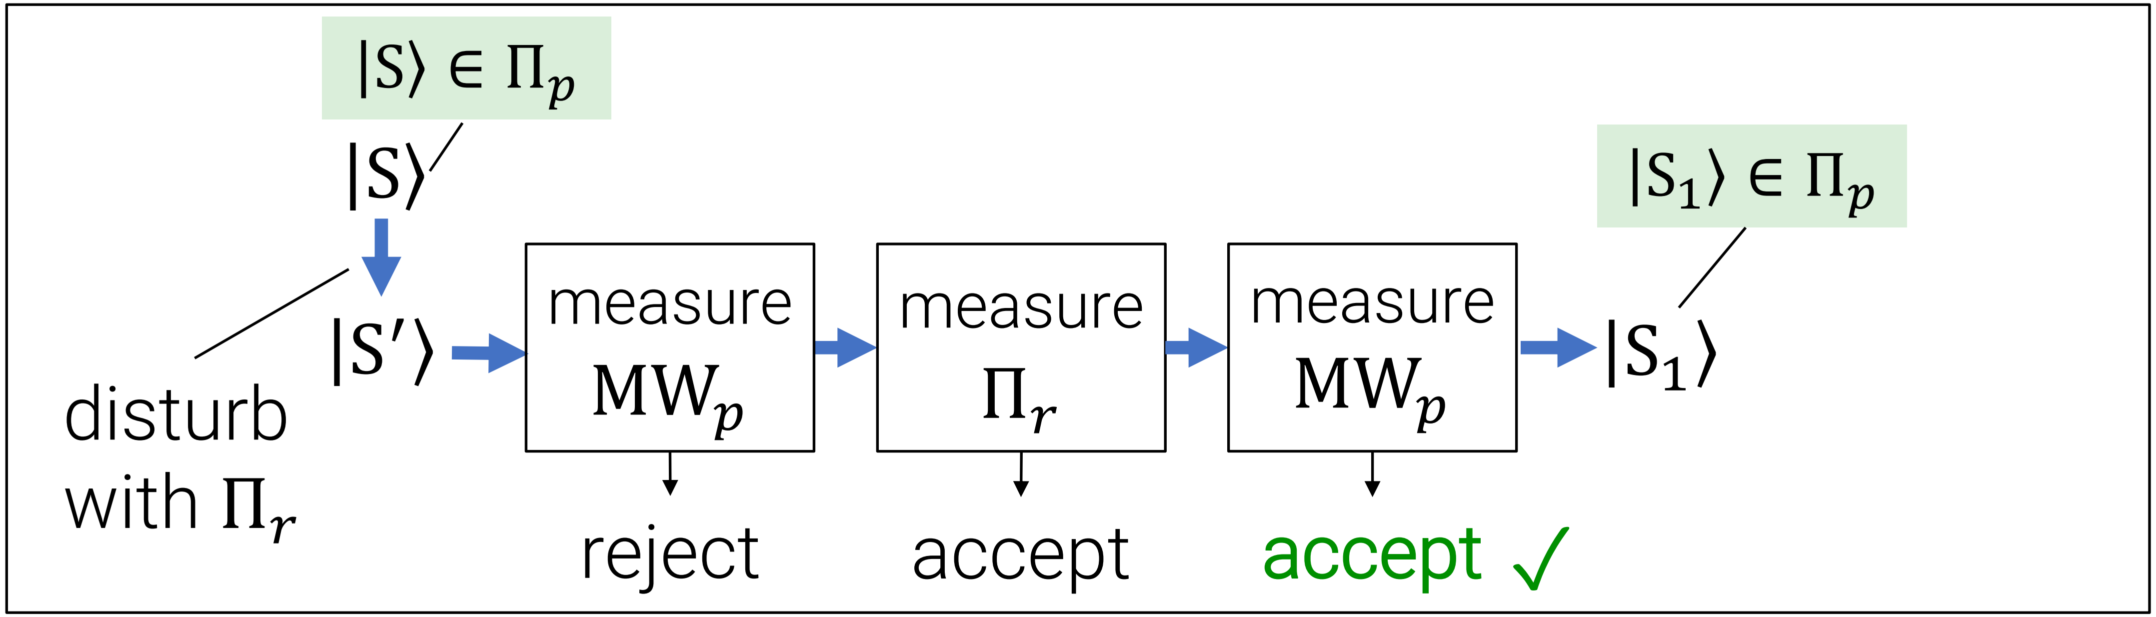

In reality, the closest thing we have is a binary measurement MW_p that runs MW and accepts if the output is $\geq p$.

Fortunately, there's a natural MW-analogue:

$$\mathbb{E}[\mathrm{MW}_p(|S\rangle)] = 1 - \delta.$$

This implies success prob of $|S\rangle$ is $\mathbb{E}_{q \leftarrow \mathrm{MW}(|S\rangle)}[q] \geq p - \delta$.

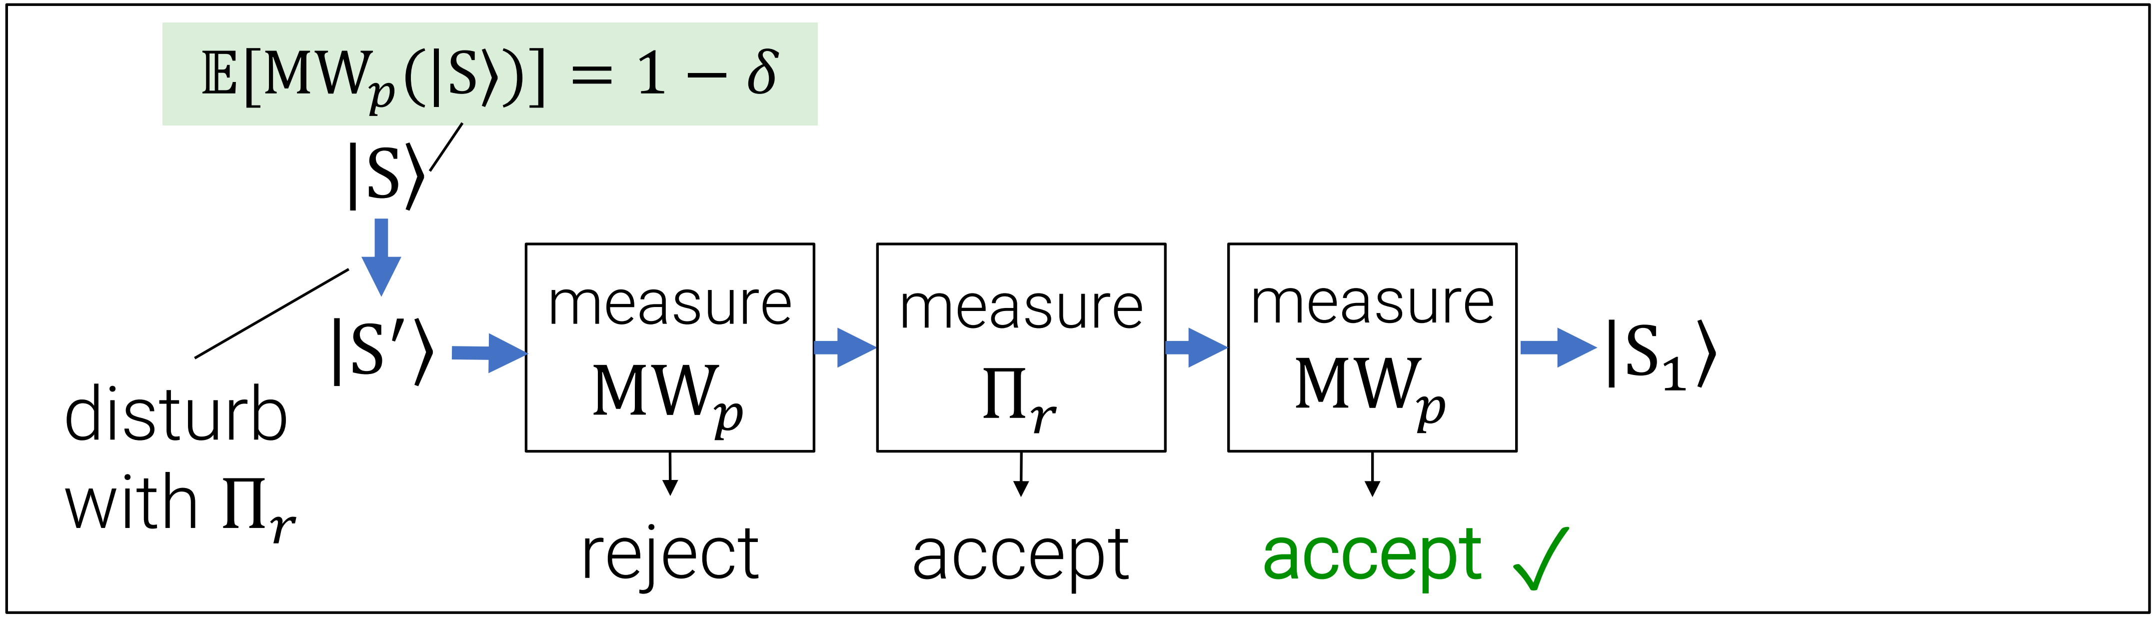

In reality, the closest thing we have is a binary measurement MW_p that runs MW and accepts if the output is $\geq p$.

Fortunately, there's a natural MW-analogue of $|S\rangle \in \Pi_p$:

$$\mathbb{E}[\mathrm{MW}_p(|S\rangle)] = 1-\delta.$$

This implies success prob of $|S\rangle$ is $\mathbb{E}_{q \leftarrow \mathrm{MW}(|S\rangle)}[q] \geq p - \delta$.

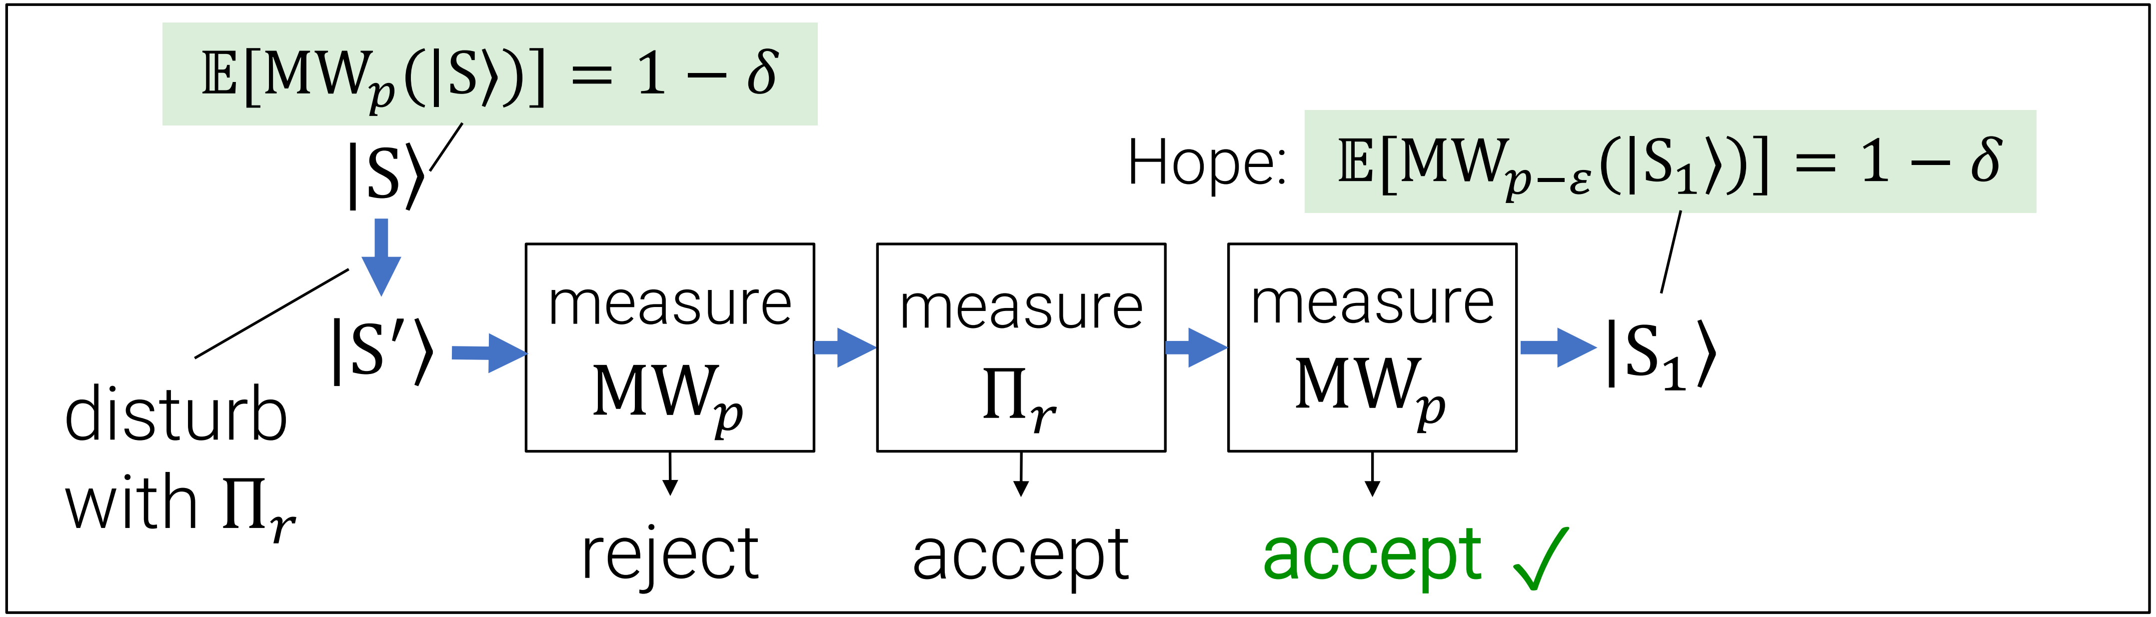

Note that the guarantee degrades: for $|\mathrm{S}_1\rangle$, the best we can hope for using (ε, δ)-almost-projectivity is $\mathbb{E}[\mathrm{MW}_{p-\varepsilon}(|\mathrm{S}_1\rangle)] = 1-\delta$.

Fortunately, there's a natural MW-analogue of $|\mathrm{S}\rangle \in \Pi_p$:

$$\mathbb{E}[\mathrm{MW}_p(|\mathrm{S}\rangle)] = 1-\delta.$$

This implies success prob of $|\mathrm{S}\rangle$ is $\mathbb{E}_{q\leftarrow\mathrm{MW}(|\mathrm{S}\rangle)}[q] \geq p-\delta$.

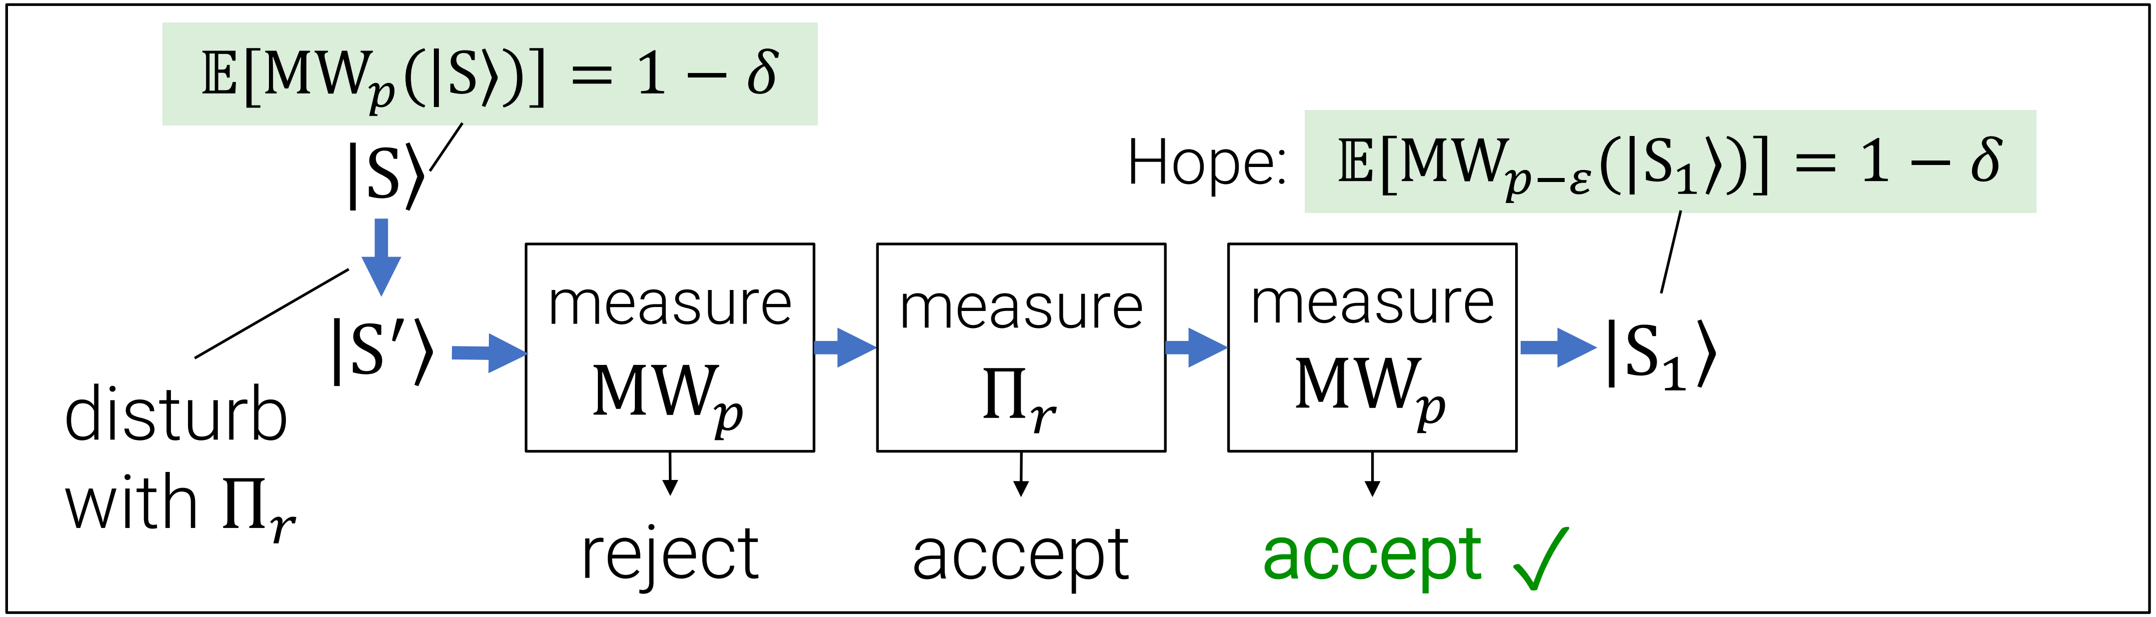

This approach seems promising, but we have a problem:

Our proof that this procedure terminates requires the measurements to be projective, but MW_p is not!

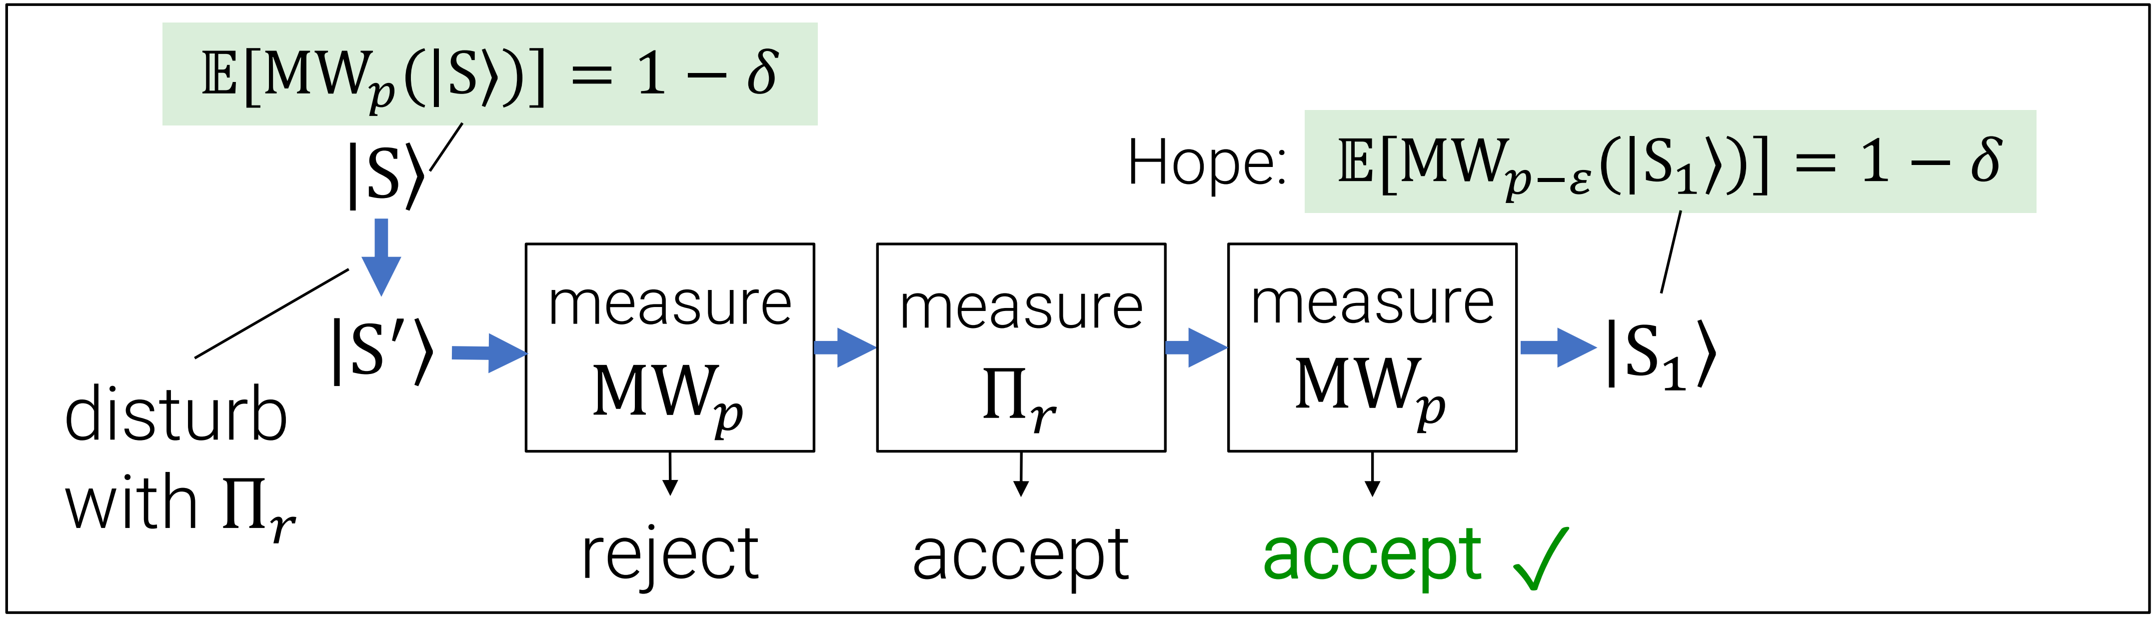

This approach seems promising, but we have a problem:

Our proof that this procedure terminates requires the measurements to be projective, but MW_p is not!

(running it twice may give different outcomes)

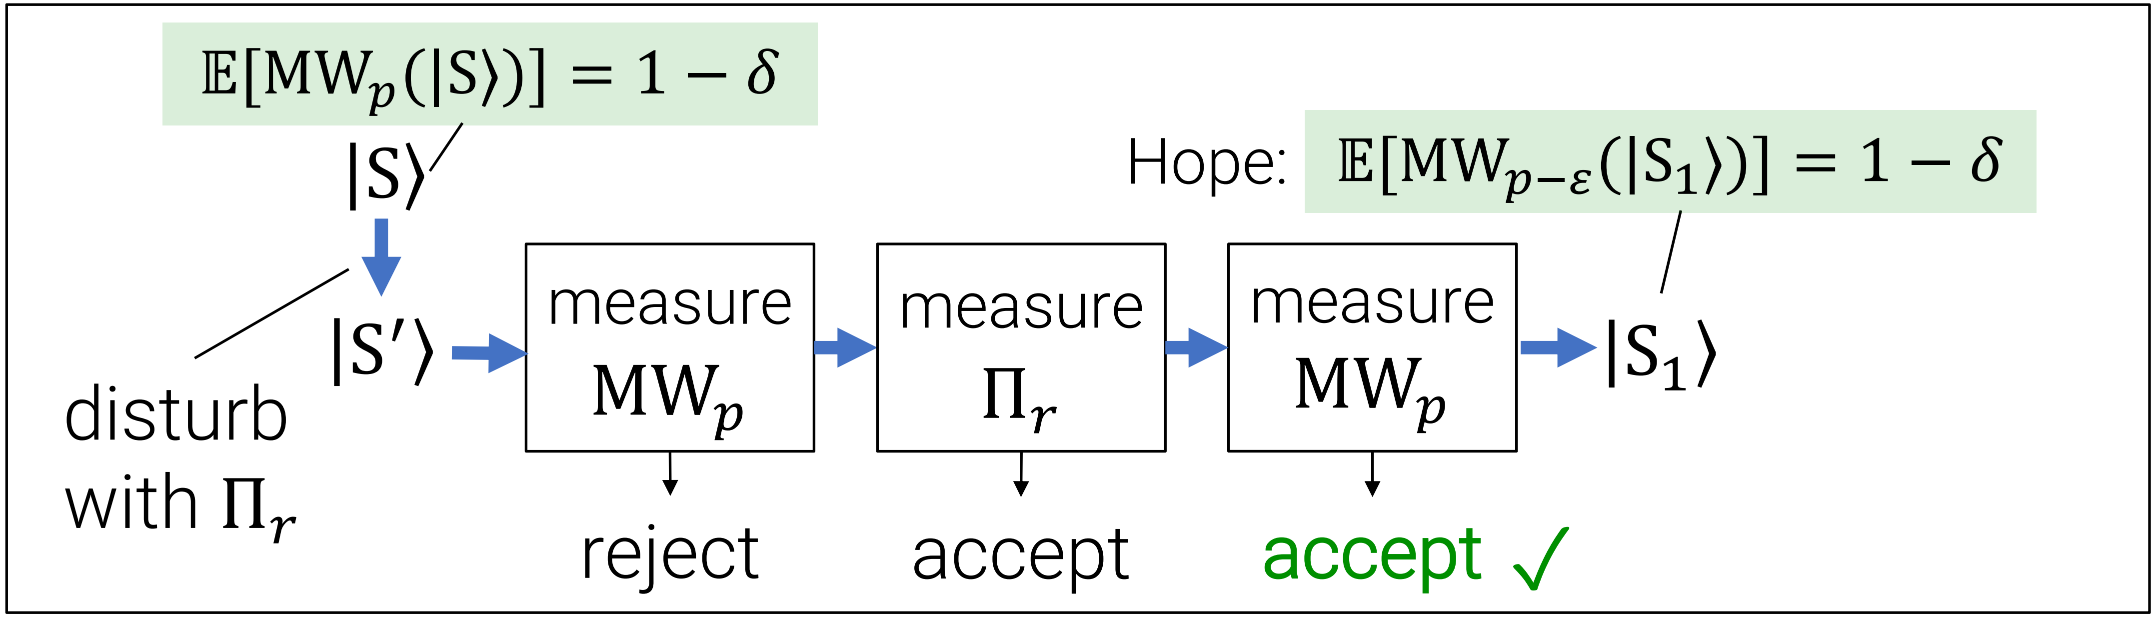

This approach seems promising, but we have a problem:

Our proof that this procedure terminates requires the measurements to be projective, but $\mathbf{MW}_p$ is not!

Easy(?) fix: Make $\mathbf{MW}_p$ projective by expanding the Hilbert space.

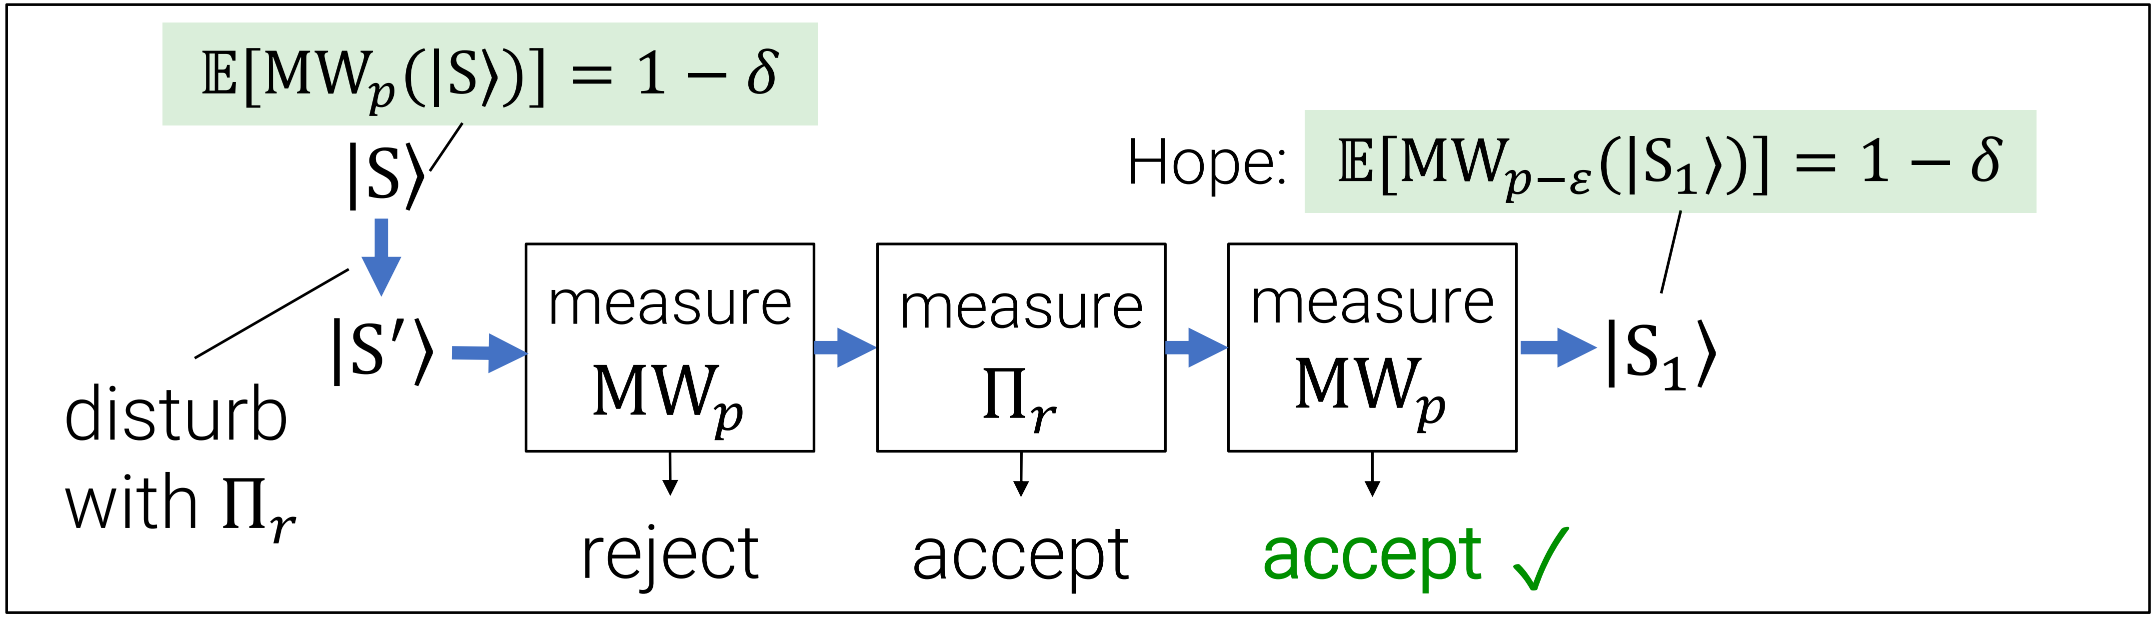

Measuring $|S'\rangle$ with MW_p can be implemented as a projective measurement of some Π_p^* on $|S'\rangle_A|0\rangle_W \in A \otimes W$.

adversary state register workspace/ancilla

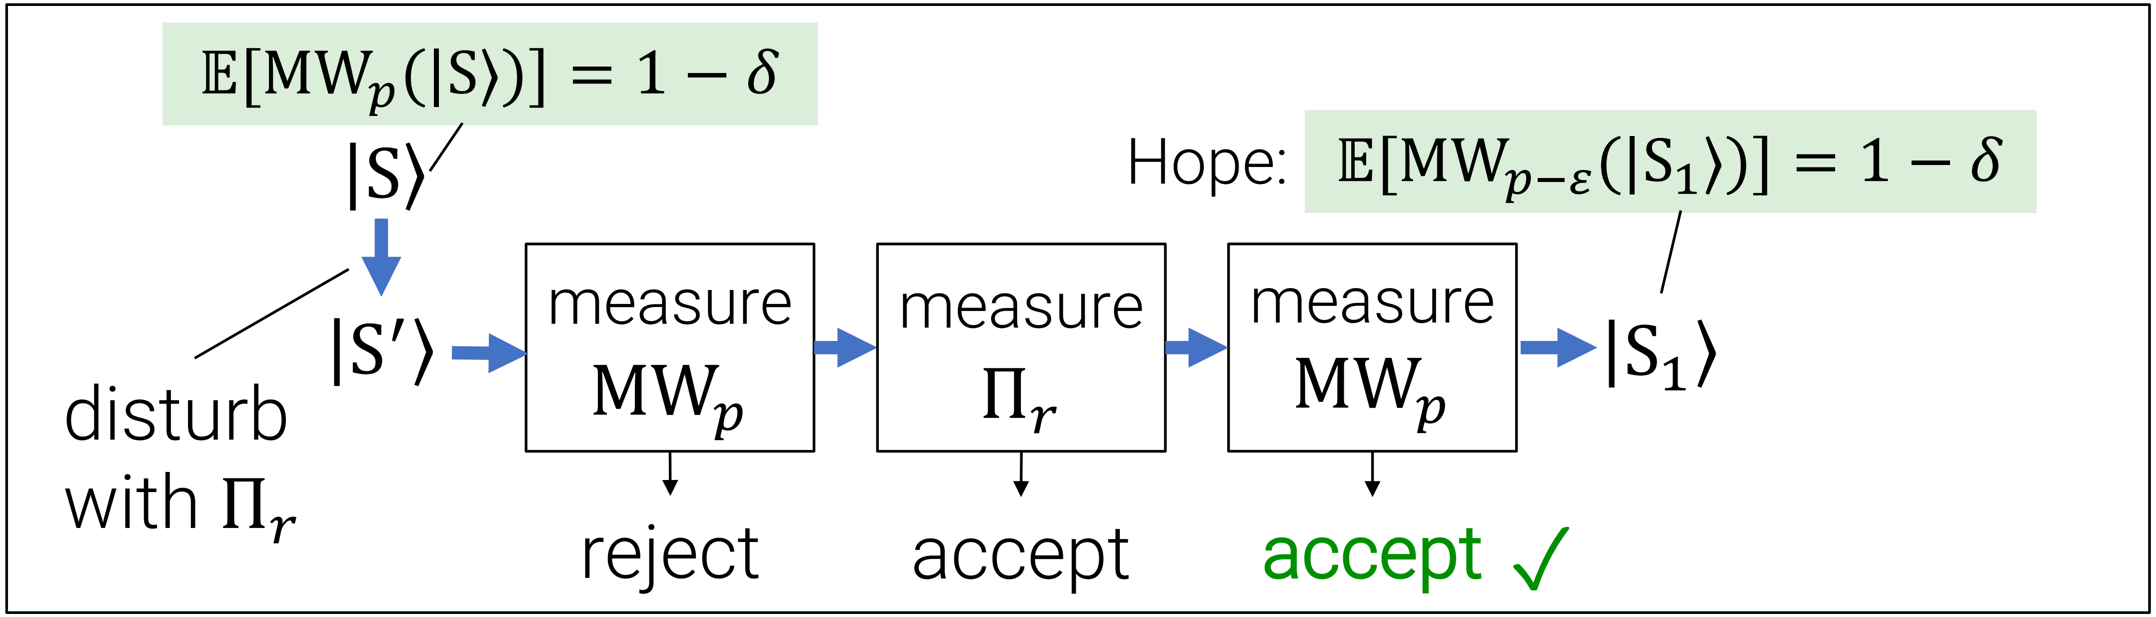

Measuring $|S'\rangle$ with MW_p can be implemented as a projective measurement of some Π_p^* on $|S'\rangle_A|0\rangle_W \in A \otimes W$.

But we need to be careful: the outcome of measuring Π_p^* only corresponds to MW_p when the W register is $|0\rangle$.

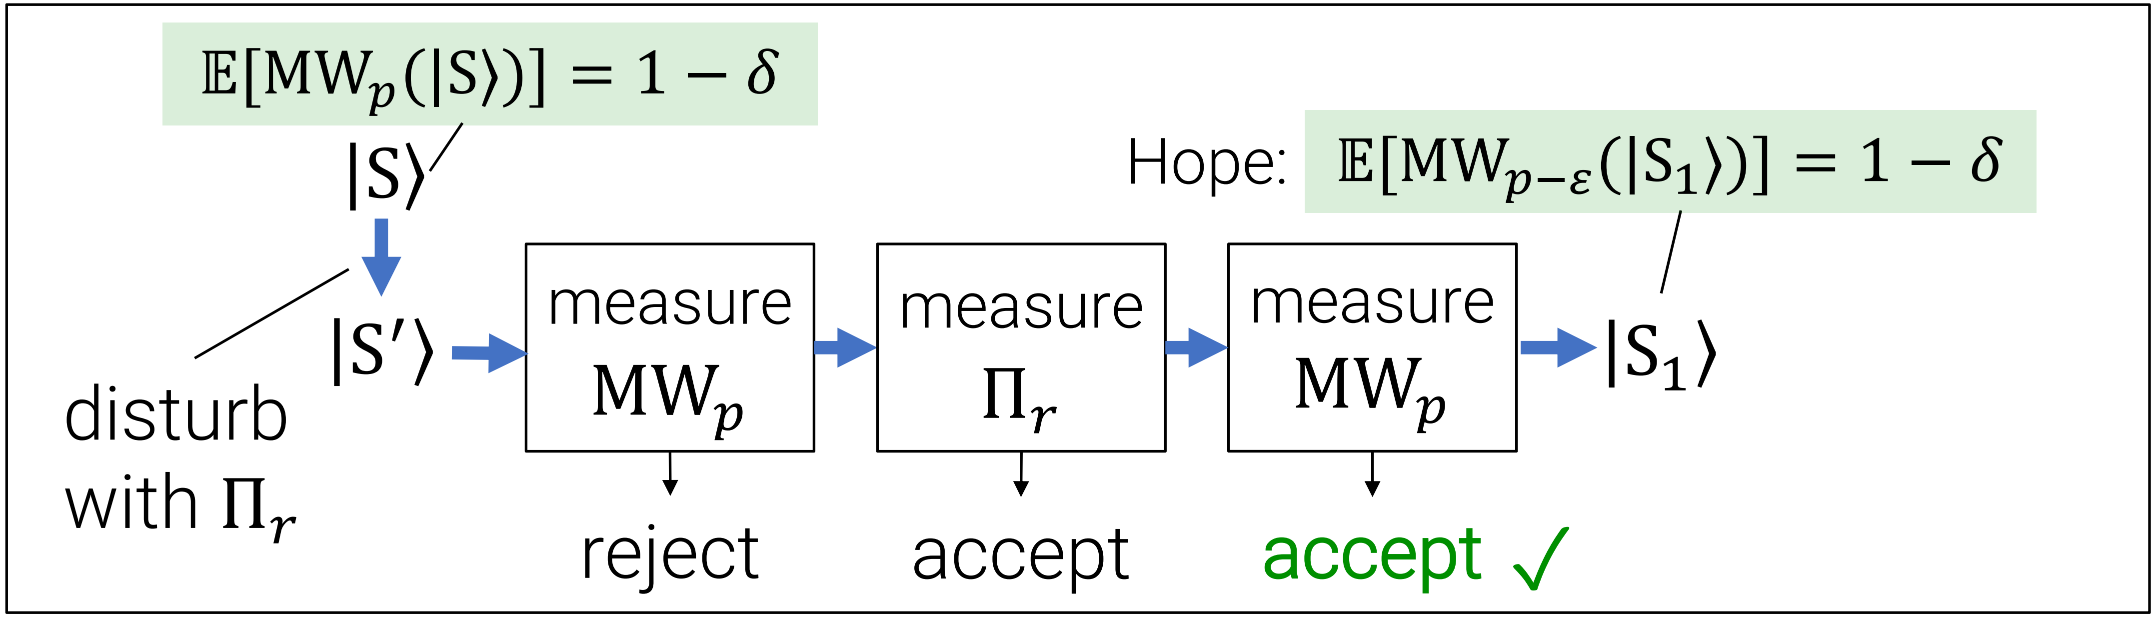

Measuring $|S'\rangle$ with MW_p can be implemented as a projective measurement of some Π_p^* on $|S'\rangle_A|0\rangle_W \in A \otimes W$.

But we need to be careful: the outcome of measuring Π_p^* only corresponds to MW_p when the W register is $|0\rangle$.

(even if we start with $|S'\rangle_A|0\rangle_W$, measuring Π_p^* once may ruin W)

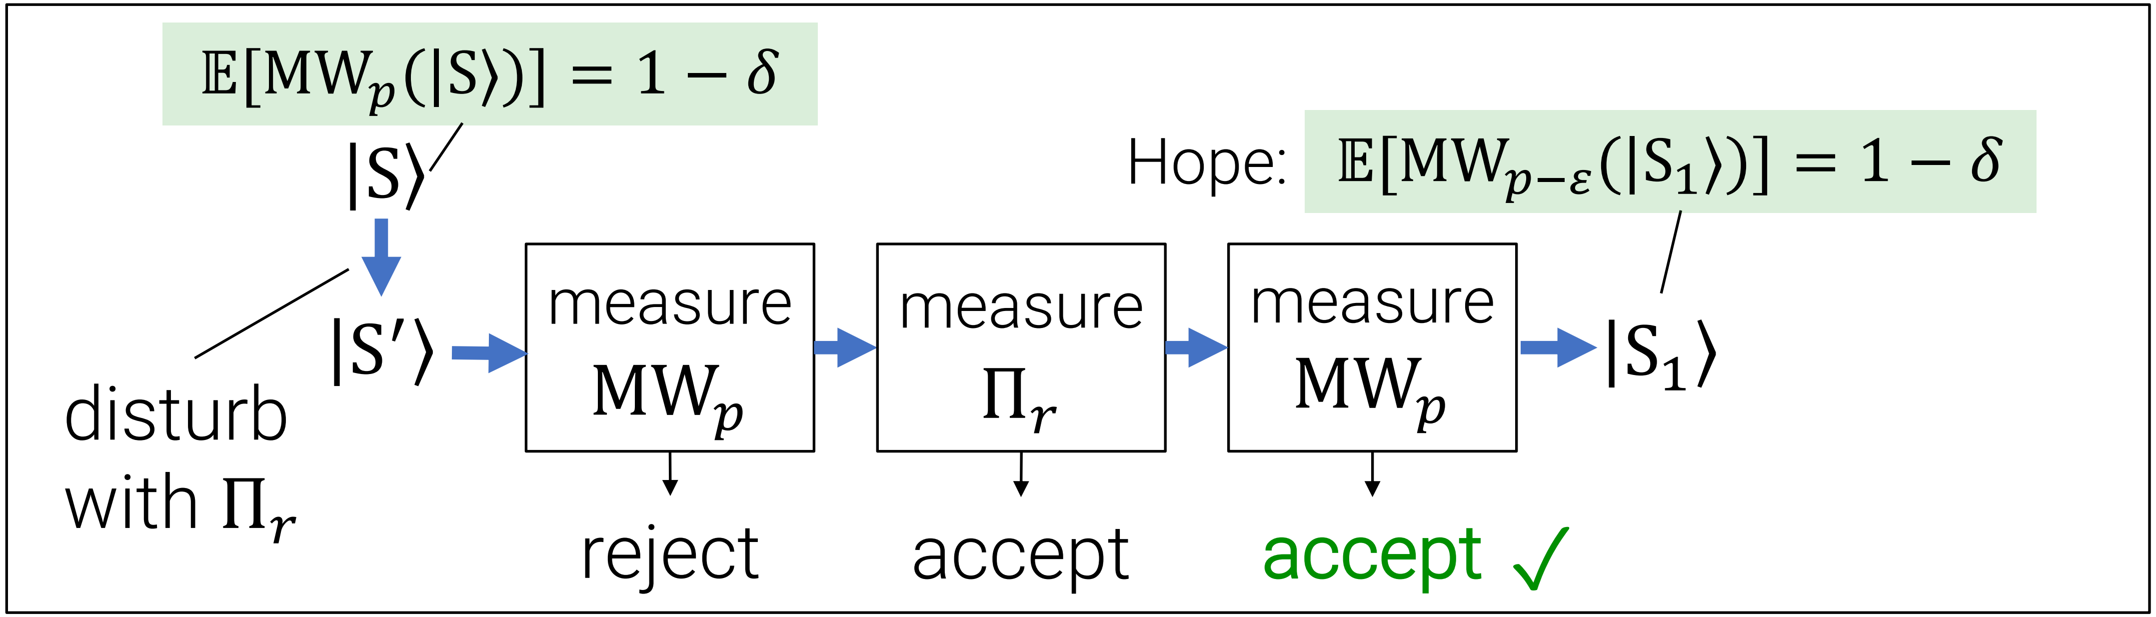

Our solution is re-define Π_r to $\Pi_r^* \coloneqq \Pi_r \otimes |0\rangle\langle 0|_W$, so that each measurement of Π_r^* attempts to “reset” the W to $|0\rangle_W$.

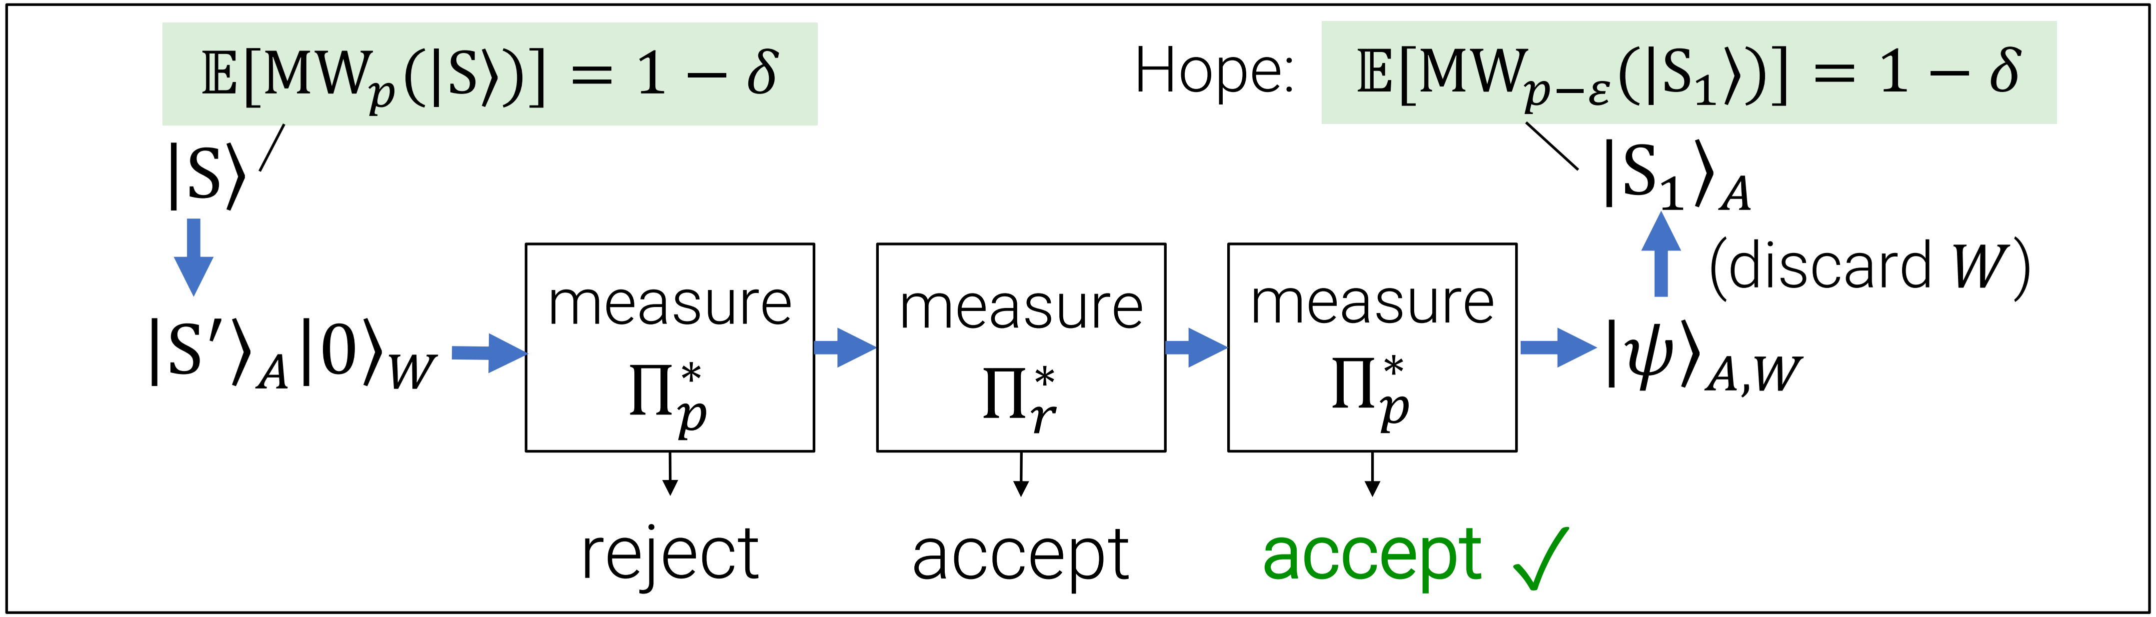

This is essentially the full repair procedure!

Our solution is re-define Π_r to $\Pi_r^* \coloneqq \Pi_r \otimes |0\rangle\langle 0|_W$, so that each measurement of Π_r^* attempts to “reset” the W to $|0\rangle_W$.

However, proving that we satisfy $\mathbb{E}[\mathrm{MW}_{p-\varepsilon}(|\mathrm{S}_1\rangle)] = 1-\delta$ requires more work (in fact, we get a weaker guarantee).

However, proving that we satisfy $\mathbb{E}[\mathrm{MW}_{p-\varepsilon}(|\mathrm{S}_1\rangle)] = 1 - \delta$ requires more work (in fact, we get a weaker guarantee).

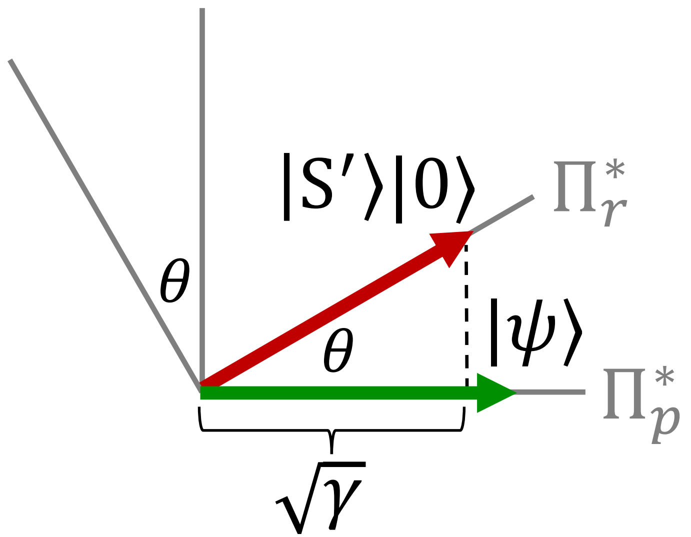

In 2-D, the guarantee depends on

$$\gamma = \cos^2\theta = \mathbb{E}[\mathrm{MW}_p(|S'\rangle)].$$

However, proving that we satisfy $\mathbb{E}[\mathrm{MW}_{p-\varepsilon}(|\mathrm{S}_1\rangle)] = 1 - \delta$ requires more work (in fact, we get a weaker guarantee).

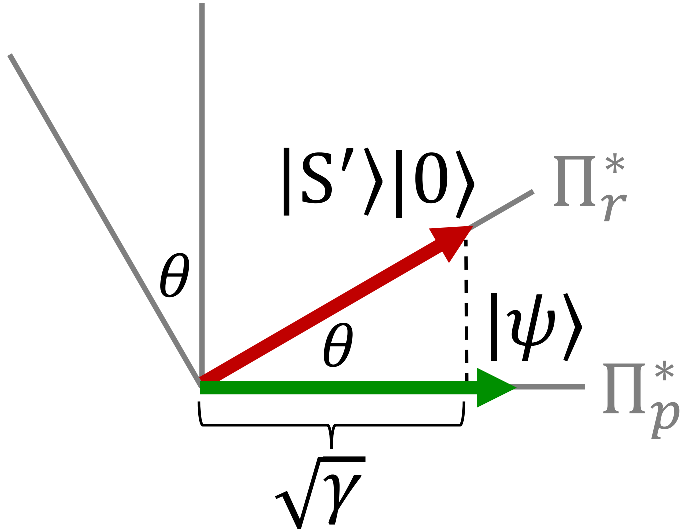

In 2-D, the guarantee depends on

$$\gamma = \cos^2\theta = \mathbb{E}[\mathrm{MW}_p(|S'\rangle)].$$

Repair outputs $|\mathrm{S}_1\rangle = \mathrm{Tr}_W(|\psi\rangle\langle\psi|)$ where

$$\mathbb{E}[\mathrm{MW}_{p-\varepsilon}(|\mathrm{S}_1\rangle)] \geq 1 - \delta/\gamma$$

However, proving that we satisfy $\mathbb{E}[\mathrm{MW}_{p-\varepsilon}(|\mathrm{S}_1\rangle)] = 1 - \delta$ requires more work (in fact, we get a weaker guarantee).

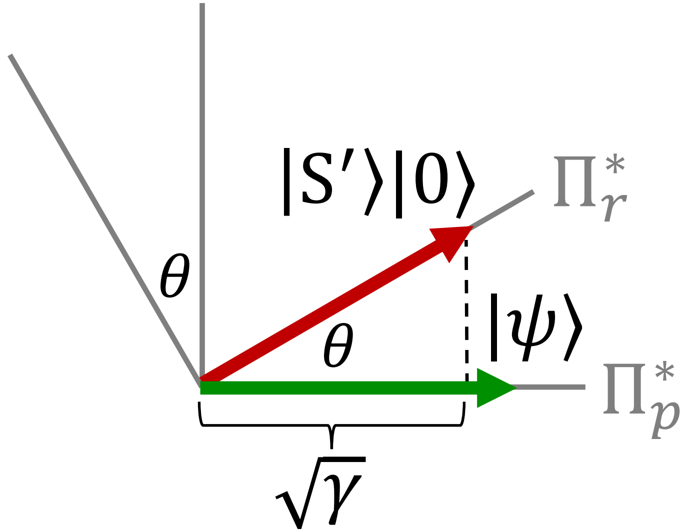

In 2-D, the guarantee depends on

$$\gamma = \cos^2\theta = \mathbb{E}[\mathrm{MW}_p(|S'\rangle)].$$

Repair outputs $|\mathrm{S}_1\rangle = \mathrm{Tr}_W(|\psi\rangle\langle\psi|)$ where

$$\mathbb{E}[\mathrm{MW}_{p-\varepsilon}(|\mathrm{S}_1\rangle)] \geq 1 - \delta/\gamma$$

For the general case, we use Jordan's lemma and prove that on most 2-D subspaces, $\gamma = \mathbb{E}[\mathrm{MW}_p(|S'\rangle)]$ is not too small (since we had $\mathbb{E}[\mathrm{MW}_p(|\mathrm{S}\rangle)] = 1 - \delta$ before disturbance).

Recap: The Full Rewinding Procedure

Recap: The Full Rewinding Procedure

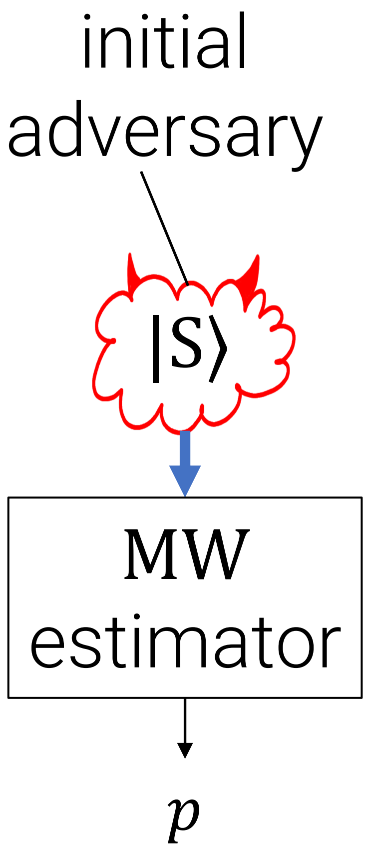

Recap: The Full Rewinding Procedure

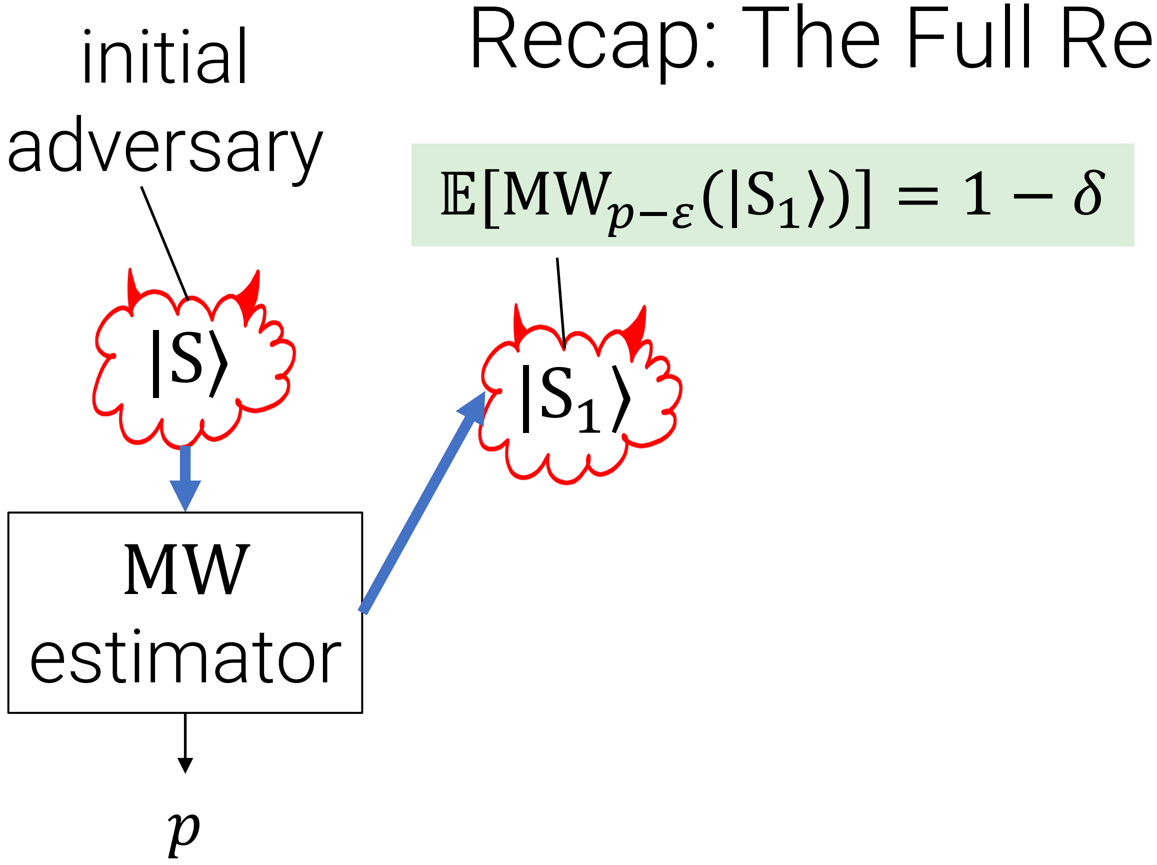

Recap: The Full Rewinding Procedure

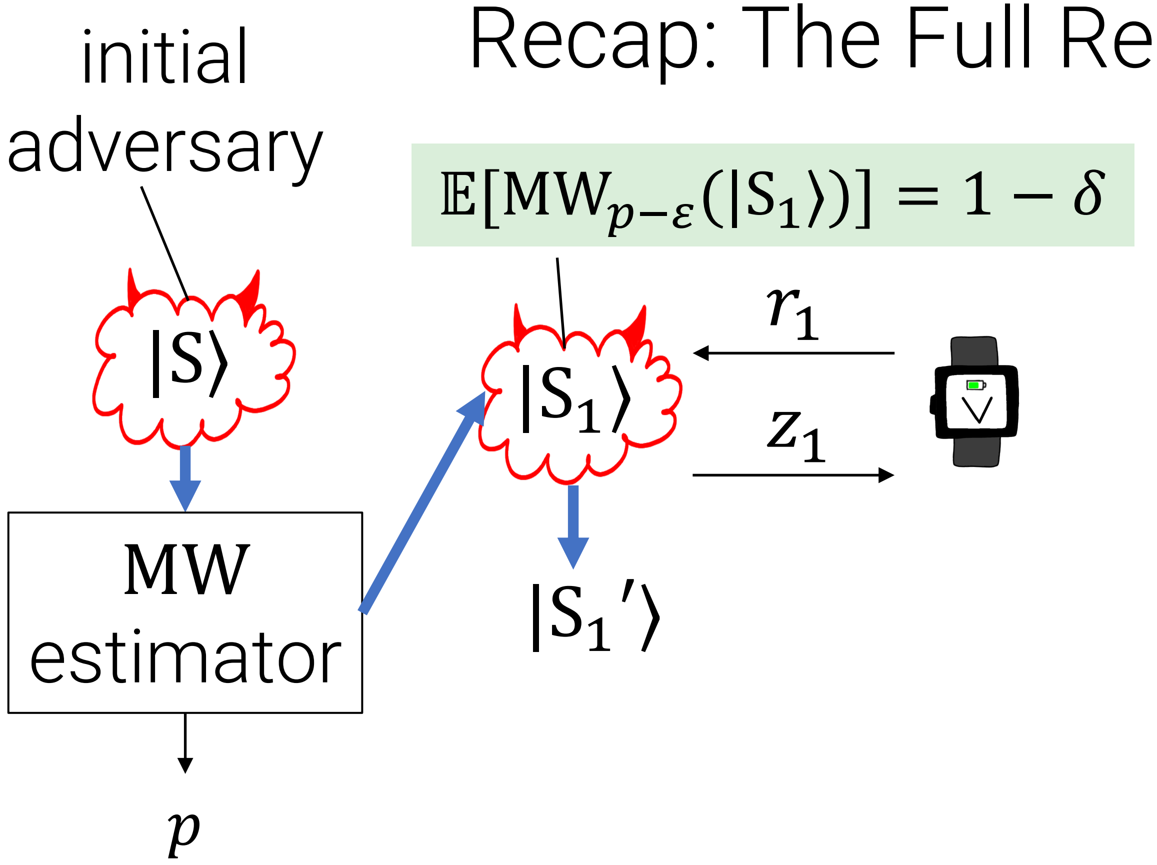

Recap: The Full Rewinding Procedure

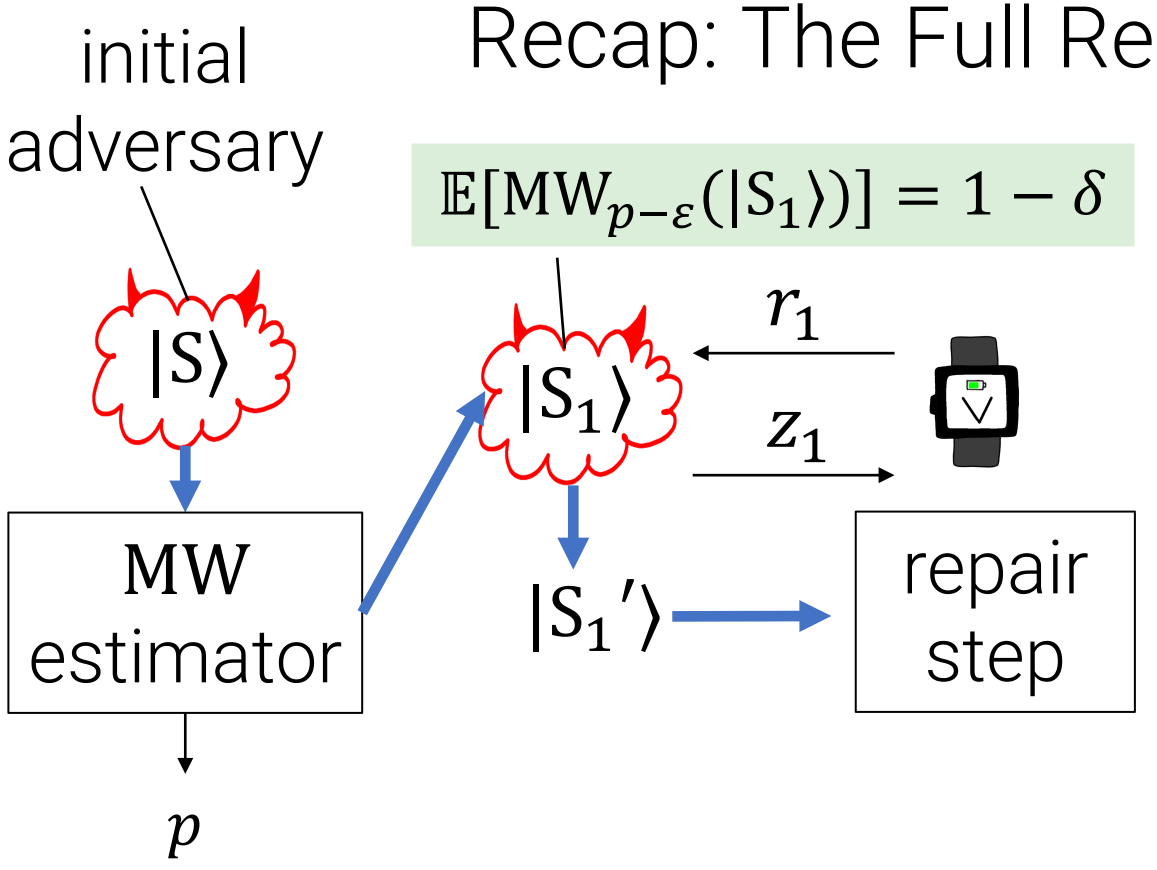

Recap: The Full Rewinding Procedure

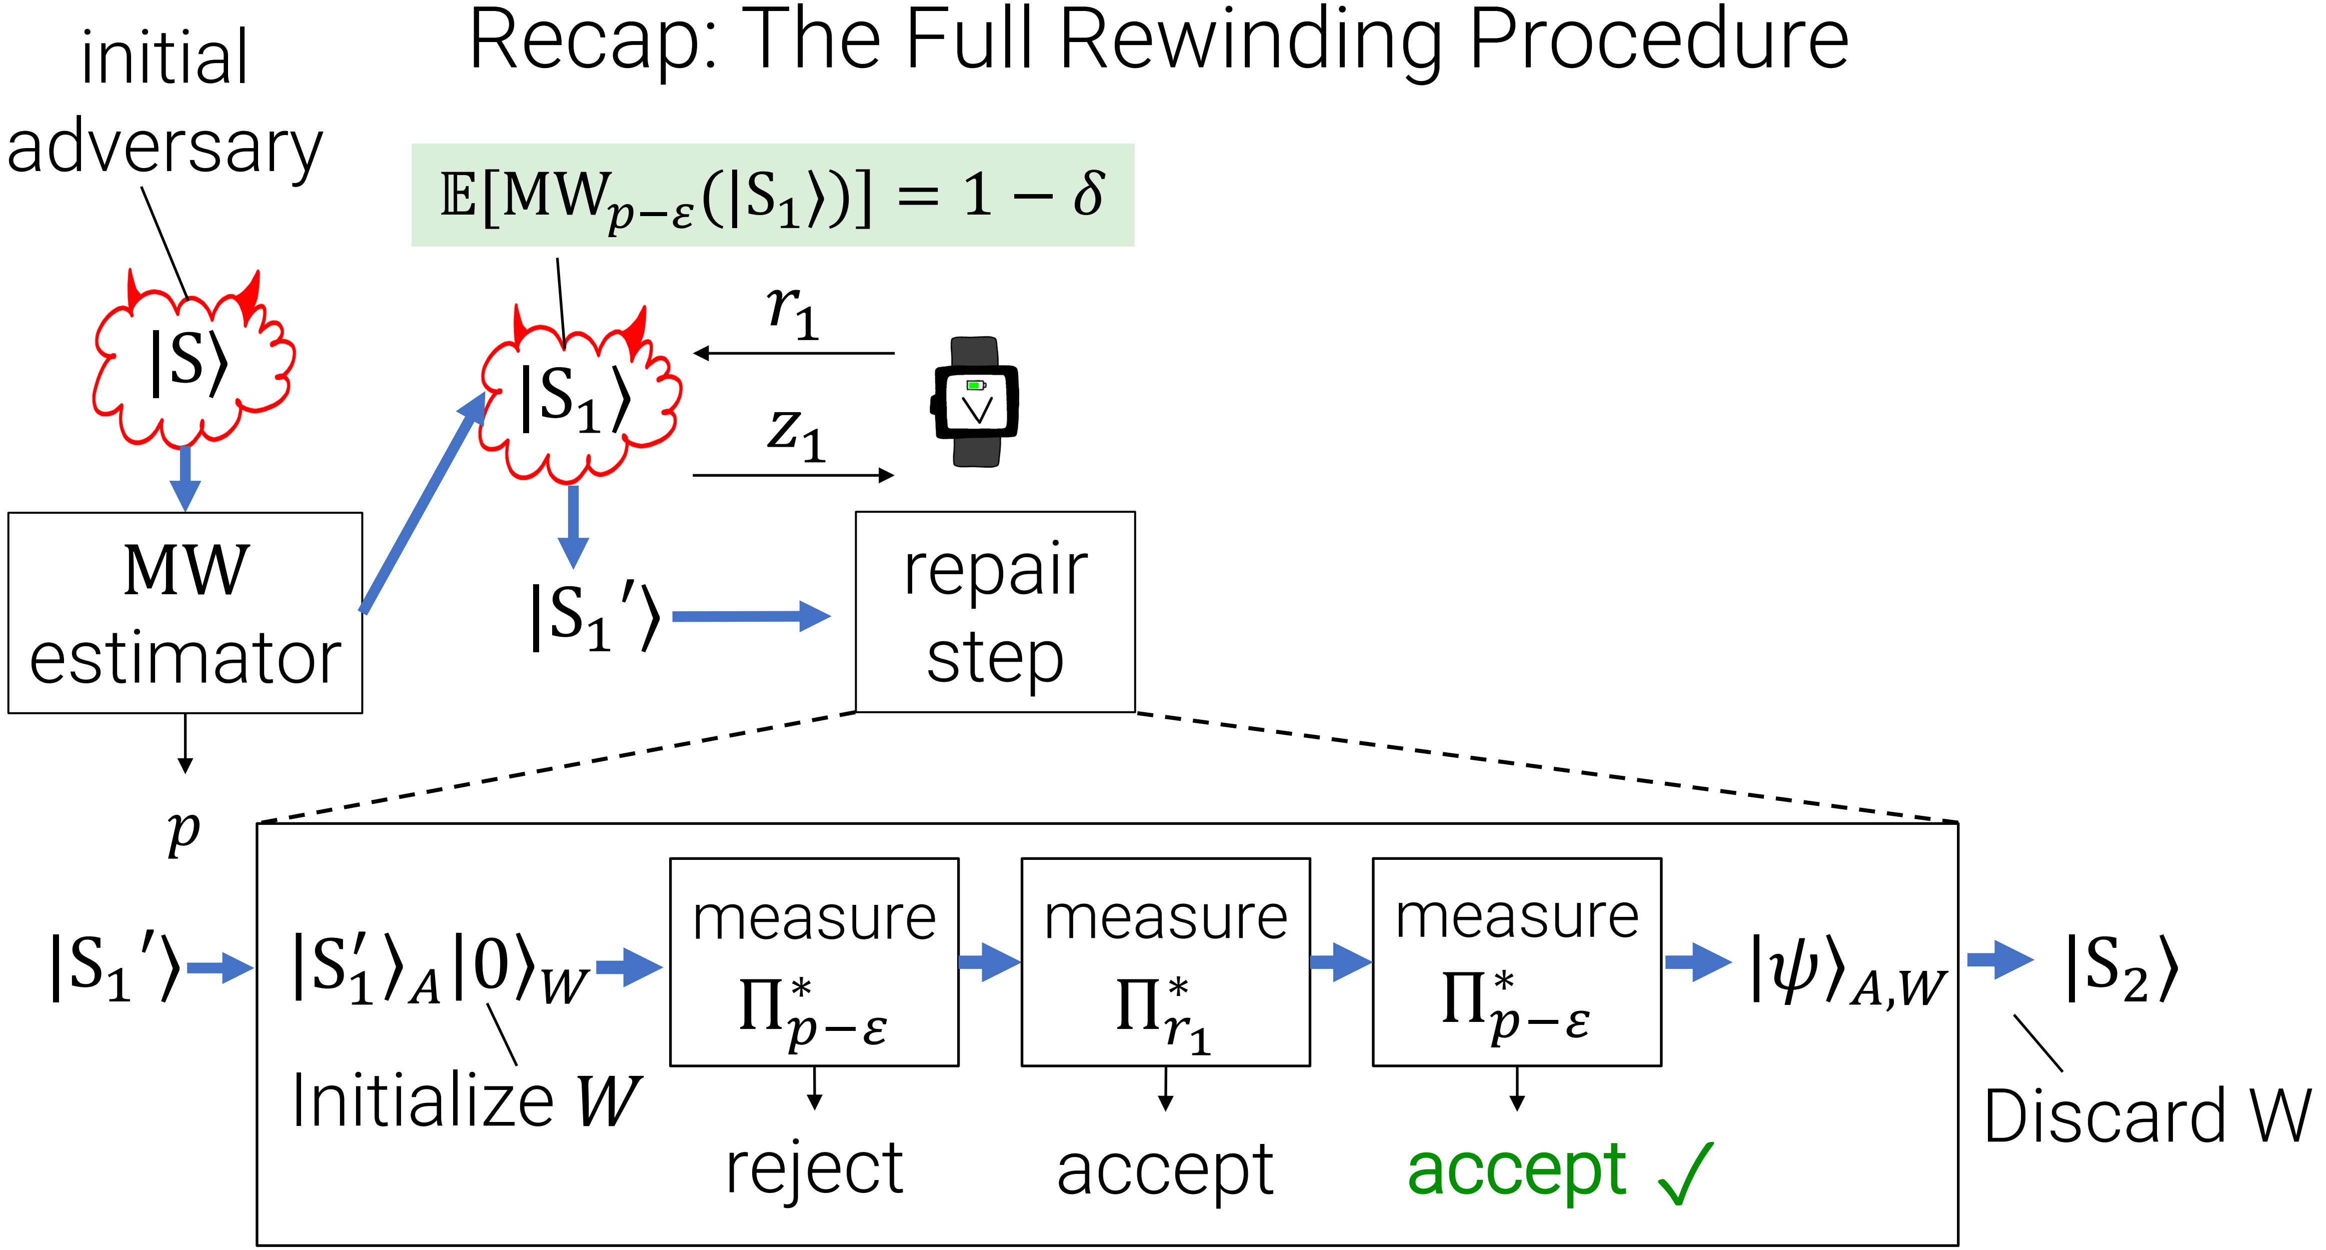

Recap: The Full Rewinding Procedure

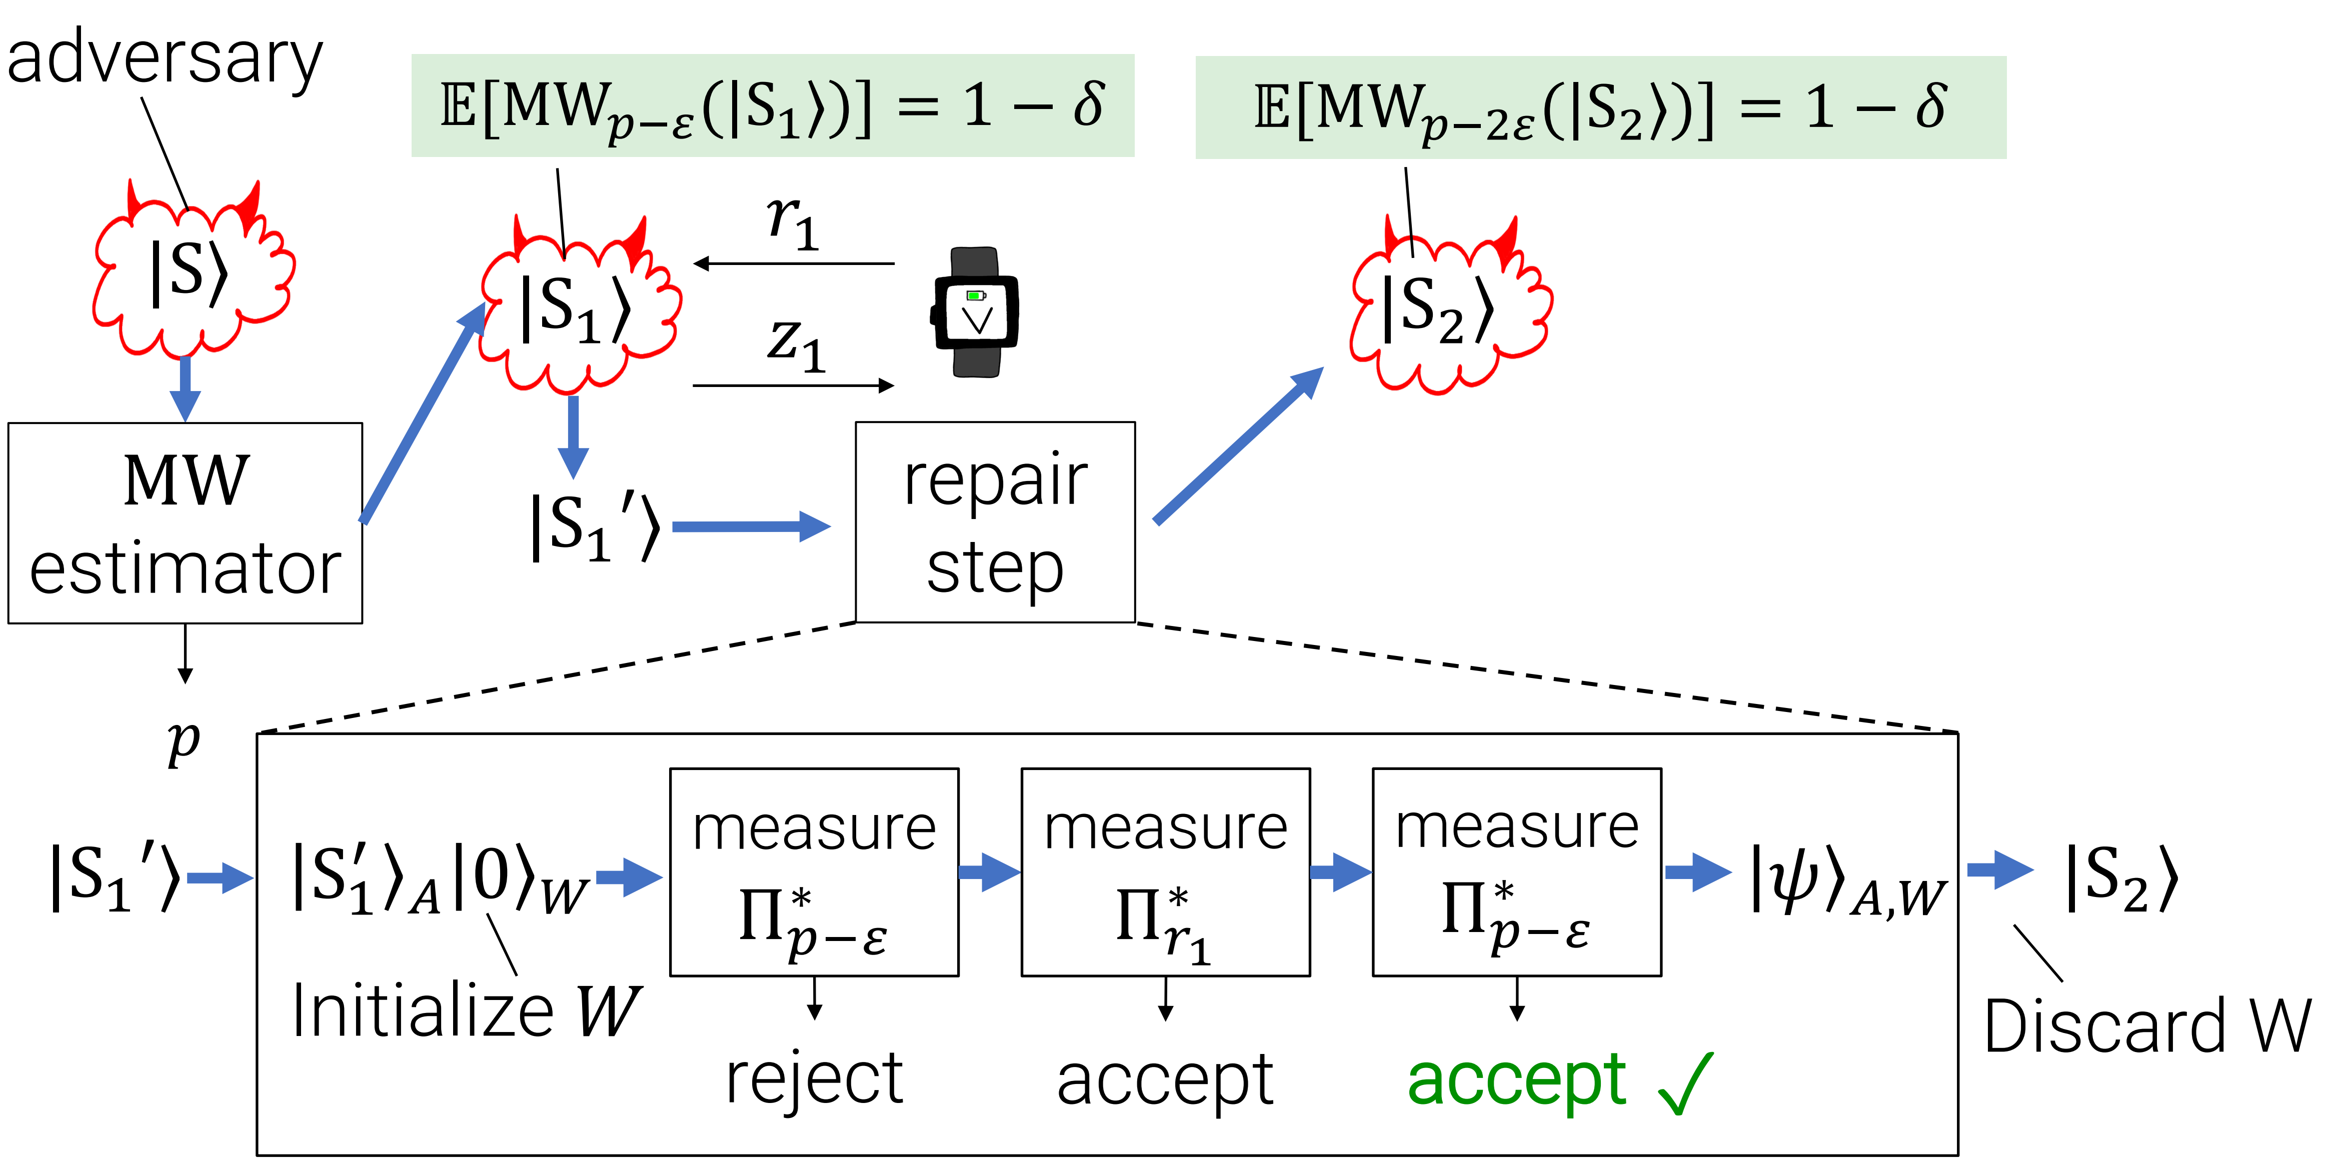

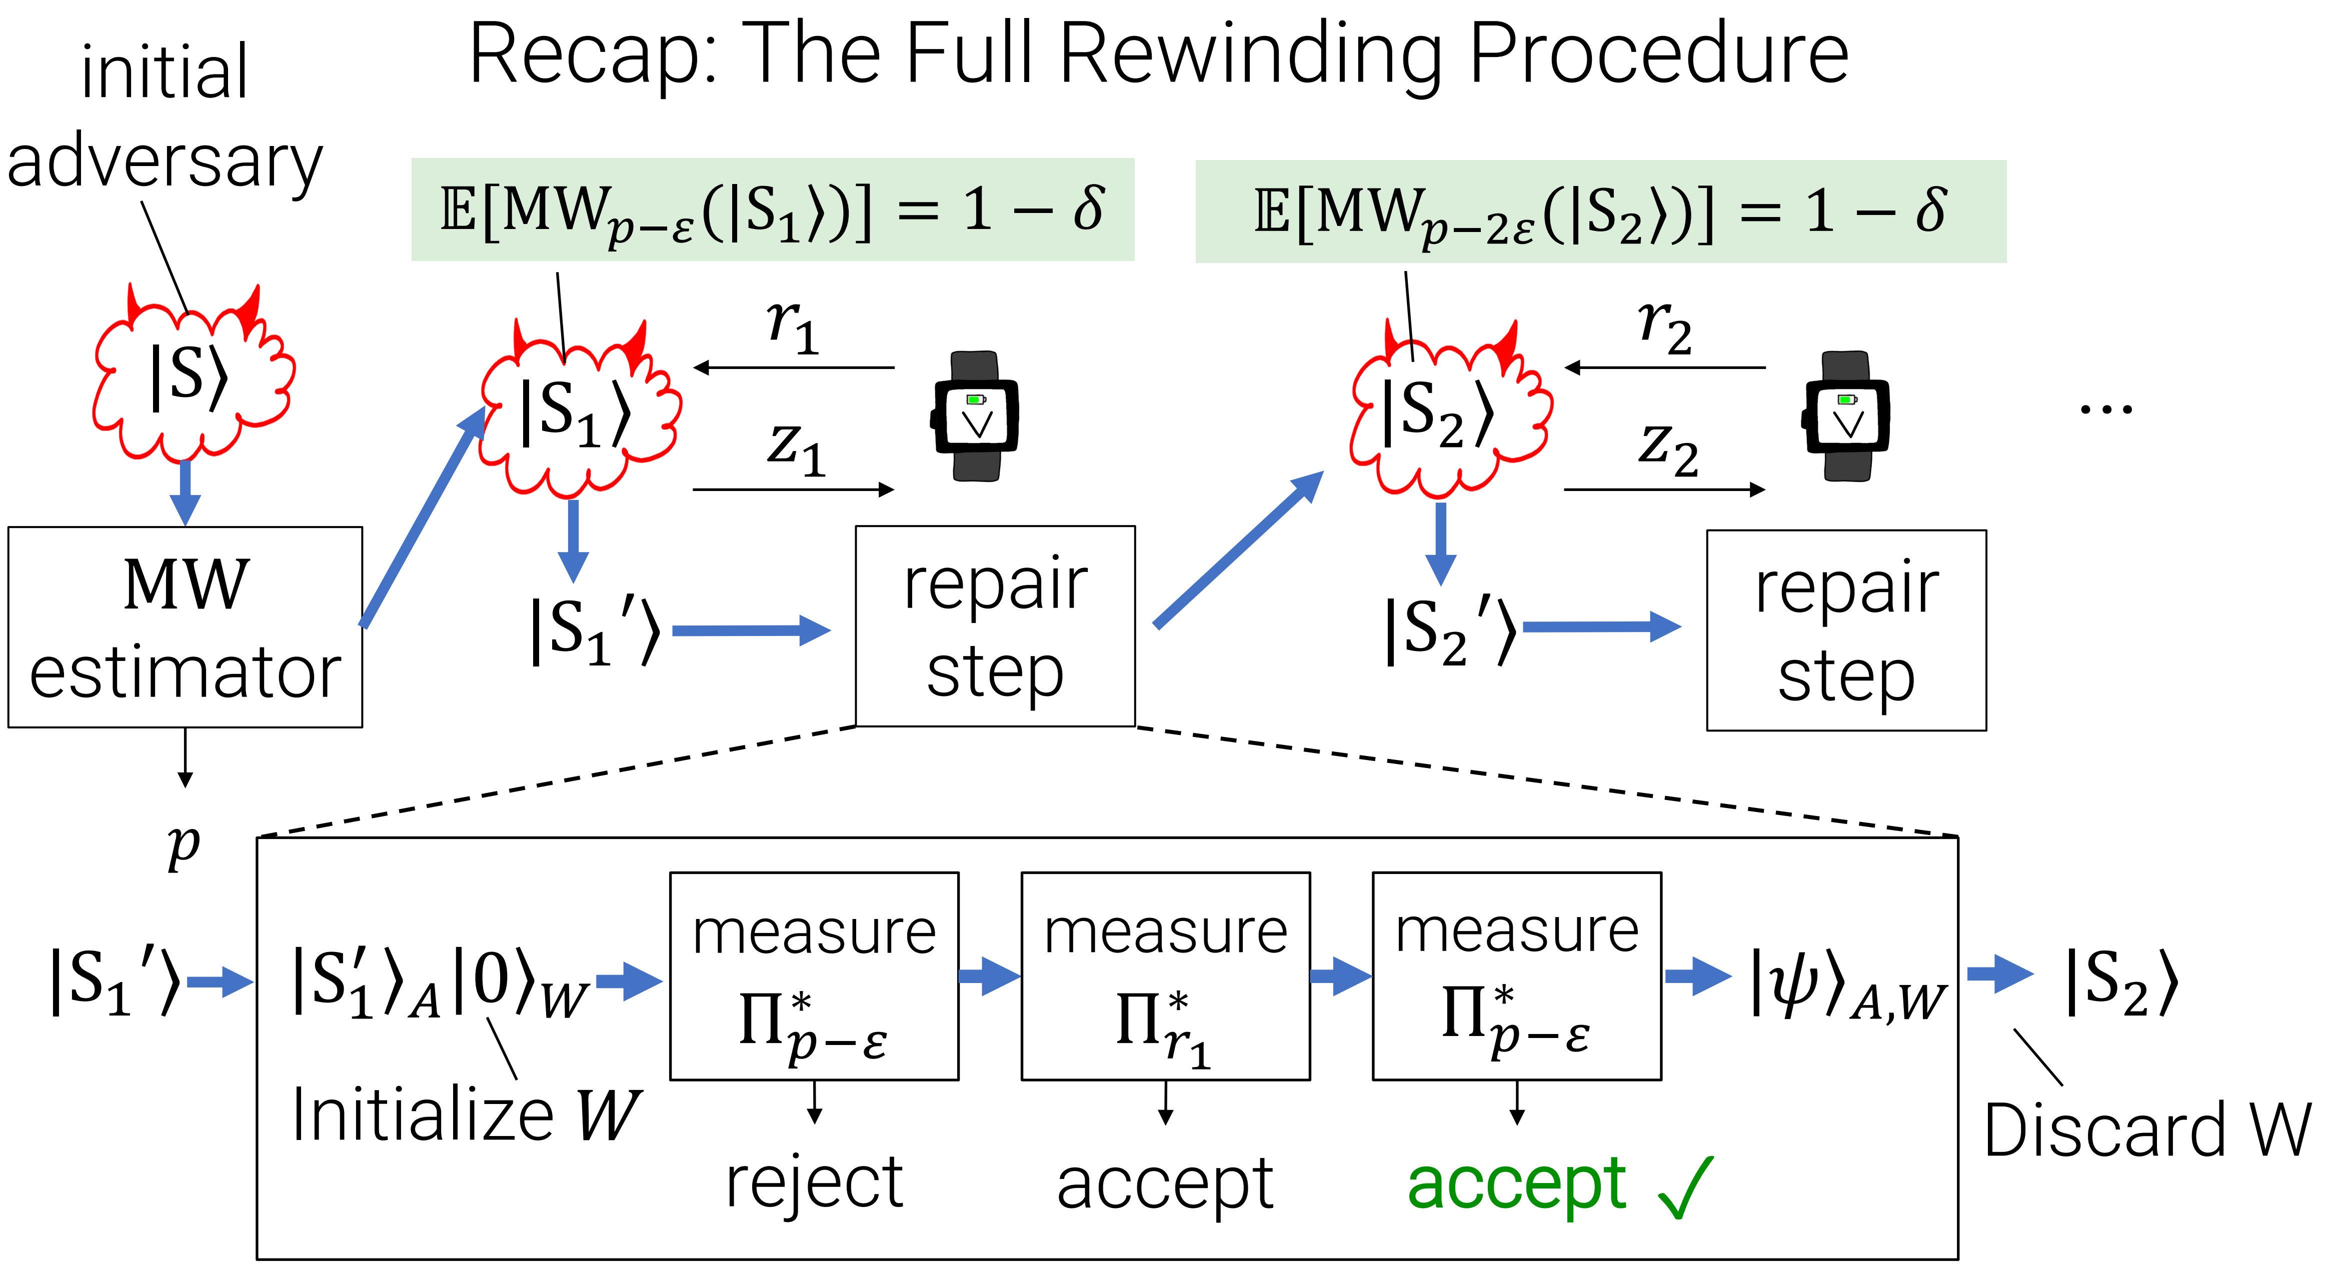

Recap: The Full Rewinding Procedure
initial adversary
$|S\rangle$
MW estimator
p
$\mathbb{E}[MW_{p-\varepsilon}(|S_1\rangle)] = 1 - \delta$
$|S_1\rangle$
r_1
z_1
$|S_1'\rangle$
repair step
$\mathbb{E}[MW_{p-2\varepsilon}(|S_2\rangle)] = 1 - \delta$
$|S_2\rangle$
r_2
z_2
$|S_2'\rangle$
repair step
...
$|S_1'\rangle$
$|S_1'\rangle_A|0\rangle_W$
Initialize W
measure $\Pi^*_{p-\varepsilon}$
reject
measure $\Pi^*_{r_1}$
accept
measure $\Pi^*_{p-\varepsilon}$
accept ✓
$|\psi\rangle_{A,W}$
$|S_2\rangle$
Discard W

Conclusions

- Much of cryptography deals with *interactive protocols*. In this setting, security is fragile in the presence of quantum adversaries because classical rewinding is inapplicable.

Conclusions

- Much of cryptography deals with *interactive protocols*. In this setting, security is fragile in the presence of quantum adversaries because classical rewinding is inapplicable.
- Rewinding is often used to *record an adversary's responses* across multiple challenges.

Conclusions

- Much of cryptography deals with *interactive protocols*. In this setting, security is fragile in the presence of quantum adversaries because classical rewinding is inapplicable.
- Rewinding is often used to *record an adversary's responses* across multiple challenges.
- We address this issue by solving an abstract problem: if a stateful quantum adversary wins a challenge-response game once, we extend it to win the game many times.

Conclusions

- Much of cryptography deals with *interactive protocols*. In this setting, security is fragile in the presence of quantum adversaries because classical rewinding is inapplicable.
- Rewinding is often used to *record an adversary's responses* across multiple challenges.
- We address this issue by solving an abstract problem: if a stateful quantum adversary wins a challenge-response game once, we extend it to win the game many times.
- **Next steps:** other use cases for rewinding? We give some answers in upcoming work [LMS21].

Thank You!

Questions?

Slide Artwork by Eysa Lee