Efficient quantum parallel repetition and applications

John Bostanci

joint work with Luowen Qian, Nick Spooner, and Henry Yuen

TLDR: The computational security of 3-message quantum interactive protocols can be amplified by parallel repetition.

Definitions & Examples

Motivating example: commitments

A quantum bit commitment is the cryptographic equivalent to a message in a locked box.

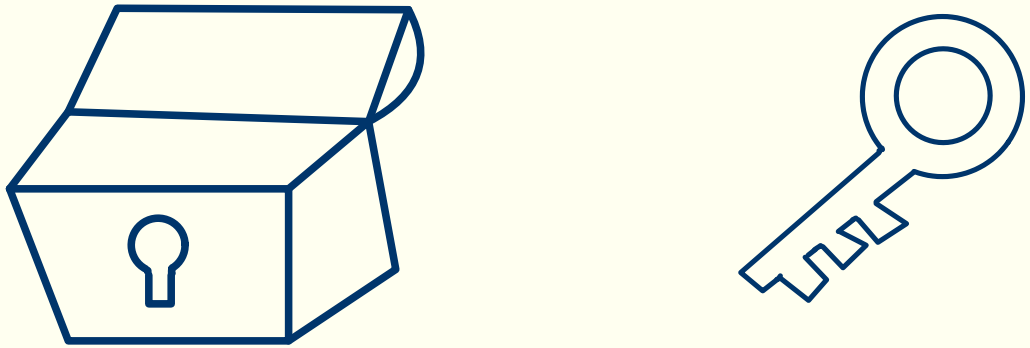

Motivating example: commitments

Quantum bit commitments are central to quantum cryptography!

- Implied by almost all other cryptographic primitives.
- Equivalent to Uhlmann transformations.
- Equivalent to Harlow-Hayden black hole radiation decoding.

Motivating example: commitments

There are two phases in a quantum bit commitment, a commit phase and a reveal phase.

Motivating example: commitments

In the commit phase, a sender generates a bipartite state $|\psi_b\rangle_{RC}$, and sends the commit register C to the receiver.

Motivating example: commitments

In the commit phase, a sender generates a bipartite state $|\psi_b\rangle_{RC}$, and sends the commit register C to the receiver.

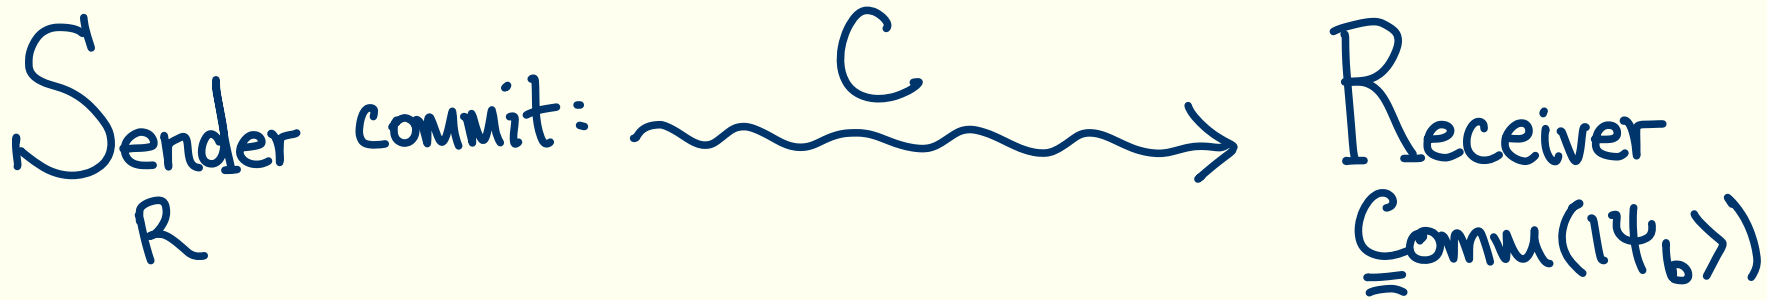

Motivating example: commitments

In the reveal phase, the sender sends the reveal register R to the receiver, as well as the bit b. The receiver measures $|\psi_b\rangle\langle\psi_b|$ to confirm.

Motivating example: commitments

In the reveal phase, the sender sends the reveal register R to the receiver, as well as the bit b. The receiver measures $|\psi_b\rangle\langle\psi_b|$ to confirm.

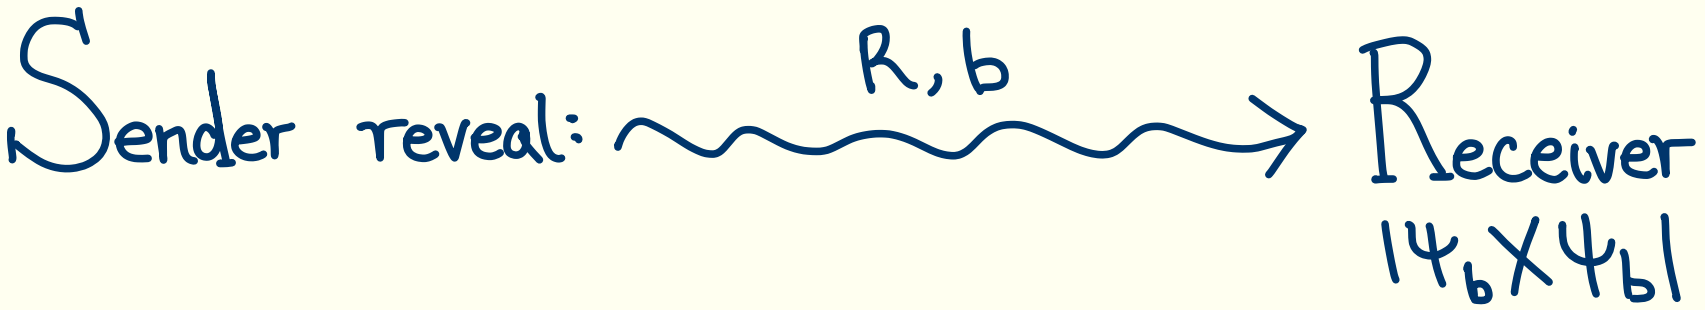

Motivating example: commitments

Security: Once the sender has committed to b, they should not be able to send $1 - b$ and have the receiver successfully measure $|\psi_{1-b}\rangle\langle\psi_{1-b}|$.

Sender
R

Receiver
Comm($|\psi_b\rangle$)

Motivating example: commitments

Security: Once the sender has committed to b, they should not be able to send $1 - b$ and have the receiver successfully measure $|\psi_{1-b}\rangle\langle\psi_{1-b}|$.

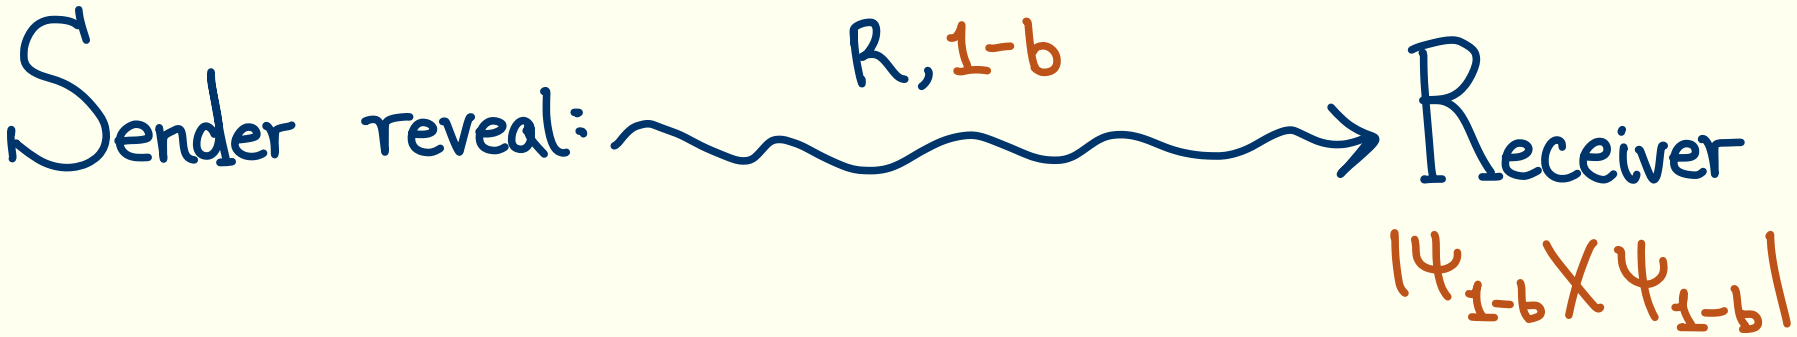

Motivating example: commitments

Question: How do we characterize whether a proposed scheme has this property?

Motivating example: commitments

Question: How do we characterize whether a proposed scheme has this property?

Answer: We define a separate computationally sound protocol called a security game.

Motivating example: binding security game

In the binding security game, a challenger first generates $|\psi_0\rangle$, and sends the reveal register to the adversary.

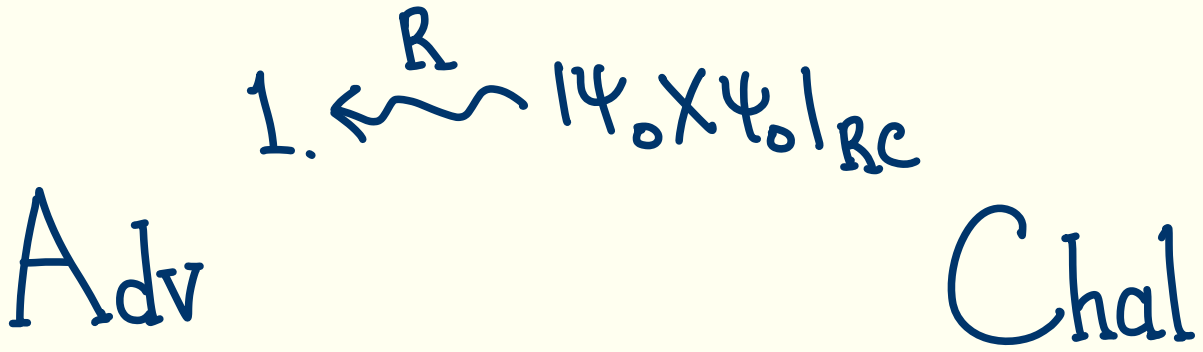

Motivating example: binding security game

The adversary performs some unitary and sends the reveal register back to the challenger.

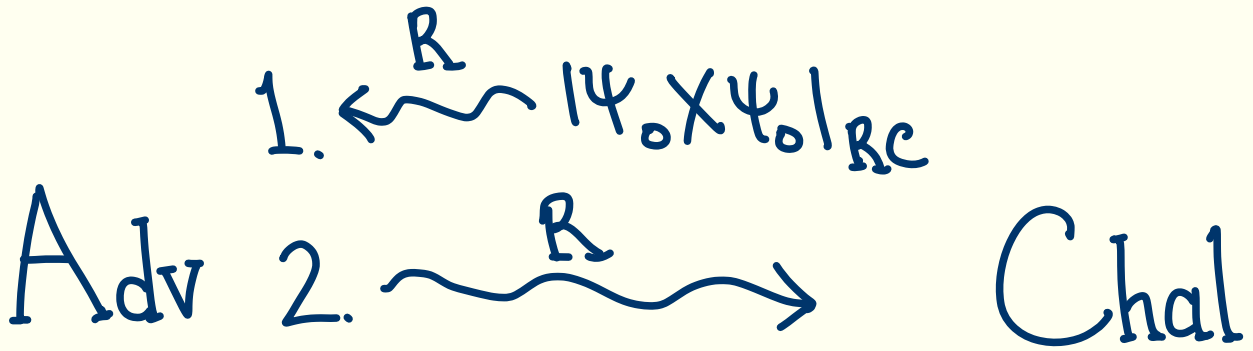

Motivating example: binding security game

The adversary wins the game if the challenger now measures $|\psi_1\rangle\langle\psi_1|$.

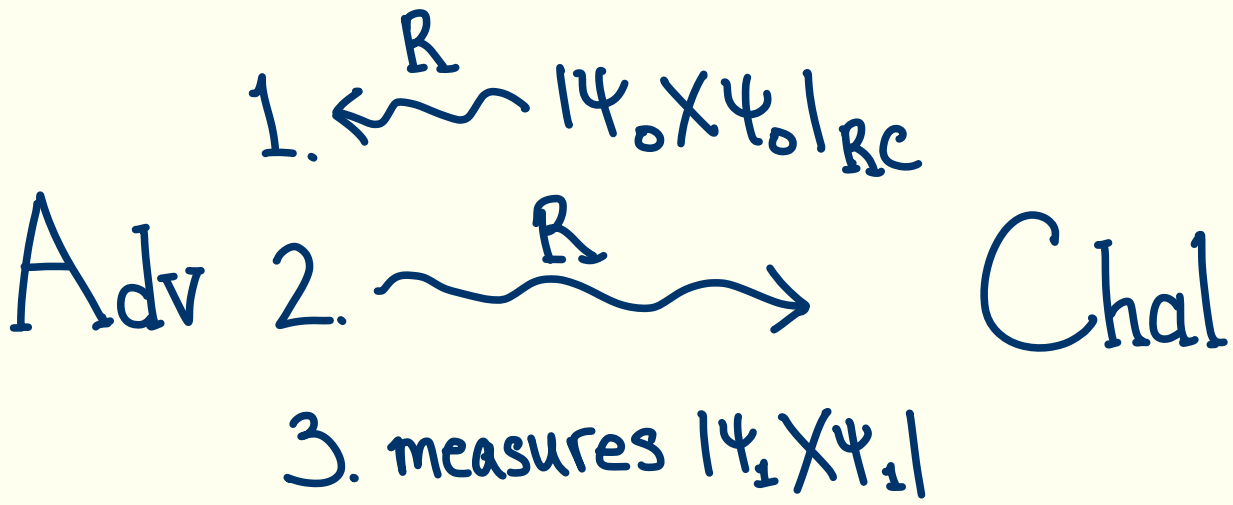

Motivating example: binding security game

We say that a quantum bit commitment is ϵ-binding if the probability that any efficient adversary can win the binding security game is at most ϵ.

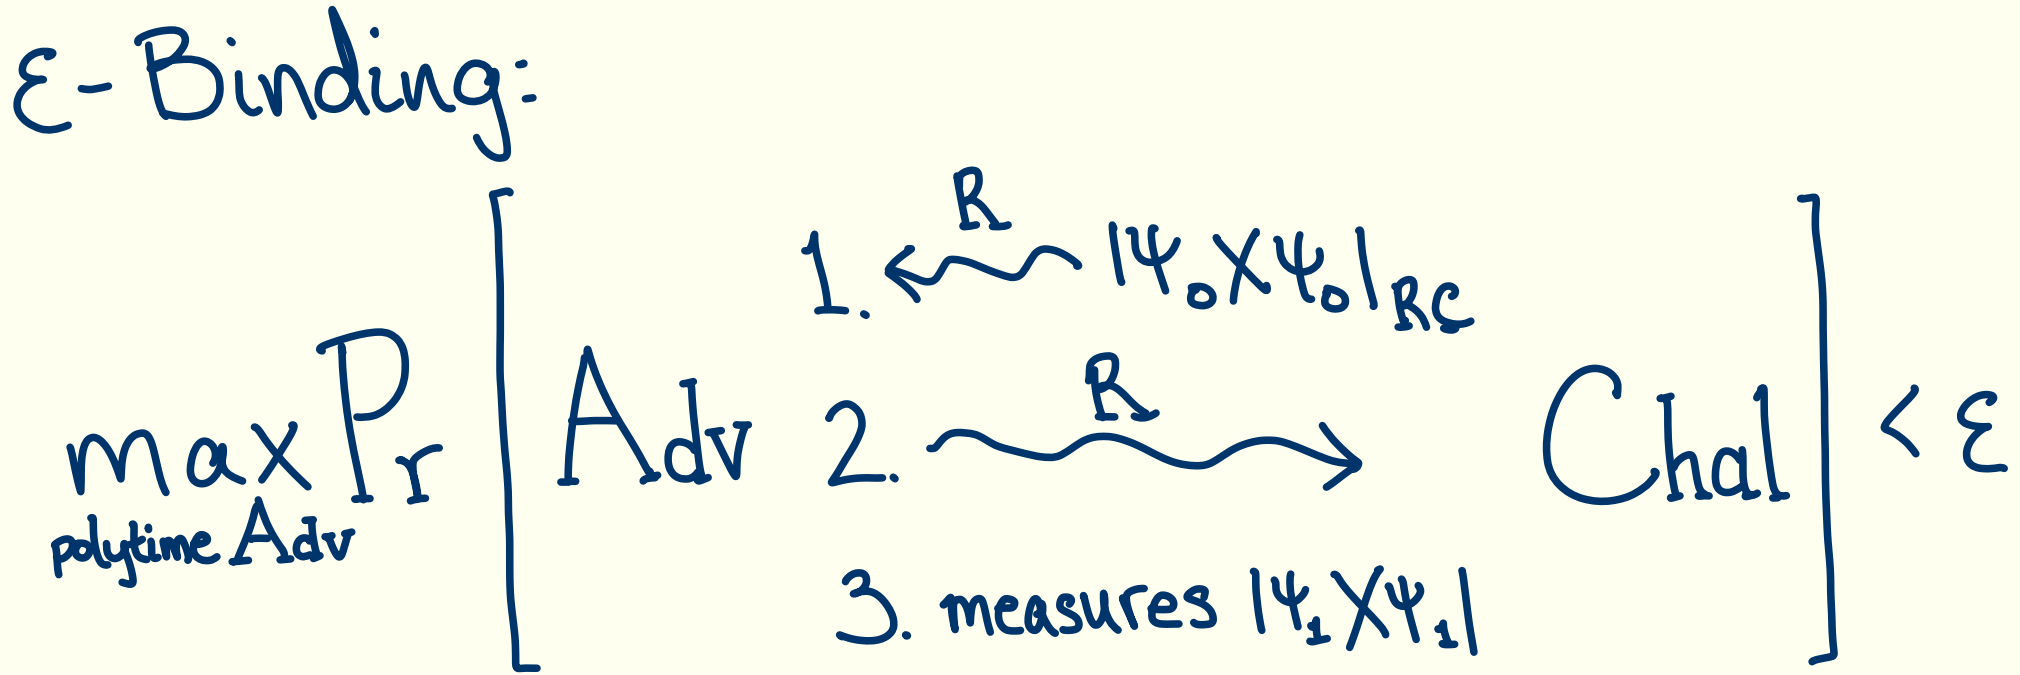

Aside: soundness

For a general computationally secure quantum interactive protocol, we call ϵ the "soundness" of the protocol.

$$\text{Sound}(\text{Chal}) := \max_{\text{polytime Adv}} \Pr\left[\begin{array}{l} 1.\ \text{Adv} \rightleftharpoons \text{Chal} \\ 2.\ \text{Chal} \xrightarrow{\text{measure}} \text{Accept} \end{array}\right]$$

A natural question: parallel repetition

Parallel repetition: Say I give you a 0.99-binding quantum bit commitment, and you instead commit using $|\psi_b\rangle^{\otimes 2}$…

Is the resulting commitment 0.99^2 binding?

A natural question: parallel repetition

Intuitively, what could the adversary in the binding security game do, other than try to flip the first copy of $|\psi_0\rangle$, and then try to flip the second copy?

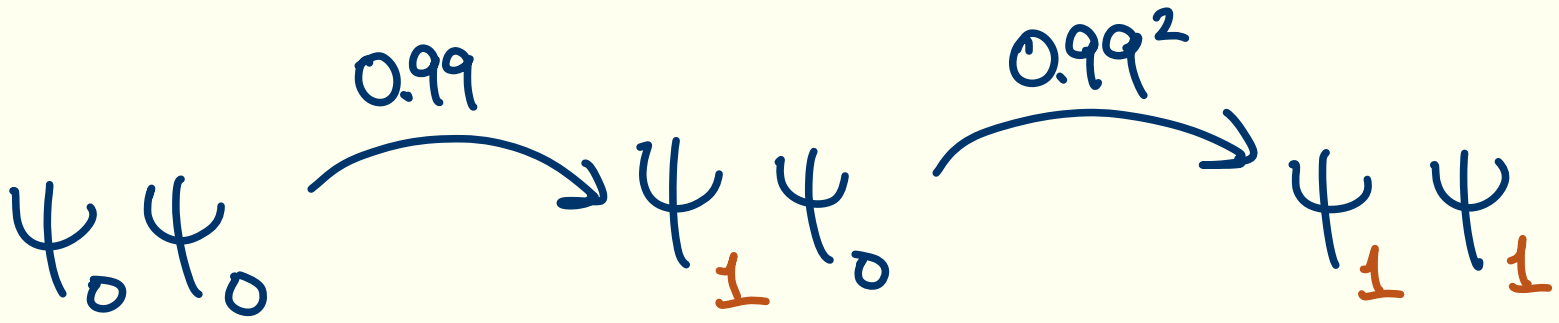

A natural question: parallel repetition

Intuitively, what could the adversary in the binding security game do, other than try to flip the first copy of $|\psi_0\rangle$, and then try to flip the second copy?

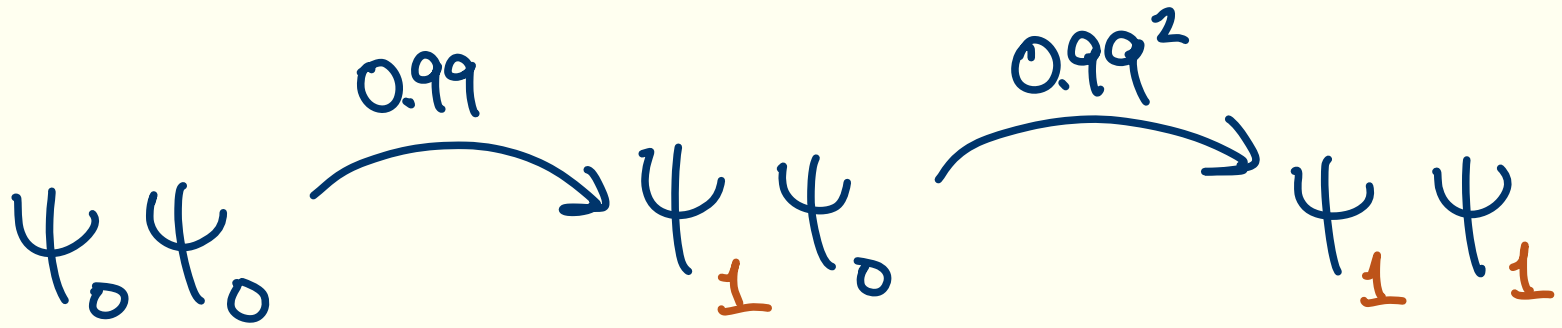

Seems like a simple question, but surprisingly difficult to prove!

Aside: parallel repetition in the wild

Although we're focusing on quantum bit commitments, the question of "does parallel repetition do what you expect" appears in many places:

Aside: parallel repetition in the wild

Although we're focusing on quantum bit commitments, the question of "does parallel repetition do what you expect" appears in many places:

- Strong amplification of bit commitments [Yan22]
- Strong amplification of Uhlmann instances [BEM+23]
- Amplification of quantum money schemes [AC13]
- 4-message (Quantum) ZK proofs for QIP (assuming EFI) [BCQ22]
- Simpler zero knowledge arguments of QMA [BG22]
- Simpler commitments from black holes [Bra23]

Results

Main result

For every 3-message computationally secure quantum interactive protocol with soundness s, the k-fold parallel repetition of the protocol has soundness $s^k + \mathrm{negl}$.

Applications

We get a lot of results for “free” from this:

Applications

We get a lot of results for “free” from this:

- Strong amplification of bit commitments [Yan22] ✓
- Strong amplification of Uhlmann instances [BEM+23] ✓
- Amplification of quantum money schemes [AC13] ✓
- 4-message (Quantum) ZK proofs for QIP (assuming EFI) [BCQ22] ✓
- Simpler zero knowledge arguments of QMA [BG22] ✓
- Simpler commitments from black holes [Bra23] ✓

Back to quantum bit commitments…

Applications: quantum bit commitments

Say that you told me about a quantum bit commitment, and told me that it satisfies 0.99-binding:

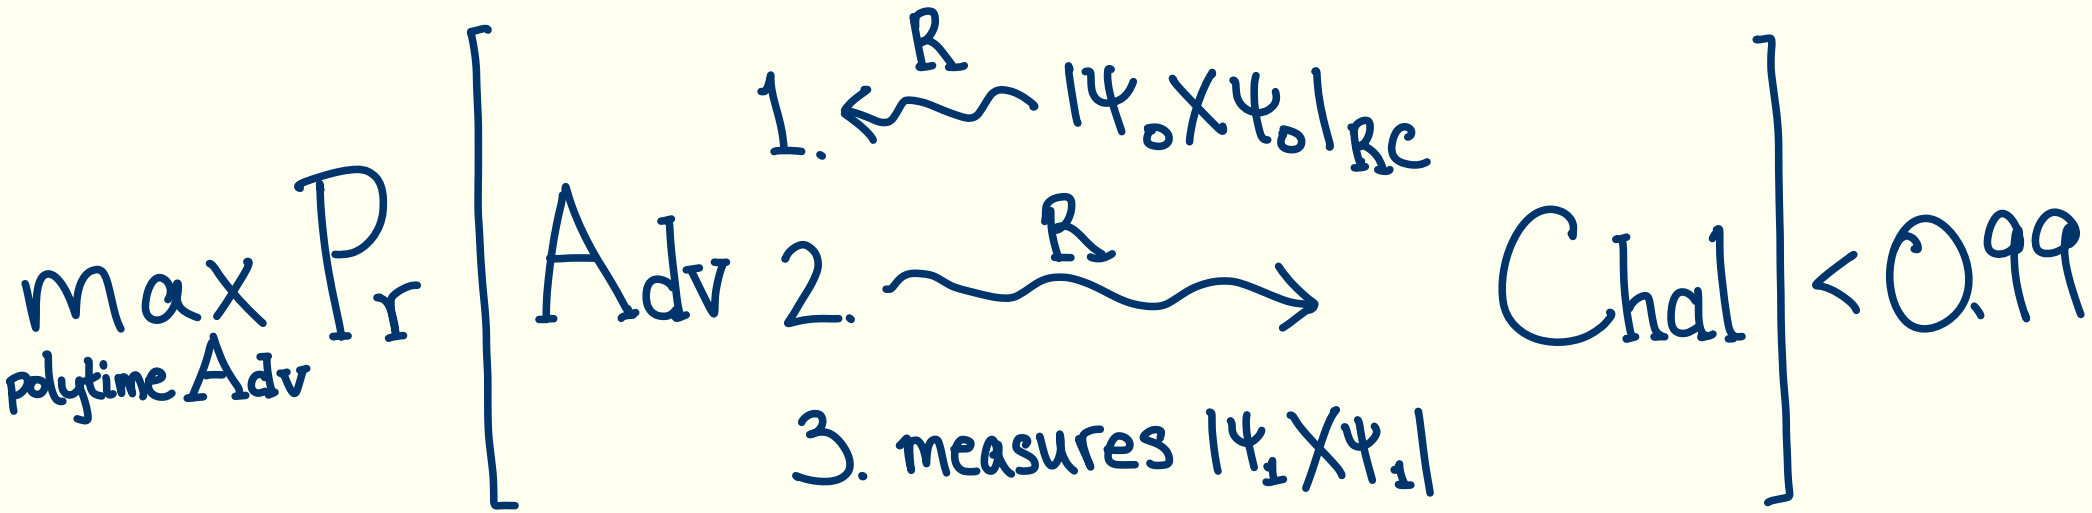

Applications: quantum bit commitments

Then if I want to use a commitment scheme that satisfies negl-binding, I can use $|\psi_b\rangle^{\otimes \mathrm{poly}}$ for some polynomial!

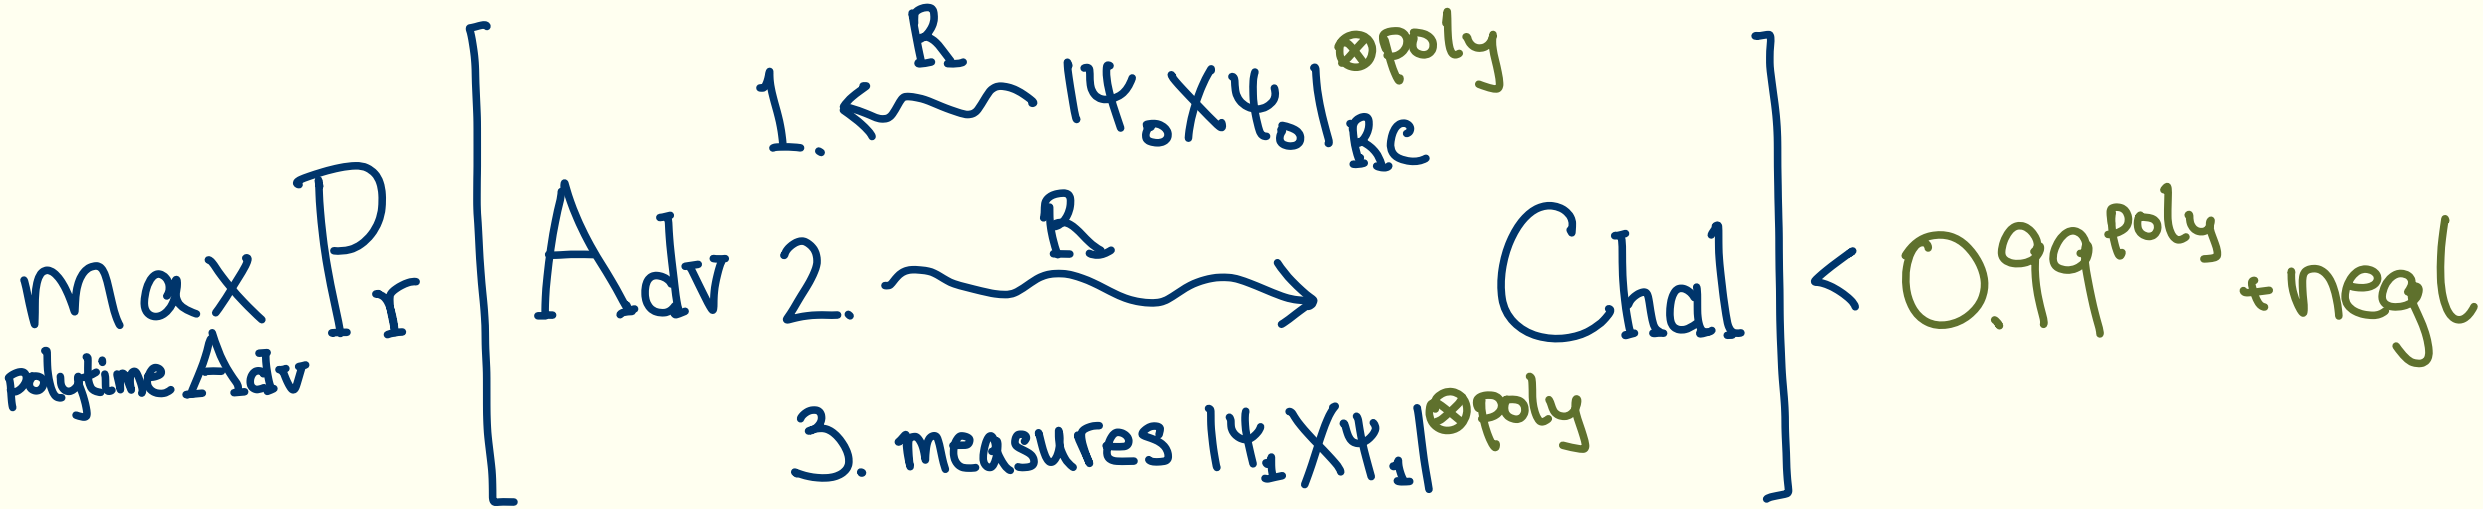

Applications: quantum bit commitments

Then if I want to use a commitment scheme that satisfies negl-binding, I can use $|\psi_b\rangle^{\otimes \mathrm{poly}}$ for some polynomial!

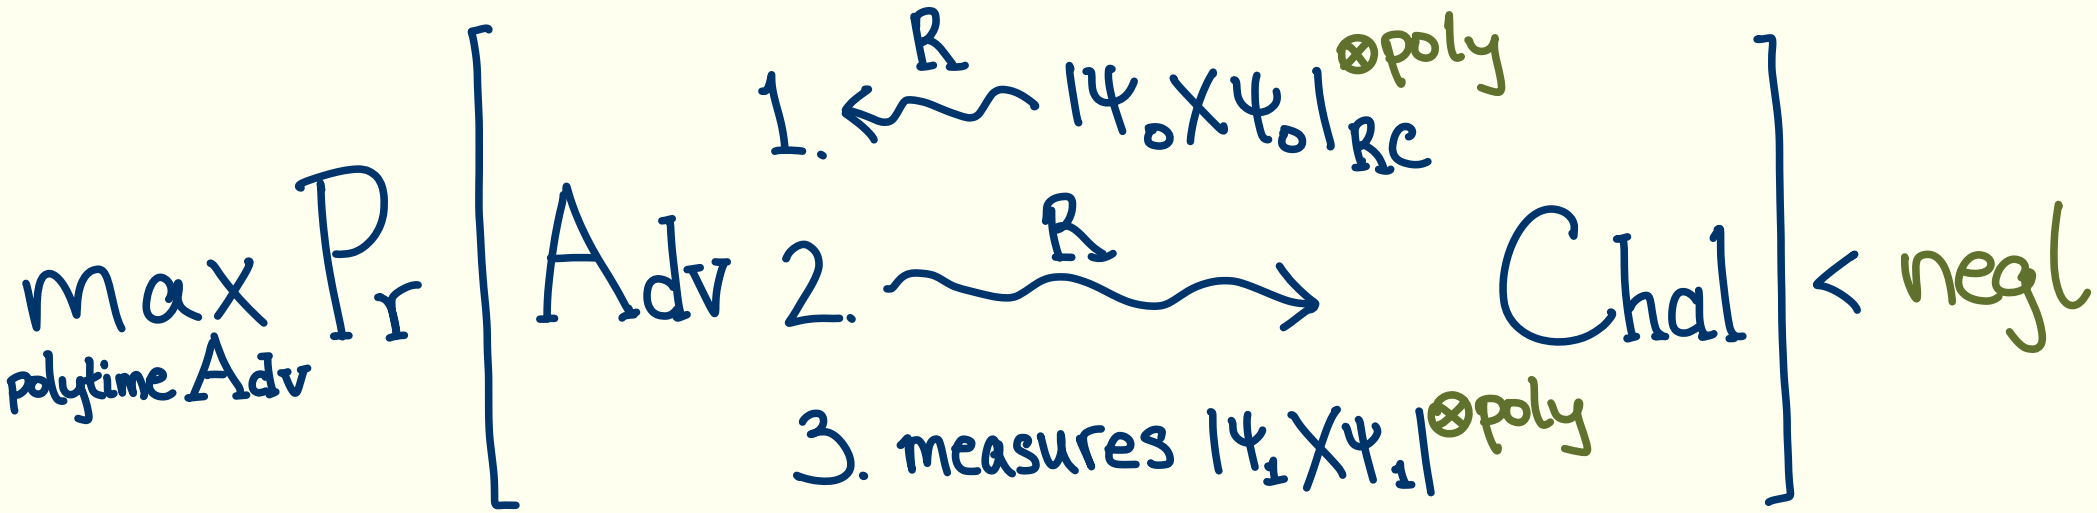

Proof strategy and Techniques

Proving parallel repetition

Let’s go through a proof of 2-fold parallel repetition.

Proving parallel repetition

Assume that we are given a computationally sound protocol.

$$\mathrm{Sound}(\mathrm{Chal}) < s$$

We want to show that the 2-fold parallel repetition satisfies:

$$\mathrm{Sound}(\mathrm{Chal}^{\otimes 2}) < s^2 + \mathrm{negl}$$

Proving parallel repetition

Instead, we proceed by contradiction. Assume that there is a super adversary that wins the 2-fold protocol with high probability.

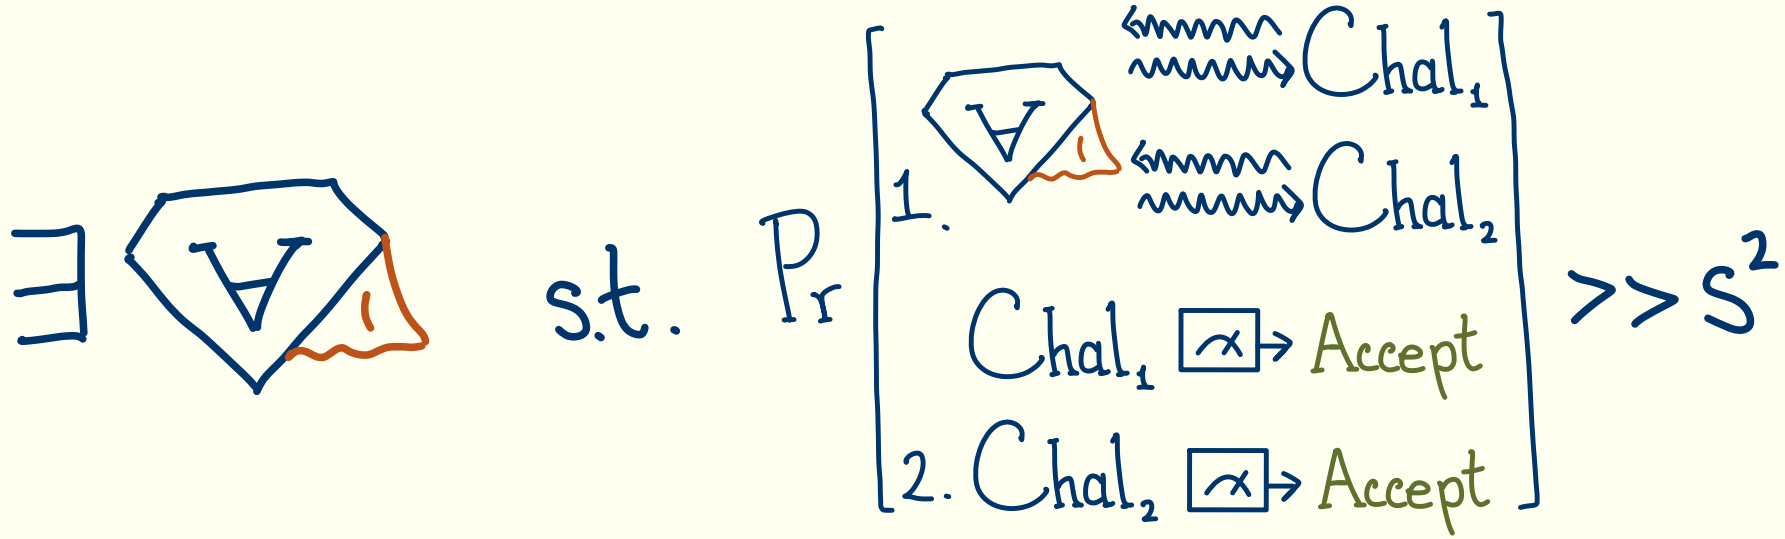

We want to design an 1-fold adversary that wins the 1-fold protocol with probability greater than s.

Designing our adversary

Our super adversary expects to play against two challengers, so our adversary will simulate one of them.

Designing our adversary

Our super adversary expects to play against two challengers, so our adversary will simulate one of them.

First question: should it simulate Chal_1 or Chal_2?

Designing our adversary: Bayes rule

Consider the probability that the super adversary wins against two challengers:

$$\mathbb{P}\left[\mathcal{A}_1 \text{ wins } \mathrm{Chal}_1 \wedge \text{wins } \mathrm{Chal}_2\right]$$

$$= \mathbb{P}\left[\mathcal{A}_1 \text{ wins } \mathrm{Chal}_2 \,\middle|\, \mathcal{A}_1 \text{ wins } \mathrm{Chal}_1\right] \cdot \mathbb{P}\left[\mathcal{A}_1 \text{ wins } \mathrm{Chal}_1\right]$$

Designing our adversary: Bayes rule

Consider the probability that the super adversary wins against two challengers:

$$\mathbb{P}\left[\text{A}_1 \text{ wins } Chal_2 \,\middle|\, \text{A}_1 \text{ wins } Chal_1\right] \cdot \mathbb{P}\left[\text{A}_1 \text{ wins } Chal_1\right]$$

$$>> s^2$$

By assumption, this is larger than s^2!

Bayes rule, case 1

What if the second term is greater than s?

Bayes rule, easy case

What if the second term is greater than s?

We are done: The adversary simulates Chal_2, and the real challenger (who is in position 1) accepts with high probability.

Bayes rule, interesting case

If the second term is smaller than s, let's re-arrange terms:

$$\mathbb{P}[\text{A}_1 \text{ wins } Chal_2 \mid \text{A}_1 \text{ wins } Chal_1] \cdot \mathbb{P}[\text{A}_1 \text{ wins } Chal_1]$$

$$>> s^2$$

Bayes rule, interesting case

If the second term is smaller than s, let's re-arrange terms:

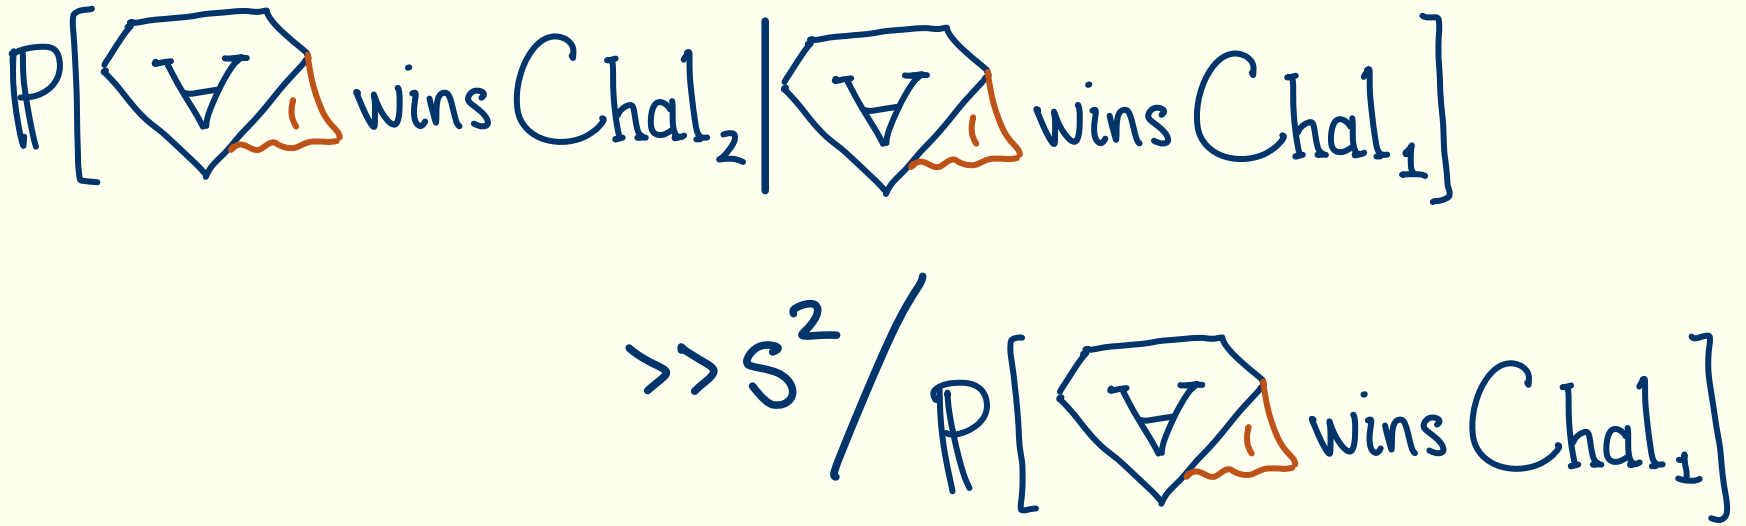

Bayes rule, interesting case

If the second term is smaller than s, let's re-arrange terms:

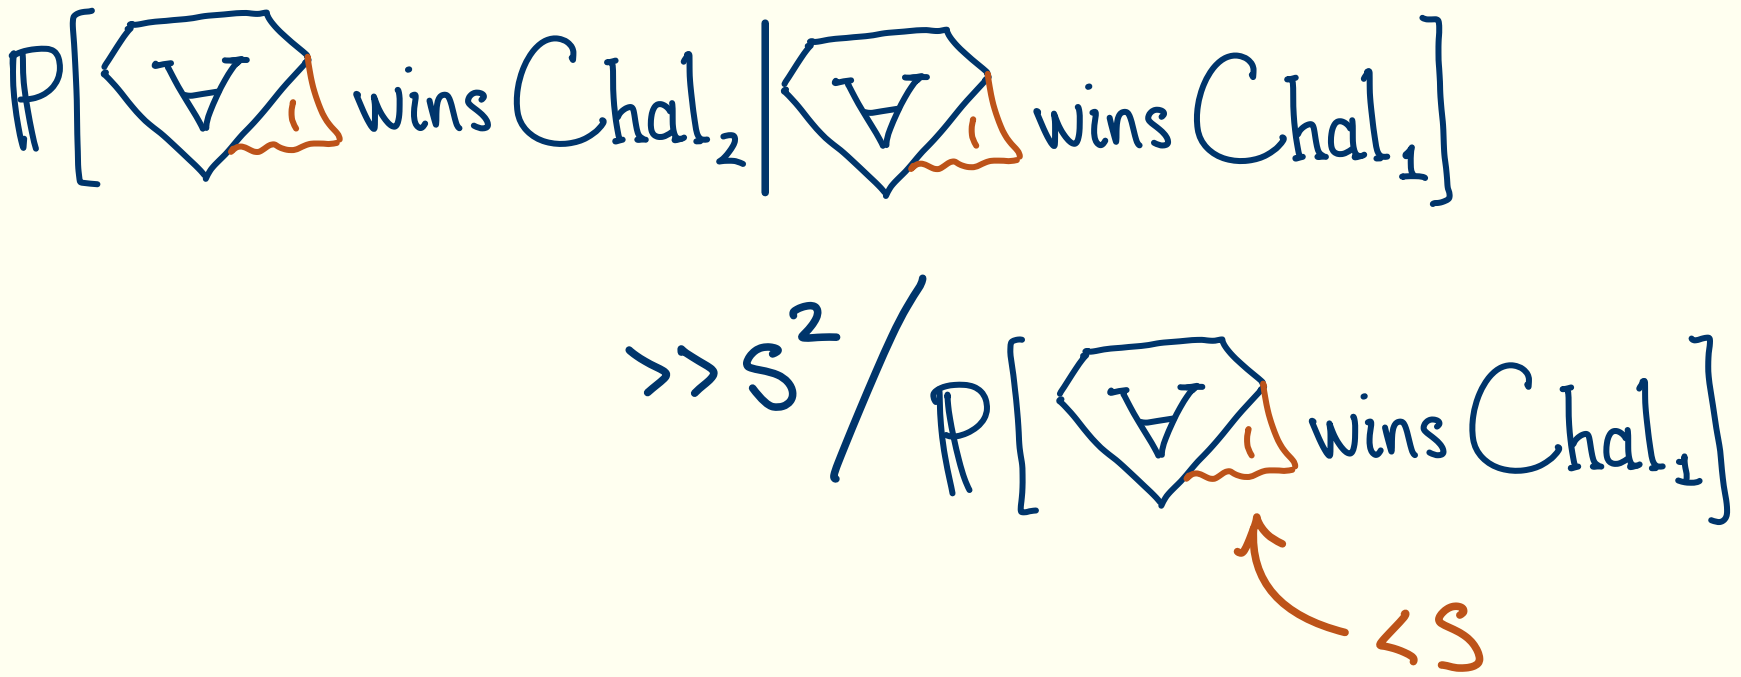

Bayes rule, interesting case

If the second term is smaller than s, let's re-arrange terms:

Bayes rule, interesting case

If we could guarantee that when the adversary simulates Chal_1, that challenger always accepts, we would be done!

Bayes rule, interesting case

If we could guarantee that when the adversary simulates Chal_1, that challenger always accepts, we would be done!

How can we post-select on Chal_1 accepting?

Classical post-selection

Classically we can keep simulating the protocol with the super adversary until we see that Chal_1 accepts (rejection sampling).

More steps involved in proving the result, but at a high level this works to prove classical parallel repetition.

Quantum rejection sampling?

Can we just do rejection sampling too?

Quantum rejection sampling?

Can we just do rejection sampling too?

No! When the challenger is quantum, we only get one copy of the challenge register. If we simulate the game and check if the first challenger accepts, we will destroy our one challenge.

Quantum post-selection

Let’s re-frame the problem a little bit.

Quantum post-selection

Let’s re-frame the problem a little bit.

From now on, we will assume that the first challenger is simulated by our adversary, and the second challenger is the real one!

Quantum post-selection

Let $|\psi\rangle_{AC_1C_2}$ be the state of the entire system (both challengers' private registers, C_1 and C_2 and adversary's register A) after the challengers send their challenge.

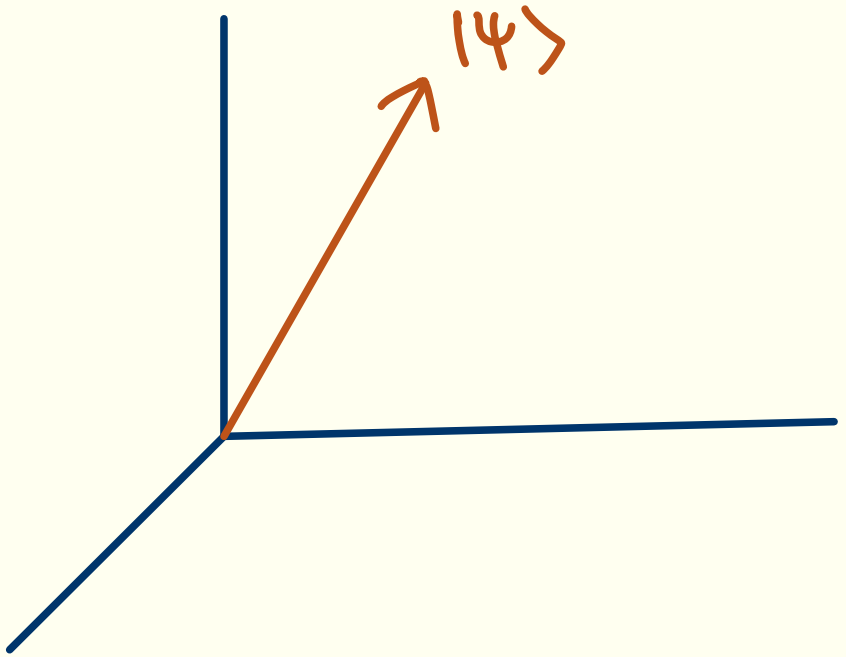

Quantum post-selection

Let $\widetilde{\Pi}_{AC_1}$ be the subspace of states accepted by the first challenger, after the adversary performs their unitary.

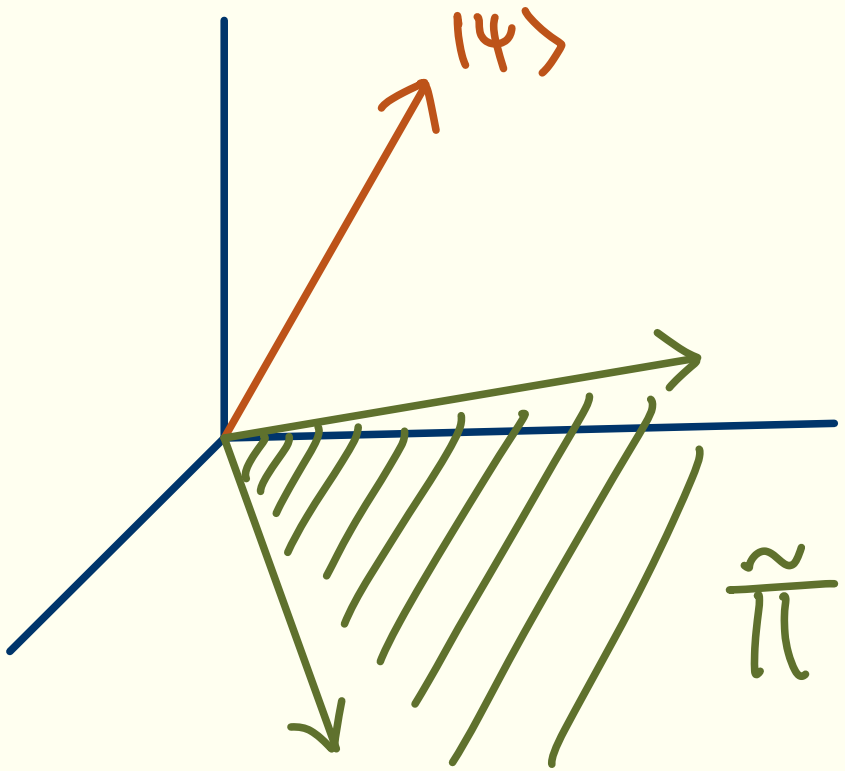

Quantum post-selection

If we could turn $|\psi\rangle_{AC_1C_2}$ into the normalized projection of $|\psi\rangle_{AC_1C_2}$ onto $\widetilde{\Pi}_{AC_1}$, we would have successfully post-selected on the first challenger accepting!

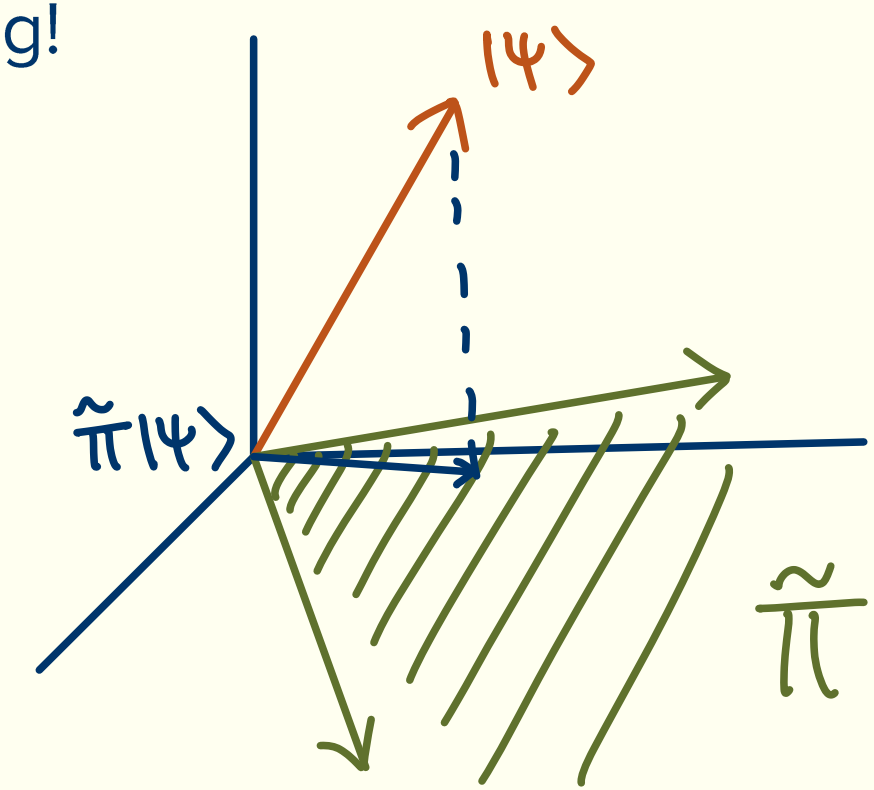

Quantum post-selection

If we could turn $|\psi\rangle_{AC_1C_2}$ into the normalized projection of $|\psi\rangle_{AC_1C_2}$ onto $\tilde{\Pi}_{AC_1}$, we would have successfully post-selected on the first challenger accepting!

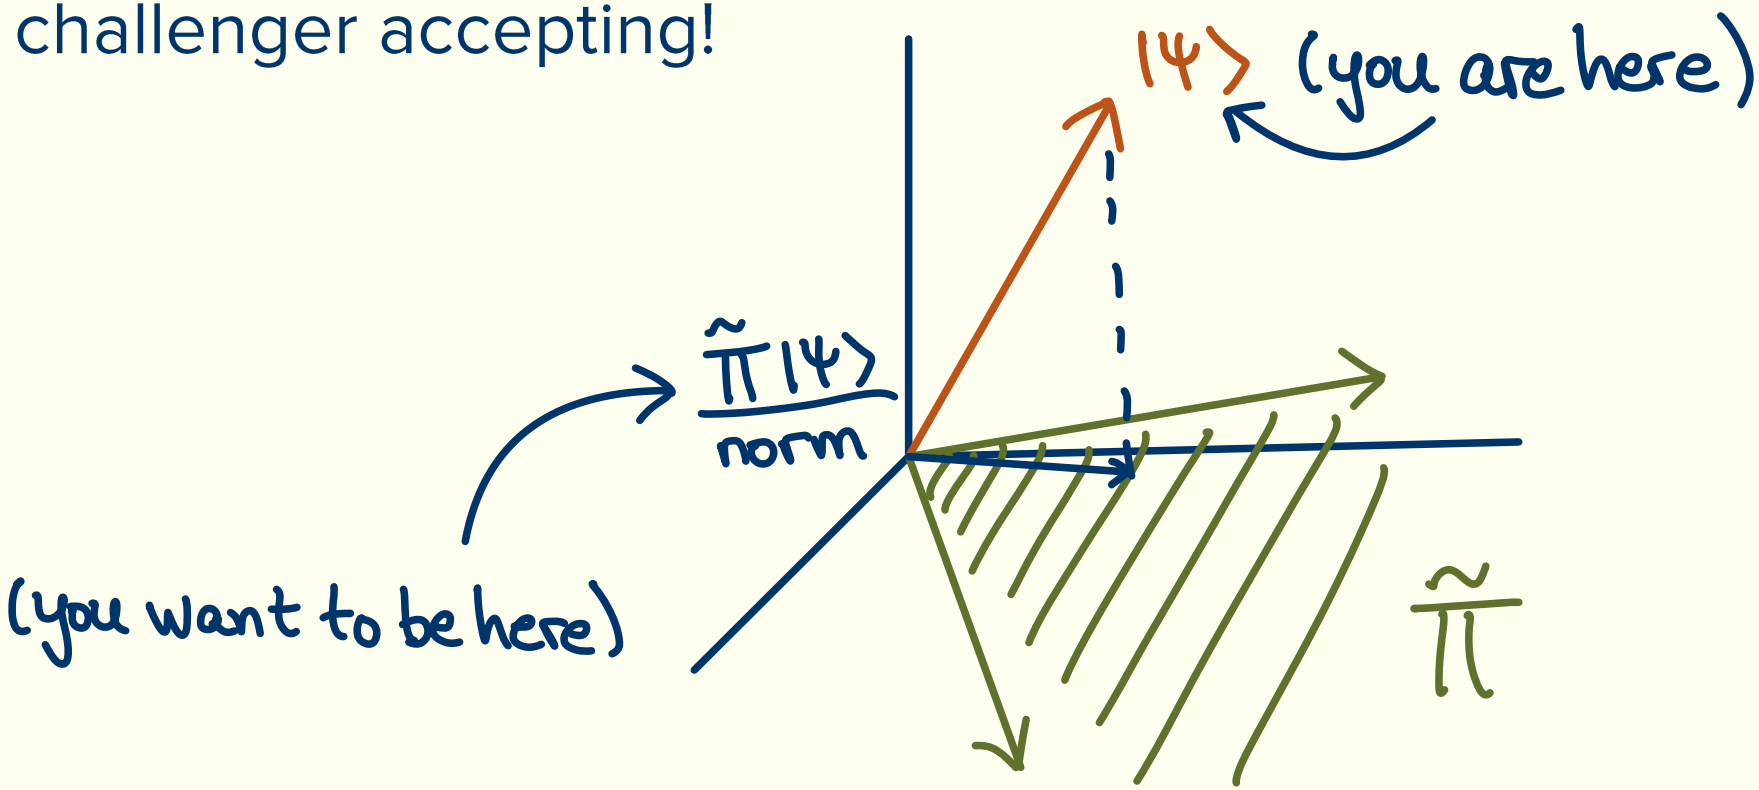

Quantum amplitude amplification?

Can we use regular amplitude amplification (i.e. Grover's search) to achieve this? Sadly no!

Quantum amplitude amplification?

Can we use regular amplitude amplification (i.e. Grover's search) to achieve this? Sadly no!

In the real protocol, our adversary does not have access to the register C_2, but amplitude amplification requires performing $\mathrm{id} - |\psi\rangle\langle\psi|$!

Quantum post-selection

We can’t perform a flip around the state $|\psi\rangle$, but we do know of a projector Π_{AC_1} that definitely contains $|\psi\rangle$!

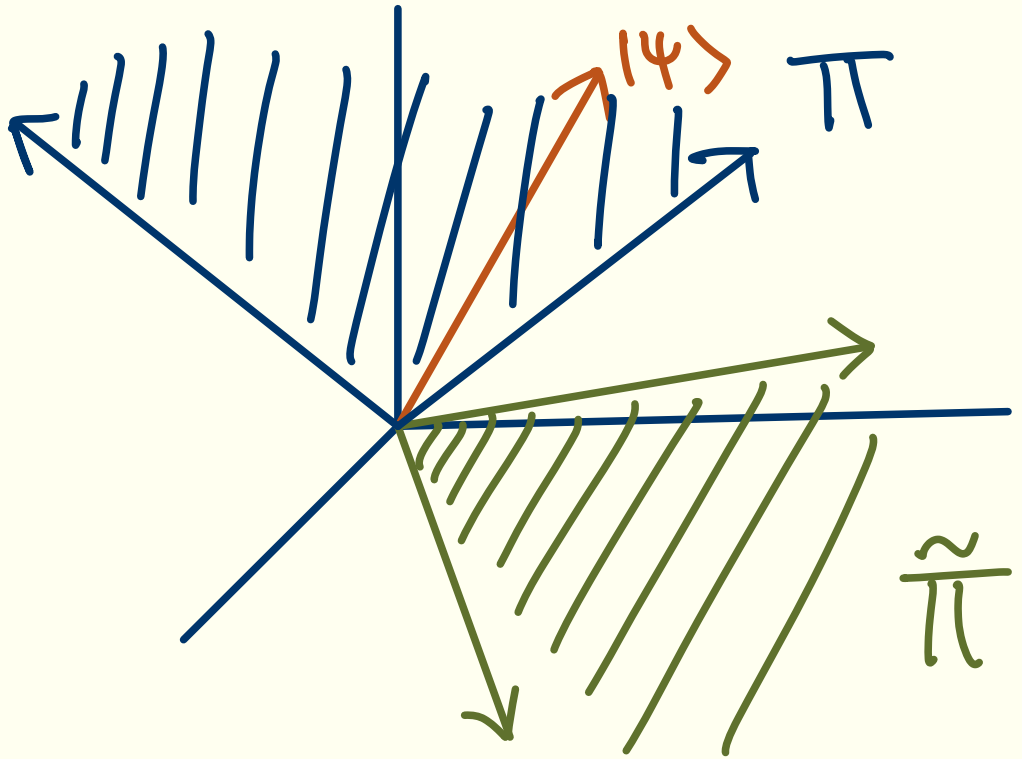

Quantum post-selection

We can't perform a flip around the state $|\psi\rangle$, but we do know of a projector Π_{AC_1} that definitely contains $|\psi\rangle$!

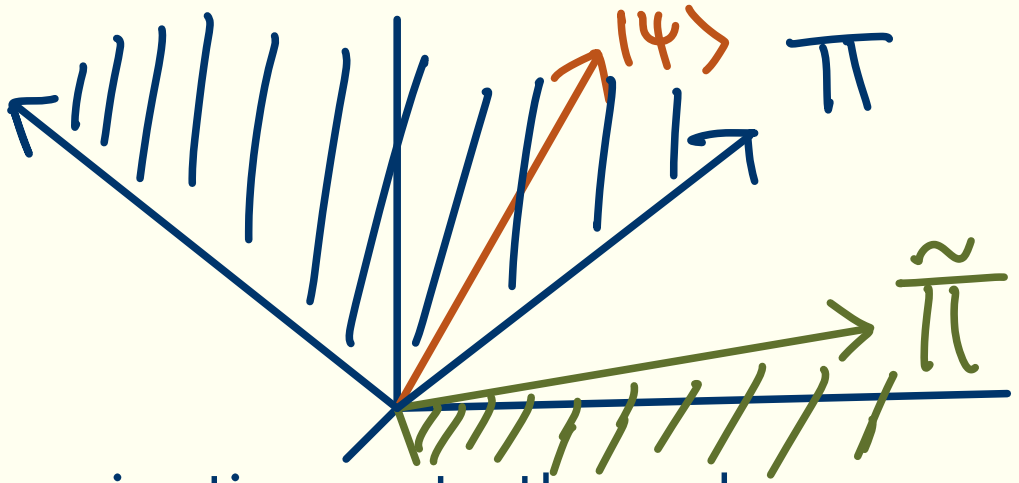

Here Π_{AC_1} is the projection onto the subspace of valid challenges from the first challenger (only).

Quantum rewinding?

Claim [CMSZ22]: Alternating Π and $\widetilde{\Pi}$ would give us some state in $\widetilde{\Pi}$.

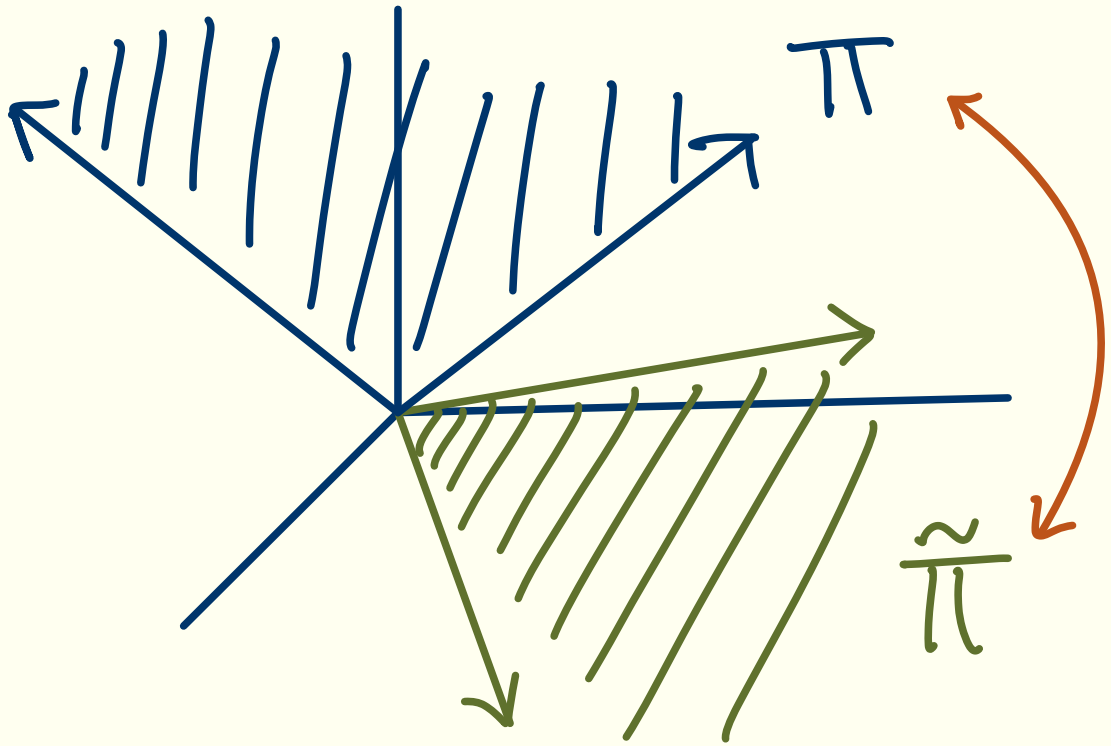

Quantum rewinding?

Claim [CMSZ22]: Alternating Π and $\widetilde{\Pi}$ would give us some state in $\widetilde{\Pi}$.

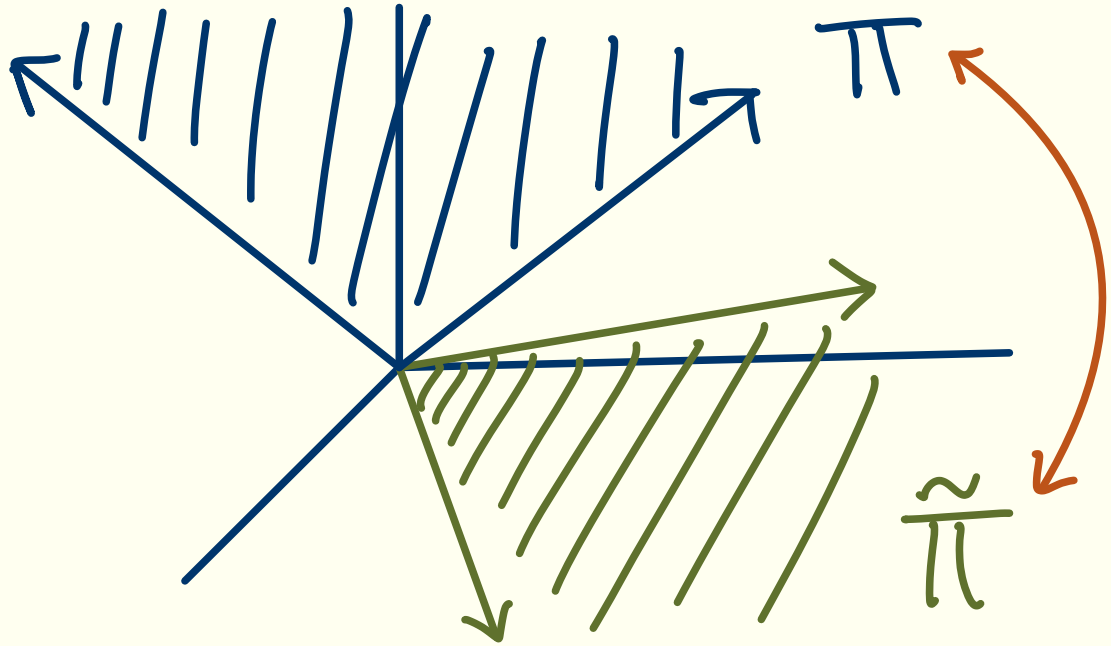

Sadly, it will not necessarily give us the projection onto $\widetilde{\Pi}$ for every state (only singular vectors)!

Quantum post-selection

Let's write $|\psi\rangle$ in the eigen-basis of Π, $\{|v_i\rangle\}$:

$$|\psi\rangle = \sum_i \alpha_i |v_i\rangle$$

Quantum post-selection

Now let's write the state that we want, in the basis of $\widetilde{\Pi}$, $\{|w_i\rangle\}$:

$$|\psi_{goal}\rangle \propto \sum_i \alpha_i \langle w_i | v_i \rangle |v_i\rangle$$

Quantum post-selection

Now let's write the state that we want, in the basis of $\tilde{\Pi}$, $\{|w_i\rangle\}$:

$$|\psi_{goal}\rangle \propto \sum_i \alpha_i \langle w_i|v_i\rangle |v_i\rangle$$

× Renormalization

Quantum post-selection

Now let's write down the singular value decomposition of $\tilde{\Pi}\Pi$:

$$\tilde{\Pi}\Pi = \sum_i c_i |w_i\rangle\langle v_i|$$

Quantum post-selection

Let's write down the singular value decomposition of $\widetilde{\Pi}\Pi$:

$$\widetilde{\Pi}\,\Pi = \sum_i c_i |w_i\rangle\langle v_i|$$

The values $|w_i\rangle$ are exactly what we want, the projection of vectors $|v_i\rangle$ from Π onto $\widetilde{\Pi}$!

Quantum singular value transform

Let's write down the singular value decomposition of $\Pi\widetilde{\Pi}$:

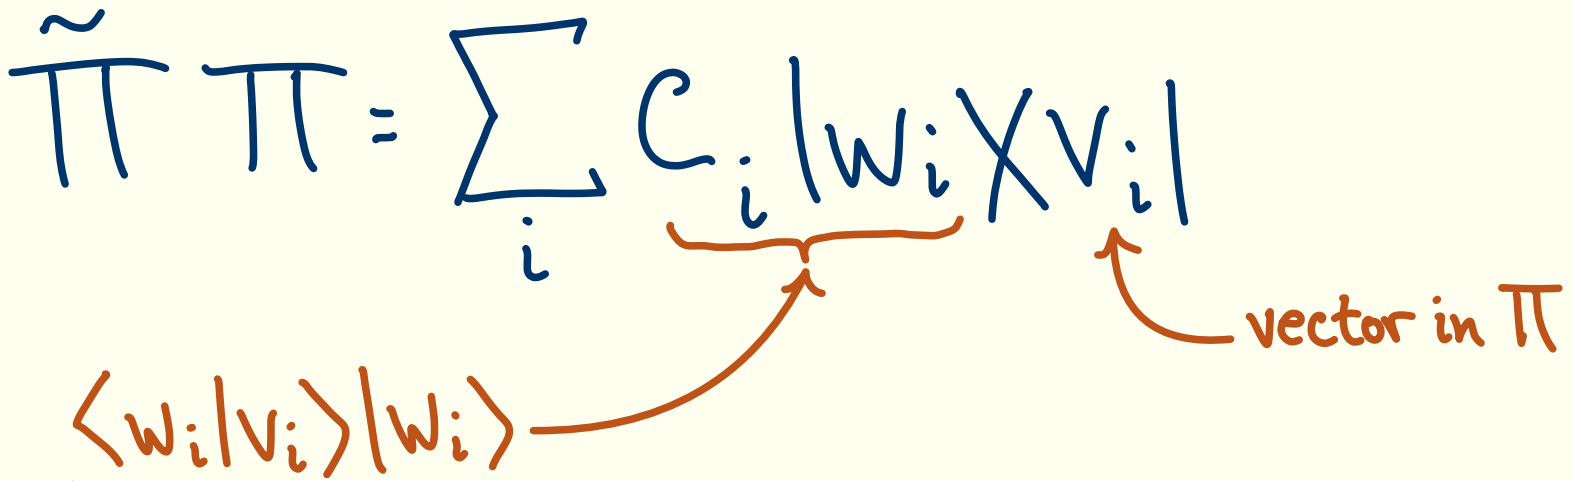

The values $|w_i\rangle$ are exactly what we want, the projection of vectors $|v_i\rangle$ from Π onto $\widetilde{\Pi}$!

Quantum singular value transform

The QSVT allows us to manipulate the singular values of $\Pi\widetilde{\Pi}$ only using those two projectors (and some phases), and thus allows us to do an approximate projection without touching C_2!

Quantum singular value transform

The QSVT allows us to manipulate the singular values of $\Pi\tilde{\Pi}$ only using those two projectors (and some phases), and thus allows us to do an approximate projection without touching C_2!

$$QSVT(\tilde{\Pi}, \Pi) \approx \sum_i c_i \gamma |w_i\rangle\langle v_i|$$

Post-selection factor.

Quantum singular value transform

The QSVT allows us to manipulate the singular values of $\Pi\tilde{\Pi}$ only using those two projectors (and some phases), and thus allows us to do an approximate projection without touching C_2!

$$QSVT(\tilde{\Pi}, \Pi) \approx \sum_i c_i \gamma |w_i\rangle\langle v_i|$$

Post-selection factor.

Details omitted, but that is the main technical idea.

Open Questions

1. Does parallel repetition decrease the soundness of arbitrary message public coin protocols exponentially?
2. Is there a modification that can be made to a protocol that makes parallel repetition decrease soundness exponentially (i.e. randomly terminating)?
3. Other applications of the parallel repetition theorem for quantum interactive arguments?

Open Questions

1. Does parallel repetition decrease the soundness of arbitrary message public coin protocols exponentially?
2. Is there a modification that can be made to a protocol that makes parallel repetition decrease soundness exponentially (i.e. randomly terminating)?
3. Other applications of the parallel repetition theorem for quantum interactive arguments?

Thanks for listening!