Separating QMA from QCMA with a classical oracle

John Bostanci, Jonas Haferkamp, Chinmay Nirkhe, and Mark Zhandry

STOC 2026

Some images generated using ChatGPT images 2.0


I love you!


HOMER HERODOTUS SOPHOCLES PLATO ARISTOTLE DEMOSTHENES CICERO VERGIL


HOMER HERODOTUS SOPHOCLES PLATO ARISTOTLE DEMOSTHENES CICERO VERGIL


Prove it!

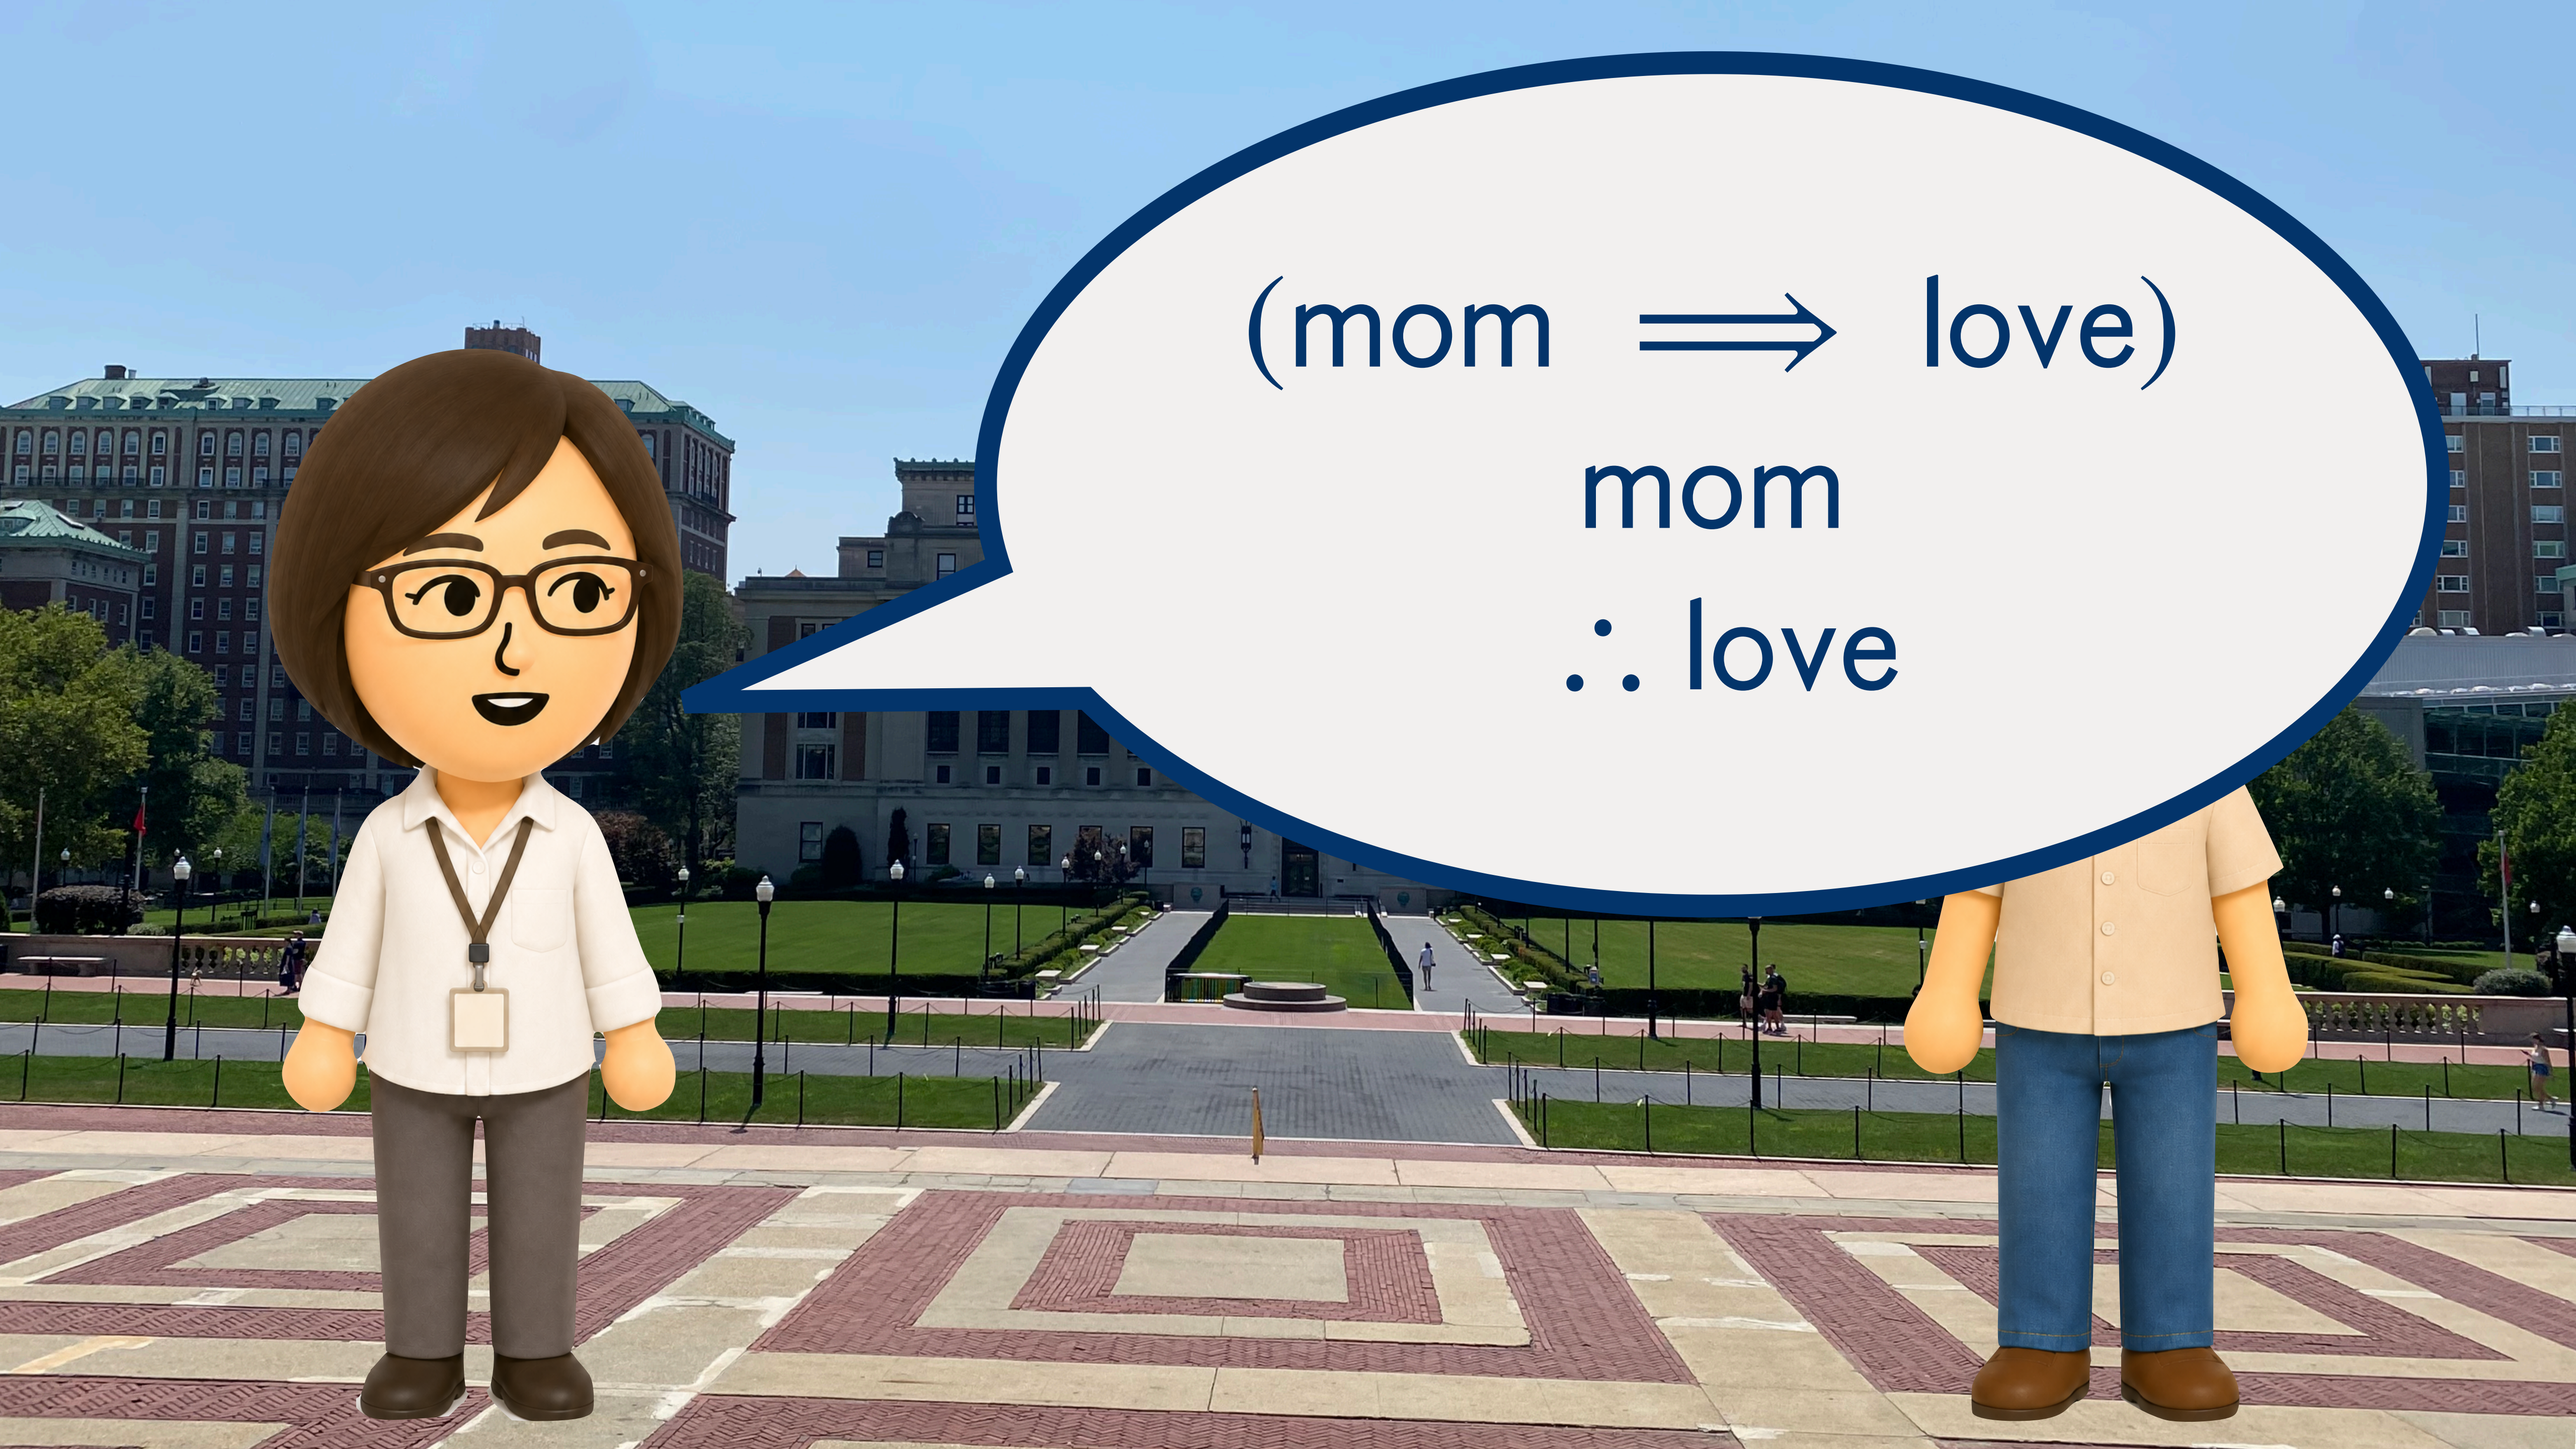
(mom ⟹ love)
mom
∴ love

Proofs in computer science

Efficient Classical Algorithm

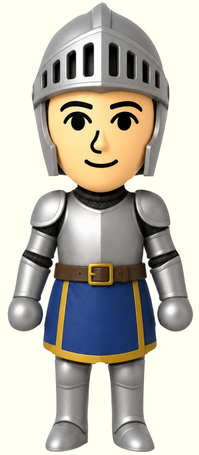

Proofs in computer science

Efficient Classical Algorithm

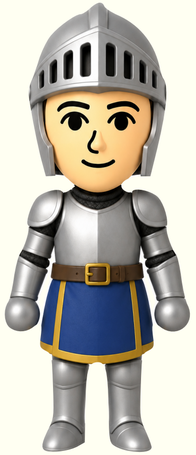

Unbounded prover

Classical proof

Efficient Classical Algorithm

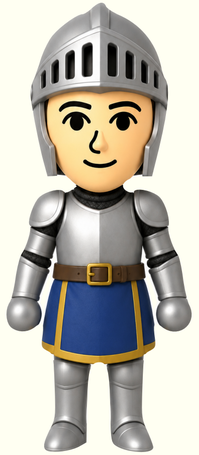

Proofs in computer science

Efficient Classical Algorithm

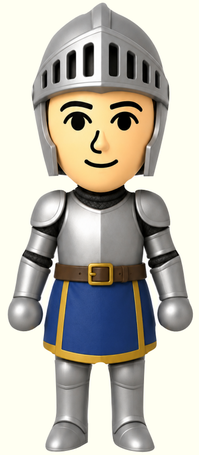

P vs NP

Unbounded prover

Classical proof

Efficient Classical Algorithm

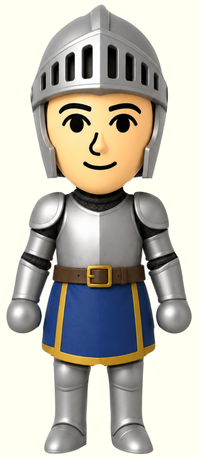

Proofs in quantum computer science

Efficient Quantum Algorithm

vs

Unbounded prover

Classical proof

Efficient Quantum Algorithm

Proofs in quantum computer science

Unbounded prover

Classical proof

Efficient Quantum Algorithm

vs

Unbounded prover

Quantum proof

Efficient Quantum Algorithm

Proofs in quantum computer science

Unbounded prover

Classical proof

Efficient Quantum Algorithm

QCMA vs QMA

Unbounded prover

Quantum proof

Efficient Quantum Algorithm

Quantum versus classical proofs

QMA captures a lot of things we want to know about the world:

Quantum versus classical proofs

QMA captures a lot of things we want to know about the world:

- Is there a thing that satisfies some physical property?
- Are two different physical processes doing the same transformation?
- Are local views of a system consistent with some larger quantum state?

Quantum versus classical proofs

QMA captures a lot of things we want to know about the world:

- Is there a thing that satisfies some physical property?
- Are two different physical processes doing the same transformation?
- Are local views of a system consistent with some larger quantum state?

→ If there were ways to efficiently verify these with a classical proof, quantum systems have a lower “complexity” than we thought!

Quantum versus classical proofs

QMA captures a lot of things we want to know about the world:

- Is there a thing that satisfies some physical property?
- Are two different physical processes doing the same transformation?
- Are local views of a system consistent with some larger quantum state?

→ If there were ways to efficiently verify these with a classical proof, quantum systems have a lower “complexity” than we thought!

→ If not, then there are problems that you can only check them with a quantum proof!

History of the QMA versus QCMA problem

- First proposed in ’02 by **Aharanov and Naveh**.
- **Aaronson & Kuperberg ’06**: Quantum oracle separation. Each $\mathcal{O}_n = \mathrm{id} - 2\,|\psi_n\rangle\langle\psi_n|$
- **Lutomirski ’11**: Proposed the expander mixing problem as a candidate classical oracle separation.
- **Fefferman & Kimmel ’15**: In-place permutation oracle. Problem corresponds to set size estimation.
- **Natarajan & Nirkhe ’22**: Distribution testing oracle. Problem corresponds to size estimation of an expander graph.
- **Li, Liu, Pelecanos, Yamakawa ’23**: Separation assuming only classical queries. Based on “Verifiable Quantum Advantage without Structure”.
- **Ben-David & Kundu ’24**: Bounded adaptivity, based on “Verifiable Quantum Advantage without Structure”.
- **Liu, Mutreja, Yuen ’25**: Aaronson-Ambainis-like conjecture implies QMA QCMA separation. Problem corresponds to size estimation of an expander graph.

Why is this problem so hard?

Why is this problem so hard?

→ We have few techniques that would differentiate quantum and classical proofs!

VS

Why is this problem so hard?

→ We have few techniques that would differentiate quantum and classical proofs!

- The typical trick is to “guess the witness”, but this works for a quantum witness.

VS

Why is this problem so hard?

→ We have few techniques that would differentiate quantum and classical proofs!

- The typical trick is to “guess the witness”, but this works for a quantum witness.

VS

→ A quantum verifier must use their proof in an “interesting” way, can’t just measure their proof, because otherwise I could send the measurement result as a proof!

Why is this problem so hard?

→ We have few techniques that would differentiate quantum and classical proofs!

- The typical trick is to “guess the witness”, but this works for a quantum witness.

VS

→ A quantum verifier must use their proof in an “interesting” way, can’t just measure their proof, because otherwise I could send the measurement result as a proof!

- Quantum algorithms are really good at detecting global structure, so it’s even hard to rule out BQP algorithms for some candidate oracle separations!

A hard problem for QCMA [FK'15]

A hard problem for QCMA [FK'15]

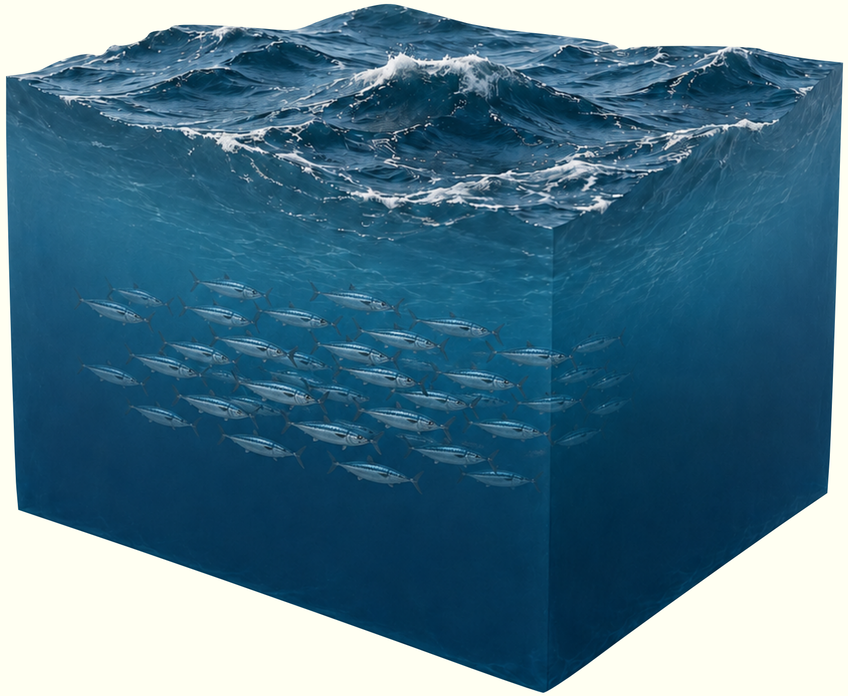

VS

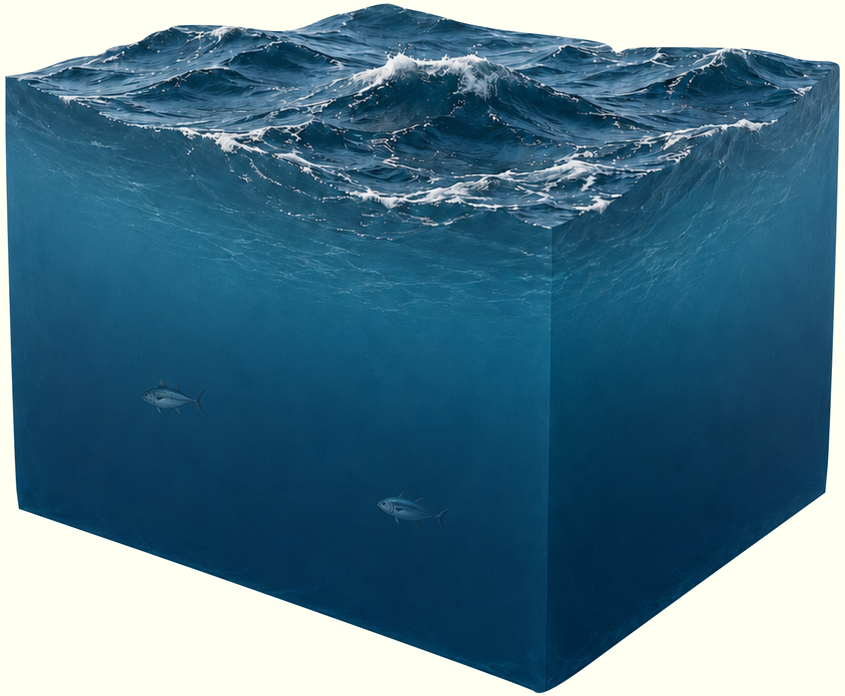

A hard problem for QCMA [FK'15]

In slightly more detail, we are considering the following problem:

Input: Oracle access to a set S (where fish are), and size ℓ (huge!)

Output: Is $|S| \geq \ell$ (YES) or $|S| \leq \ell/2$ (NO), promised one of the two is the case.

VS

A new approach for ruling out QCMA

Assume for the sake of contradiction that the sailor (with classical witness) would successfully solve this problem, then he definitely distinguishes between many fish and no fish…

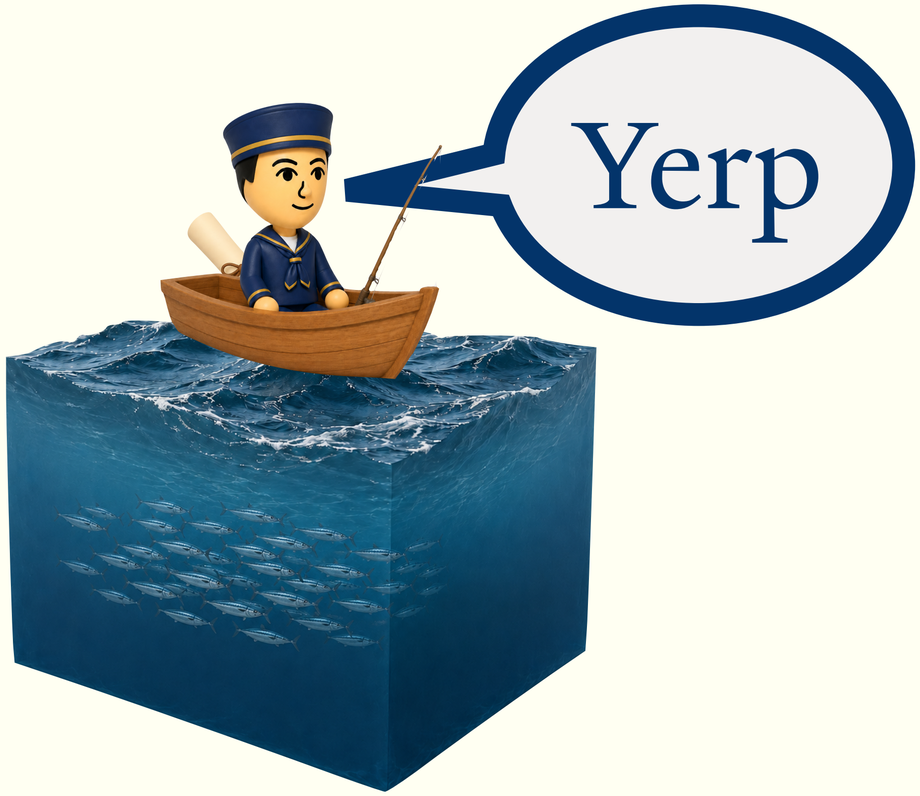

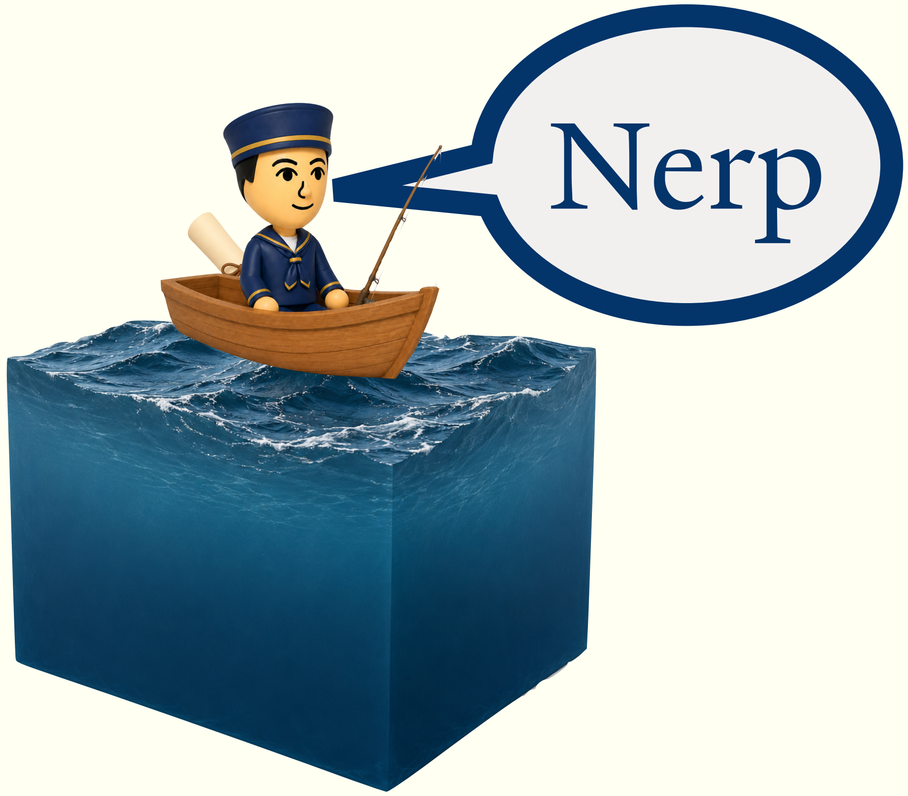

A new approach for ruling out QCMA

Assume for the sake of contradiction that the sailor (with classical witness) would successfully solve this problem, then he definitely distinguishes between many fish and no fish…

If the sailor never catches a fish, they would not distinguish between the two cases!

A new approach for ruling out QCMA

Imagine ``simulating'' the problem in a pool: Let the sailor move around as if they were in the ocean, but any time they fish they won't catch anything.

A new approach for ruling out QCMA

Imagine ``simulating'' the problem in a pool: Let the sailor move around as if they were in the ocean, but any time they fish they won't catch anything.

Because they (by assumption) distinguished between the two cases, if we look at where they want to fish, we will likely catch a fish! Say they fished at position x_1…

A new approach for ruling out QCMA

Now, by assumption the should distinguishes between the ocean with many fish and then ocean with one fish at location x_1…

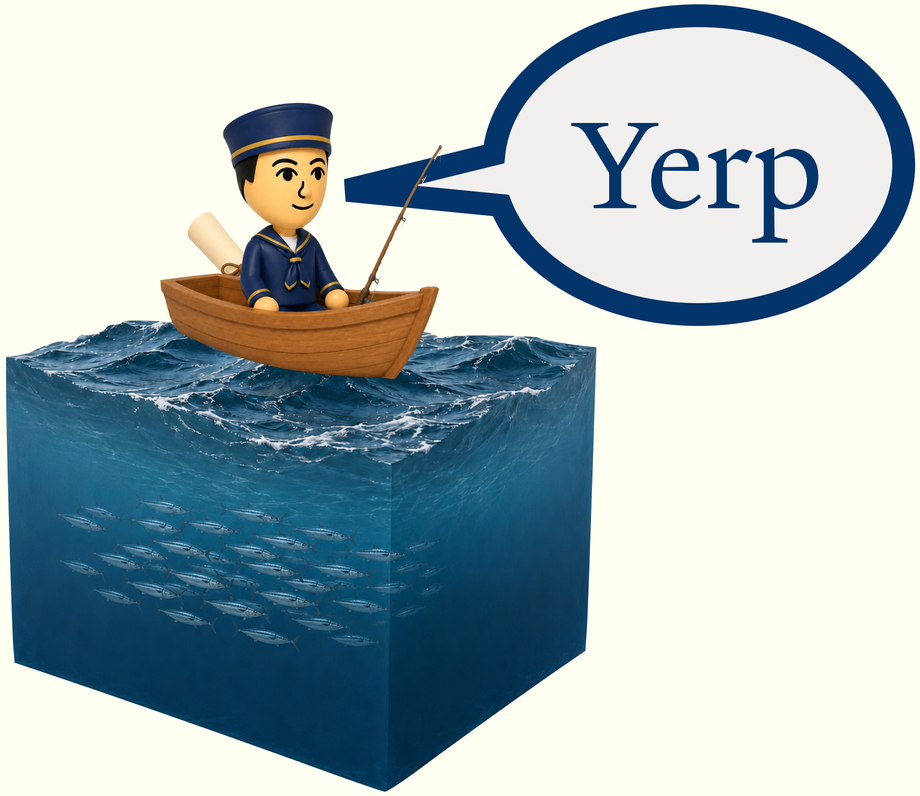

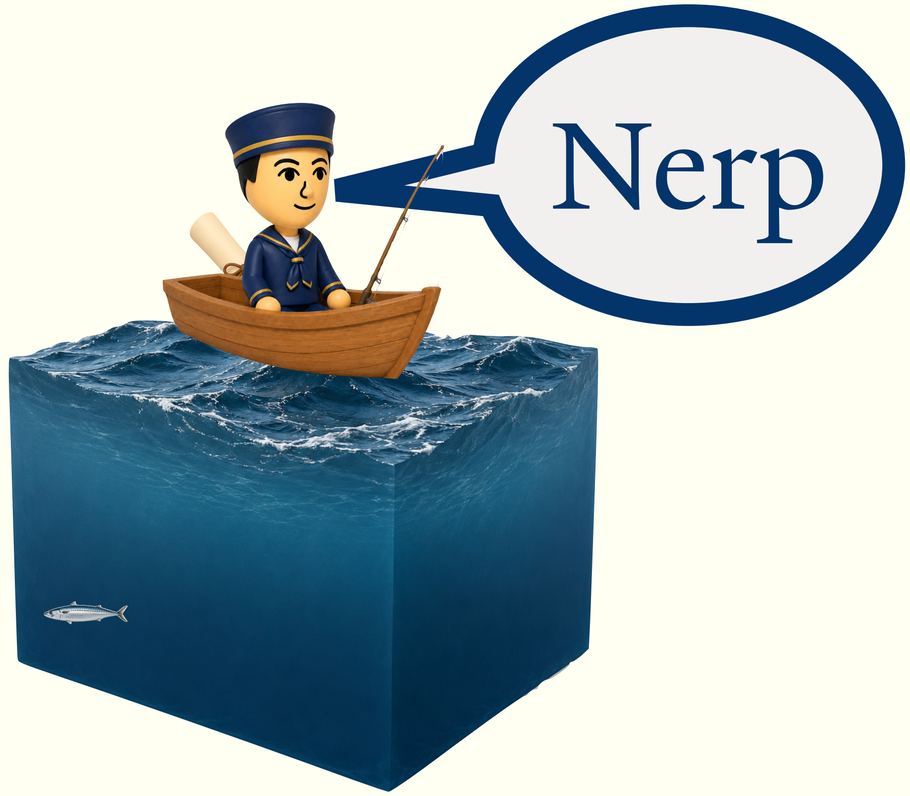

If the sailor doesn't find a fish in a different location, they can't distinguish them!

A new approach for ruling out QCMA

Now simulate this problem in a pool: Let the sailor move around as if they were in the ocean, but any time they fish they won't catch anything, unless they fish at x_1.

A new approach for ruling out QCMA

Now simulate this problem in a pool: Let the sailor move around as if they were in the ocean, but any time they fish they won't catch anything, unless they fish at x_1.

If we look at where they want to fish, we will likely catch a fish in a different location! Say they fished at position x_2…

A new approach for ruling out QCMA

We can keep repeating this process: Simulate the sailor in a pool to get another position where a fish will be, then re-do the simulation with more fake fish!

A new approach for ruling out QCMA

We can keep repeating this process: Simulate the sailor in a pool to get another position where a fish will be, then re-do the simulation with more fake fish!

Even if the map contains enough information to tell the sailor about v fish…
We can repeat this simulation to learn where $\gg v$ fish are, without even fishing at all!

A new approach for ruling out QCMA

We can keep repeating this process: Simulate the sailor in a pool to get another position where a fish will be, then re-do the simulation with more fake fish!

Contradiction!

Even if the map contains enough information to tell the sailor about v fish…
We can repeat this simulation to learn where $\gg v$ fish are, without even fishing at all!

A new approach for ruling out QCMA

An important thing to note about this proof: It depends on the fact that the witness is “re-usable” between the different simulations!

A new approach for ruling out QCMA

An important thing to note about this proof: It depends on the fact that the witness is “re-usable” between the different simulations!

If (instead) we had a potion that the sailor had to drink, once they drink it on the first round we wouldn’t be able to run the second simulation!

Getting back to QMA

Set size verification is outside of QCMA.

VS

Getting back to QMA

Set size verification is outside of QCMA. Unfortunately, it's also not in QMA!

VS

Getting back to QMA

Set size verification is outside of QCMA. Unfortunately, it's also not in QMA!

Let's see how to put the problem back in QMA.

VS

Spectral Forrelation

We can add a second oracle that encodes information about the “Fourier transform” of the set S. This results in a problem we call spectral Forrelation.

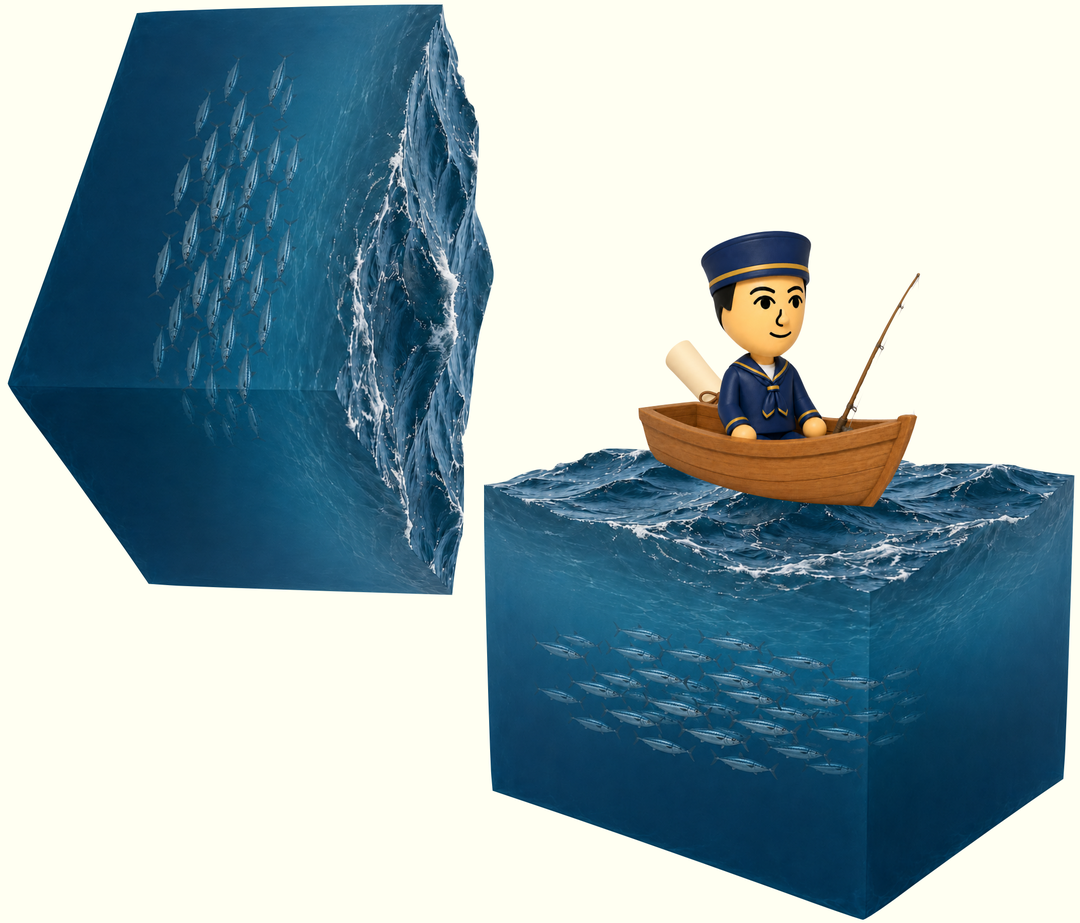

VS

The spectral Forrelation problem

The spectral Forrelation problem is a problem about pairs of sets (S, U), which we treat as oracles through the set membership functions.

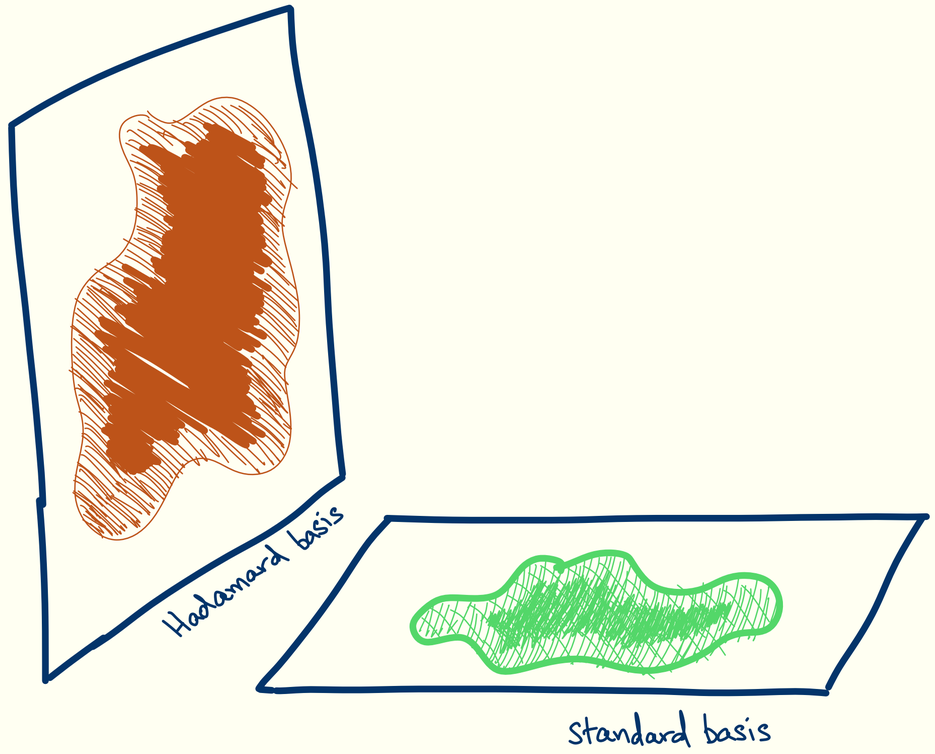

The spectral Forrelation problem

We say that two sets (S, U) are spectrally Forrelated if there is a state $|\psi\rangle$ such that
1. $|\psi\rangle$ is mostly supported on S
2. The Fourier transform of $|\psi\rangle$ is mostly supported on U.

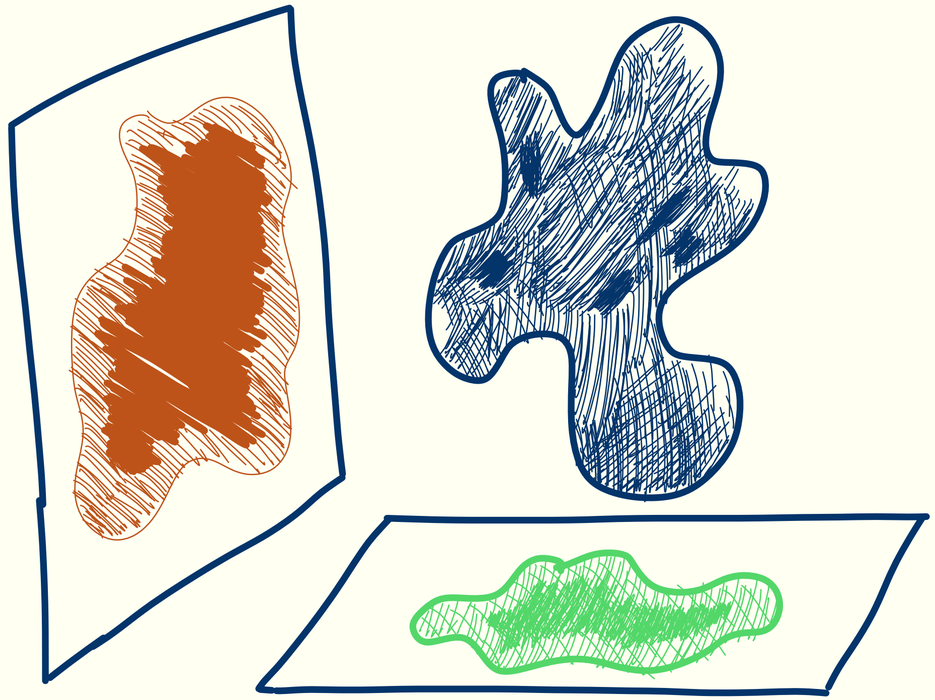

The spectral Forrelation problem

Input: Oracle access to two sets (S, U) (via set membership functions)

Output: Are they spectrally Forrelated (YES) or not (NO), promised one is the case?

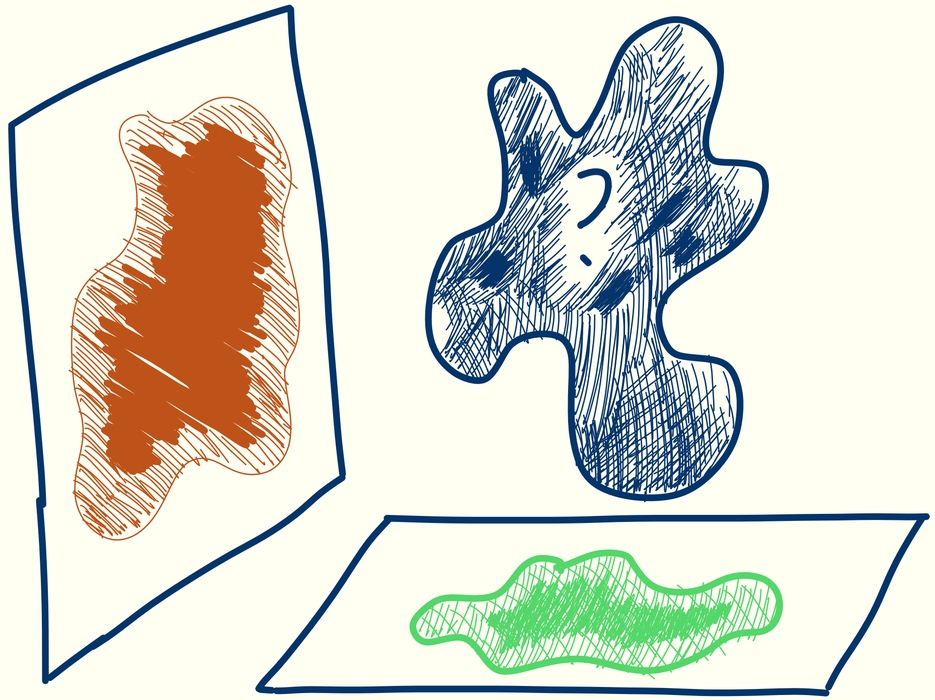

Spectral Forrelation is in QMA

Given a copy of a state $|\psi\rangle$:

- Check if the state is supported on S, reject if not.

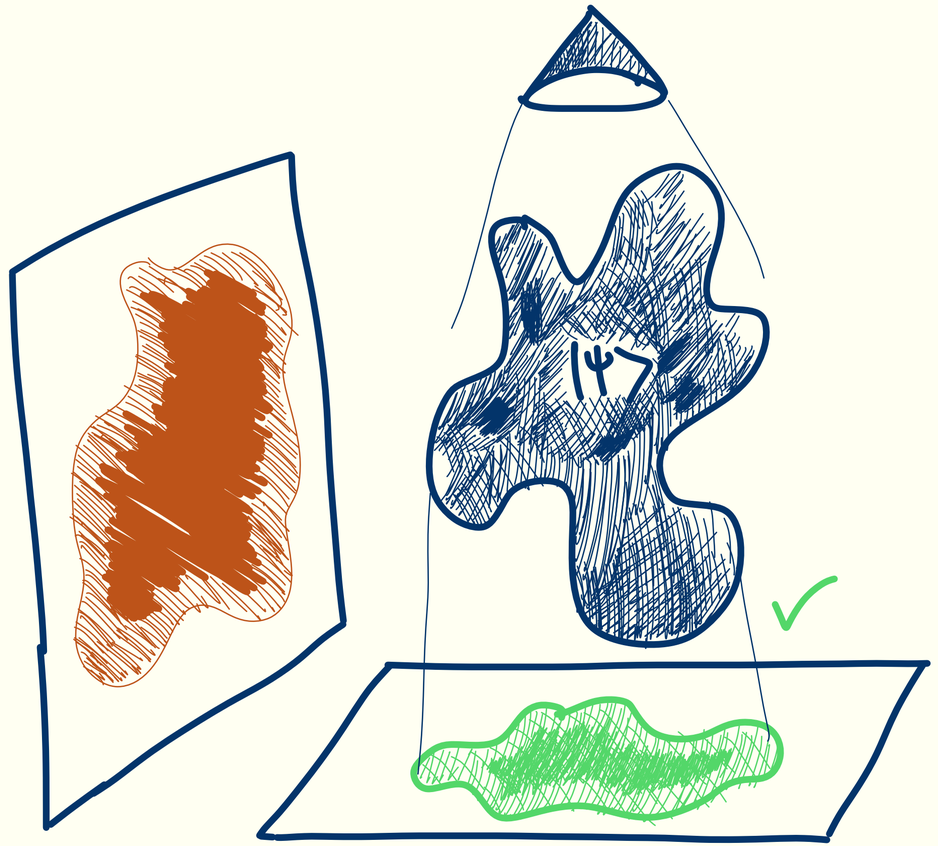

Spectral Forrelation is in QMA

Given a copy of a state $|\psi\rangle$:

- Check if the state is supported on S, reject if not.
- Fourier transform the state.

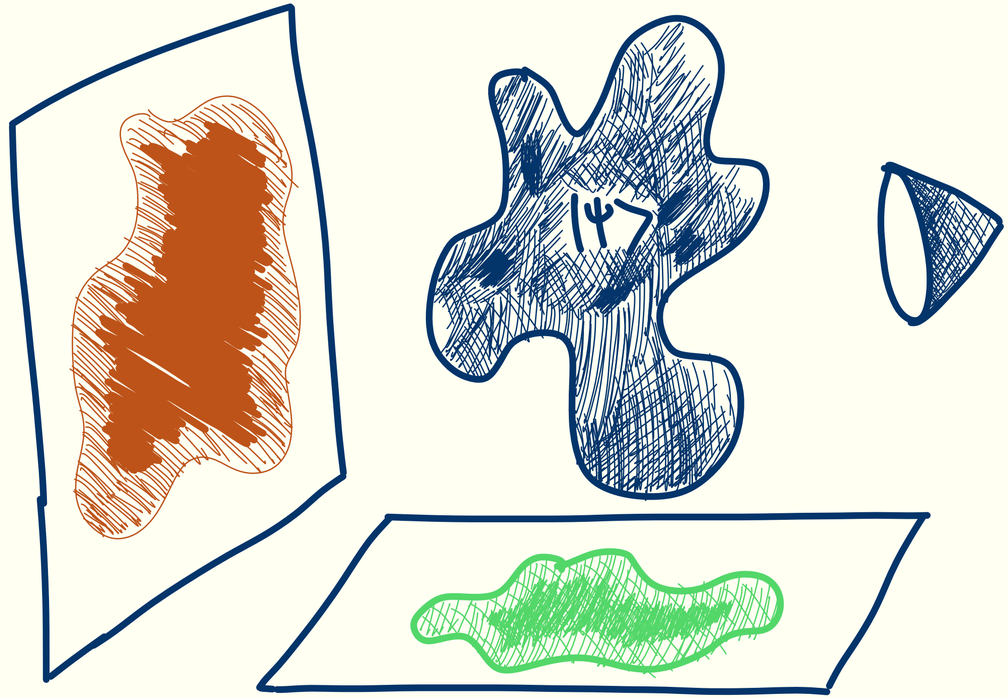

Spectral Forrelation is in QMA

Given a copy of a state $|\psi\rangle$:

- Check if the state is supported on S, reject if not.
- Fourier transform the state.
- Check if the state is supported on U, reject if not.
- Accept.

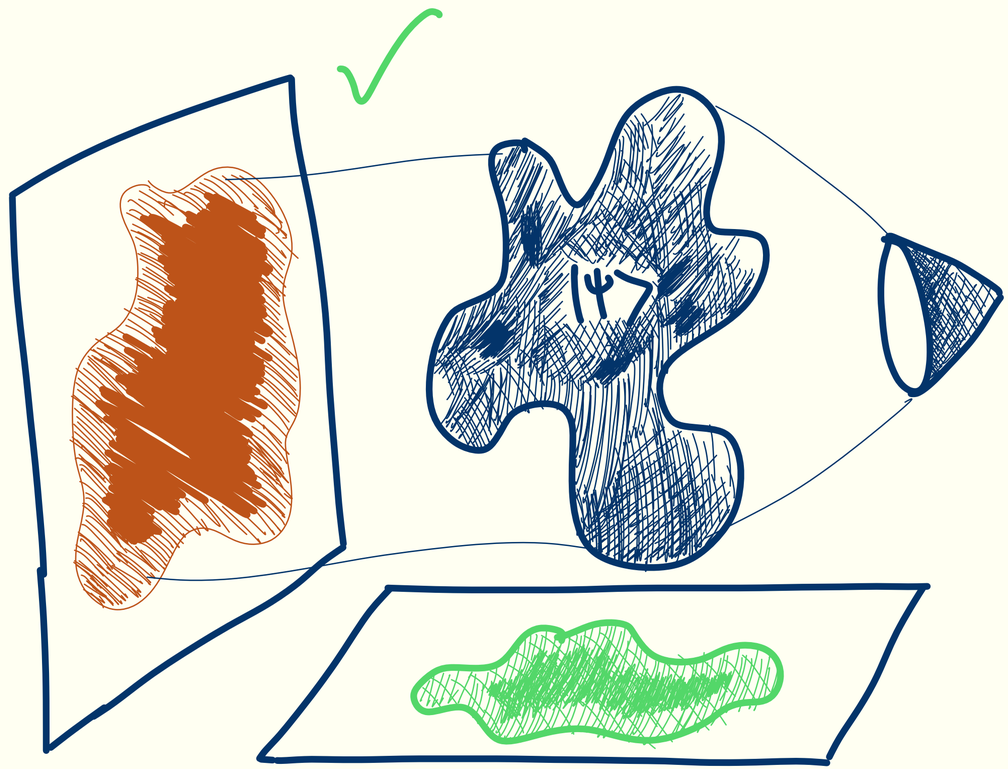

Spectral Forrelation is also not in QCMA

We rule out QCMA algorithms for spectral Forrelation in two steps:

Spectral Forrelation is also not in QCMA

We rule out QCMA algorithms for spectral Forrelation in two steps:

Theorem 1: A QCMA algorithm implies a good sampler that only looks at U.

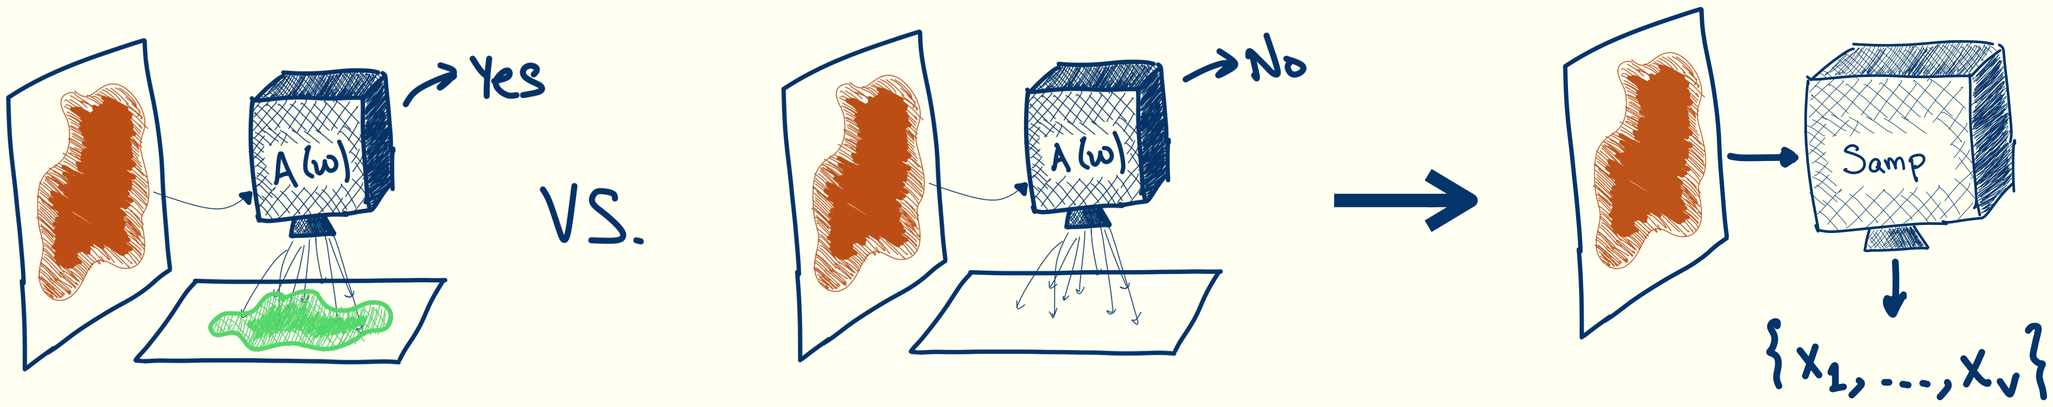

Spectral Forrelation is also not in QCMA

We rule out QCMA algorithms for spectral Forrelation in two steps:

Theorem 1: A QCMA algorithm implies a good sampler that only looks at U.

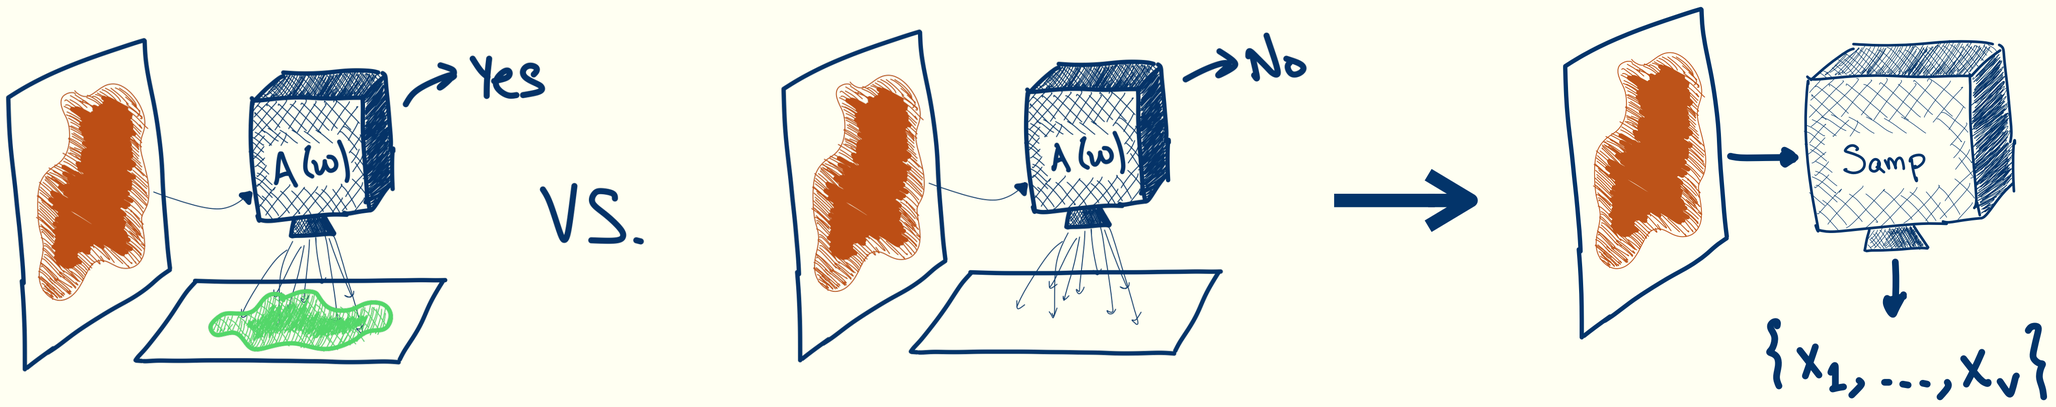

Theorem 2: A few queries to U can not help you sample points from S.

Spectral Forrelation is also not in QCMA

Theorem 2: A few queries to U can not help you sample points from S.

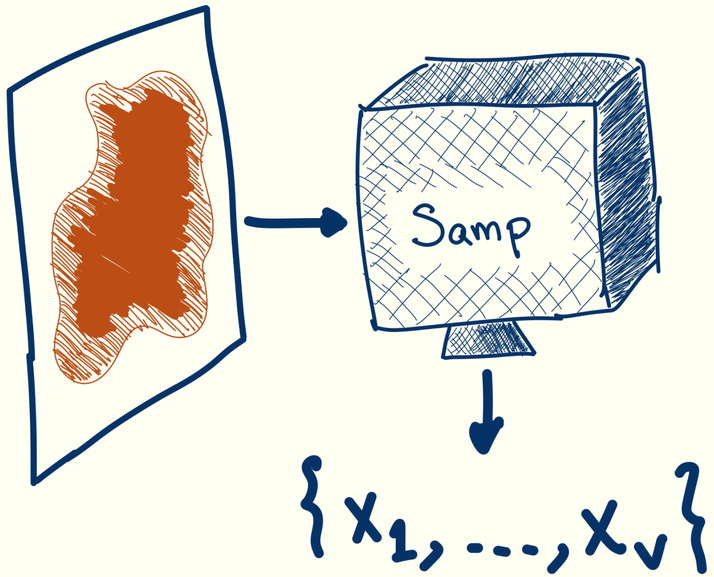

Spectral Forrelation is also not in QCMA

Theorem 2: A few queries to U can not help you sample points from S.

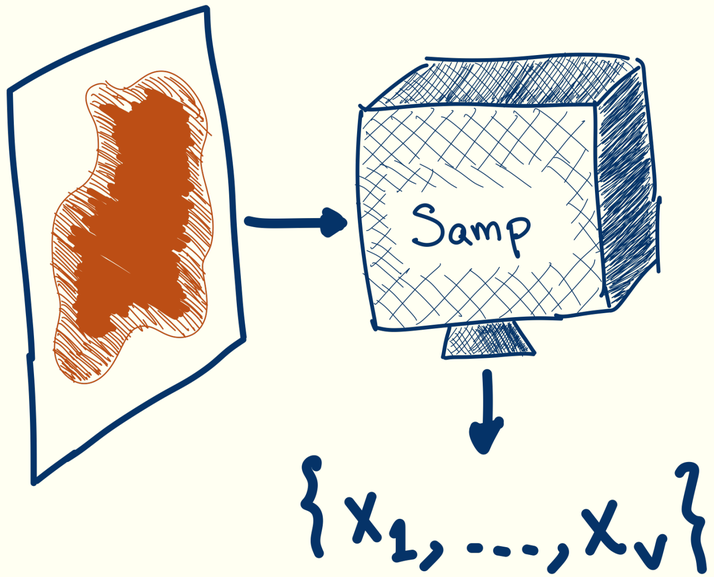

This is more tricky than before: U can be correlated to S.
→ We need to understand what the sampler knows about S when they look at U!

Lazy sampling

Classically, when someone interacts with a random function, we don't write down every entry. Lazy sampling is how we only store what the adversary knows about the function.

x	$f(x)$
z	1
w	0
y	0

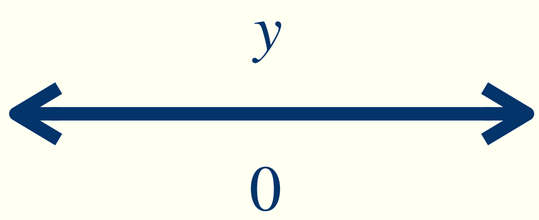

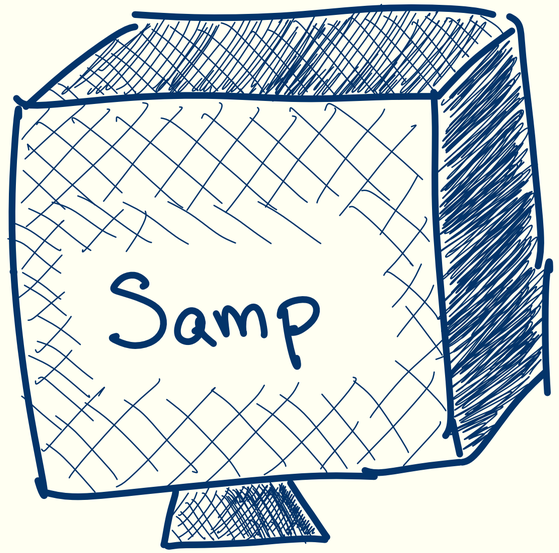

Lazy sampling

Classically, when someone interacts with a random function, we don't write down every entry. Lazy sampling is how we only store what the adversary knows about the function.

x	$f(x)$
z	1
w	0
y	0

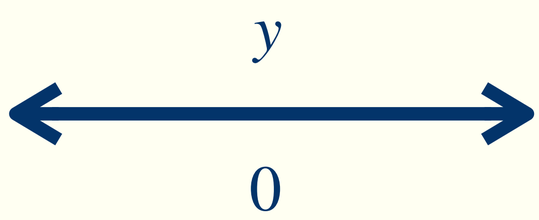

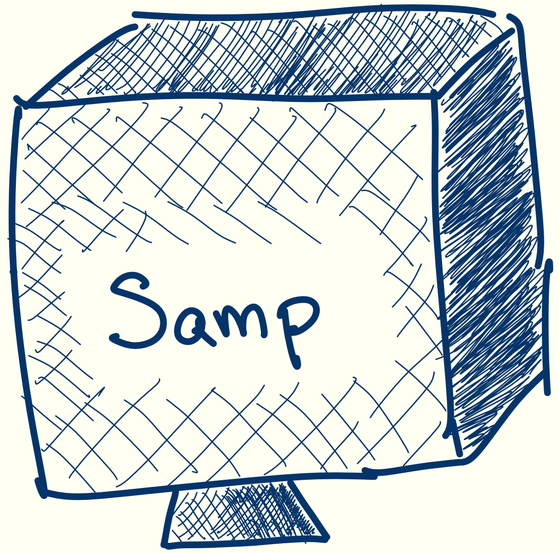

A similar technique exists quantumly, called the compressed oracle technique.

Compressed oracles

In the compressed oracle technique, the sampler interacts with an unknown quantum system. When they query, we instead "condition" on the system being consistent with the answer to their query!

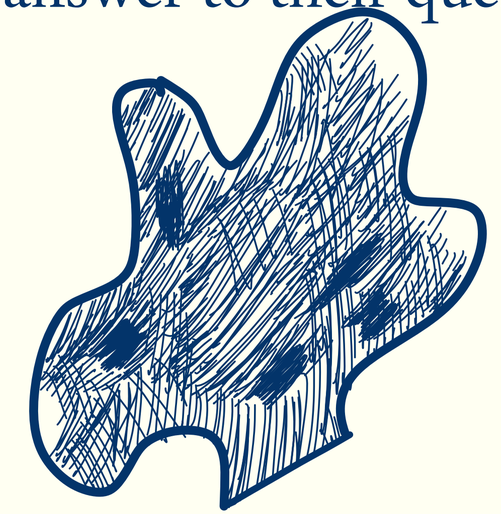

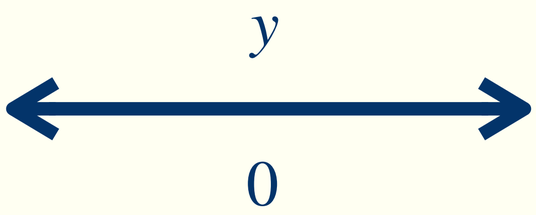

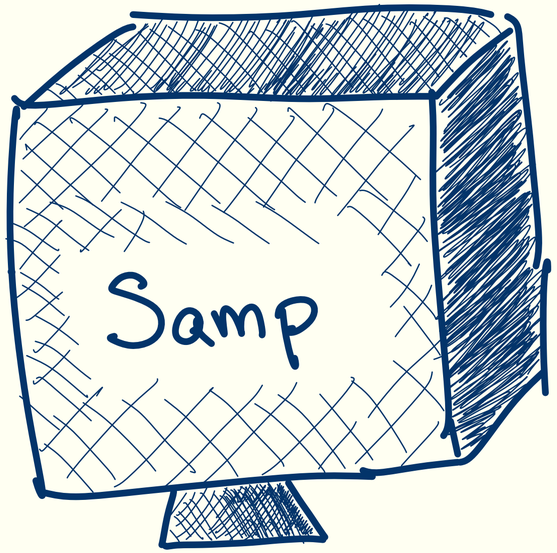

Compressed oracles for spectral Forrelation

Doing this for spectral Forrelation leads to an interesting connection with physics! We model the quantum system is as a collection of ``bosons'' (like, photons or gluons from physics) at random positions (corresponding to S)…

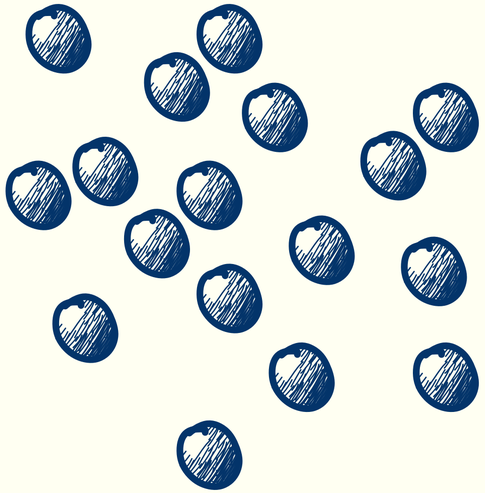

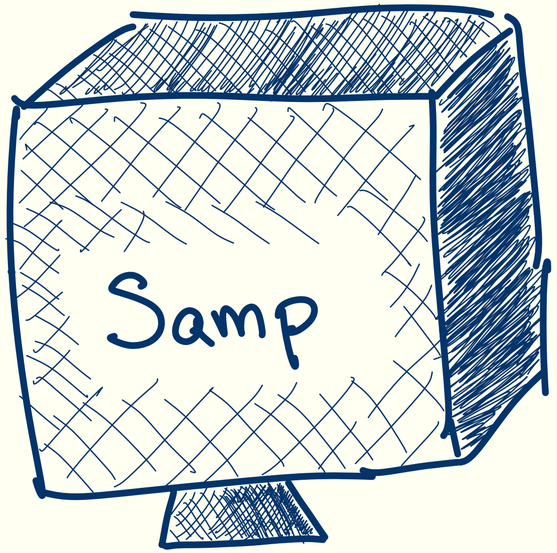

Compressed oracles for spectral Forrelation

Doing this for spectral Forrelation leads to an interesting connection with physics! We model the quantum system is as a collection of ``bosons'' (like, photons or gluons from physics) at random positions (corresponding to S)…

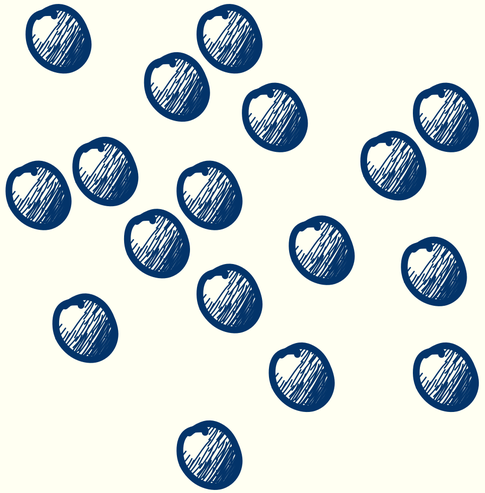

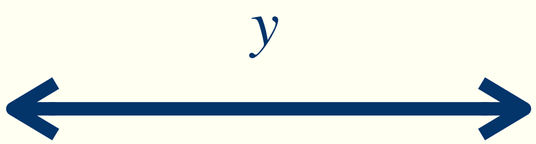

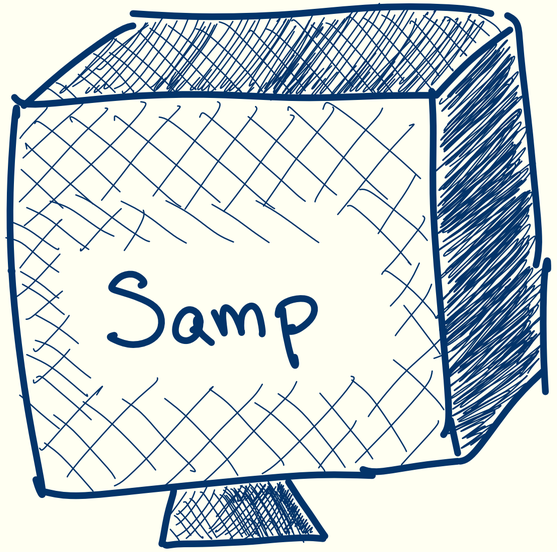

When the algorithm queries U at y, we have to give random pairs of bosons additional momentum y (these form pairs, which I've been told are like Cooper pairs).

Compressed oracles for spectral Forrelation

Doing this for spectral Forrelation leads to an interesting connection with physics! We model the quantum system is as a collection of ``bosons'' (like, photons or gluons from physics) at random positions (corresponding to S)…

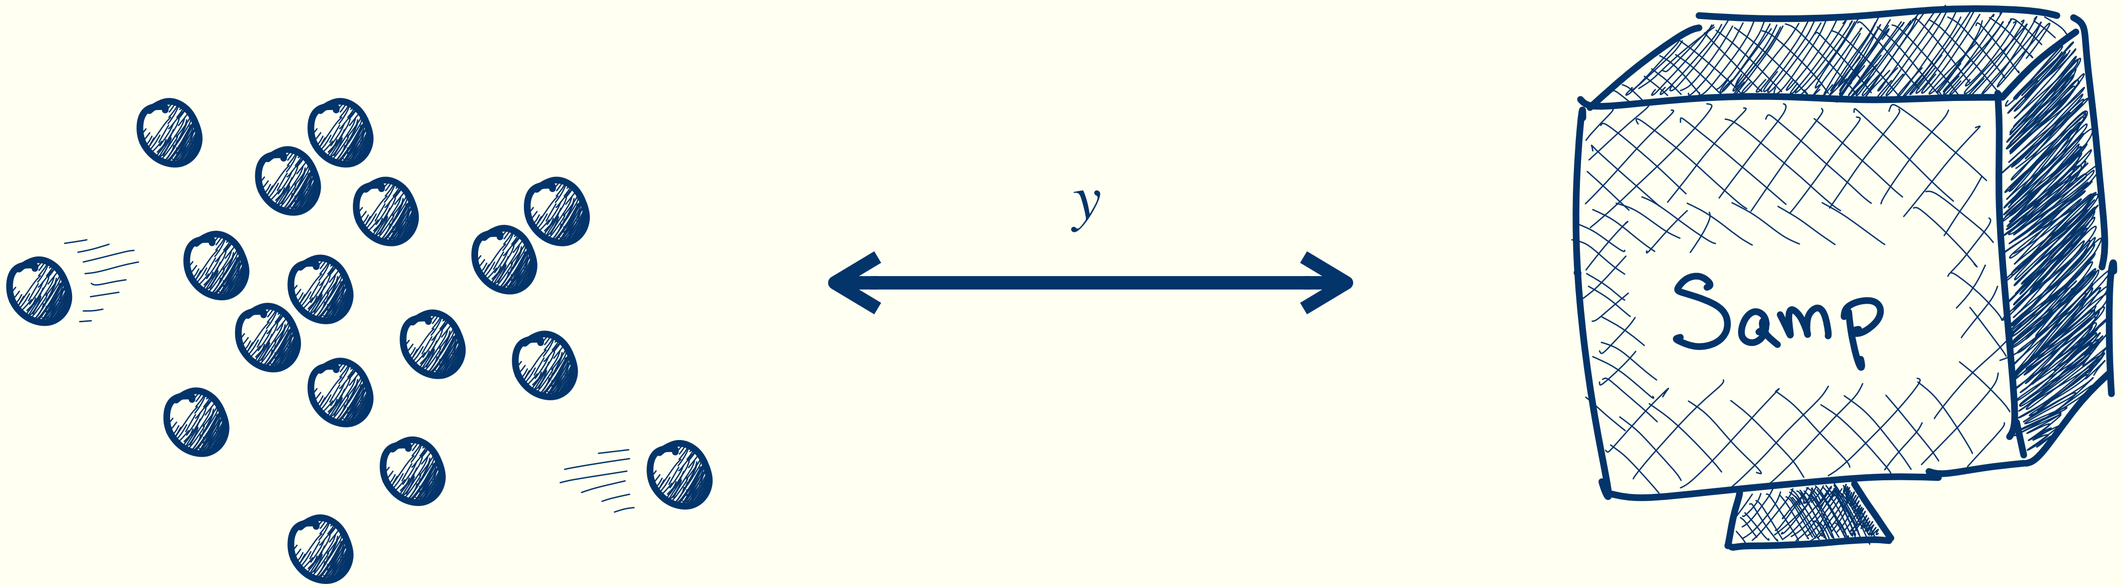

When the algorithm queries U at y, we have to give random pairs of bosons additional momentum y (these form pairs, which I've been told are like Cooper pairs).

Why can't QCMA solve spectral Forrelation?

We show that knowing about the momentum of pairs of bosons tells you nothing about the positions of those bosons.

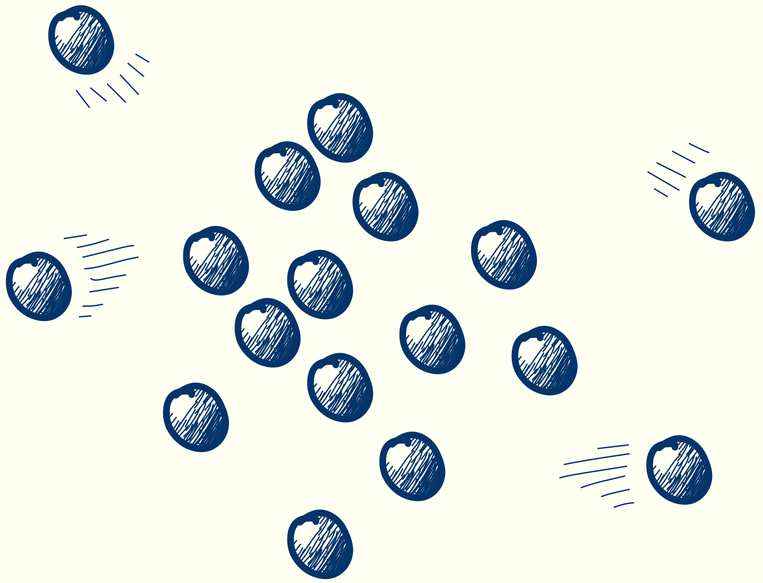

Why can't QCMA solve spectral Forrelation?

We show that knowing about the momentum of pairs of bosons tells you nothing about the positions of those bosons.

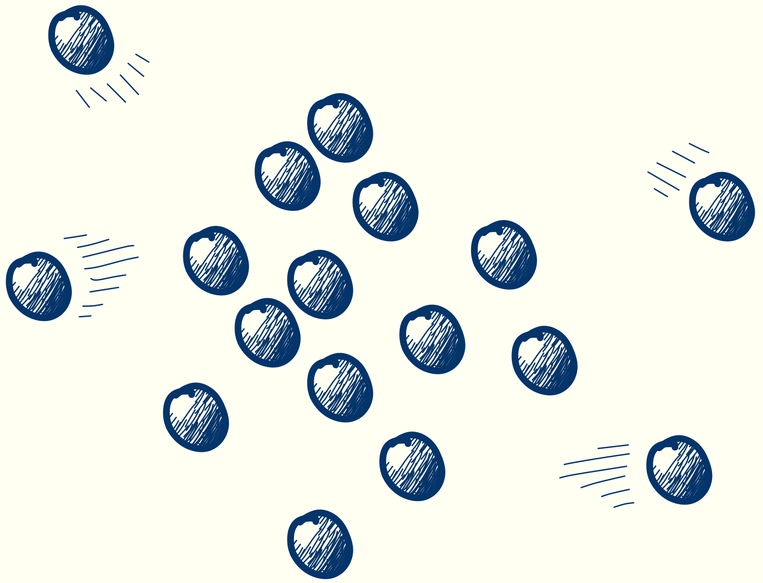

The bosons are not always paired: the sampler might ``double bounce'', i.e., add momentum to a boson that already was moving.

Why can't QCMA solve spectral Forrelation?

We show that knowing about the momentum of pairs of bosons tells you nothing about the positions of those bosons.

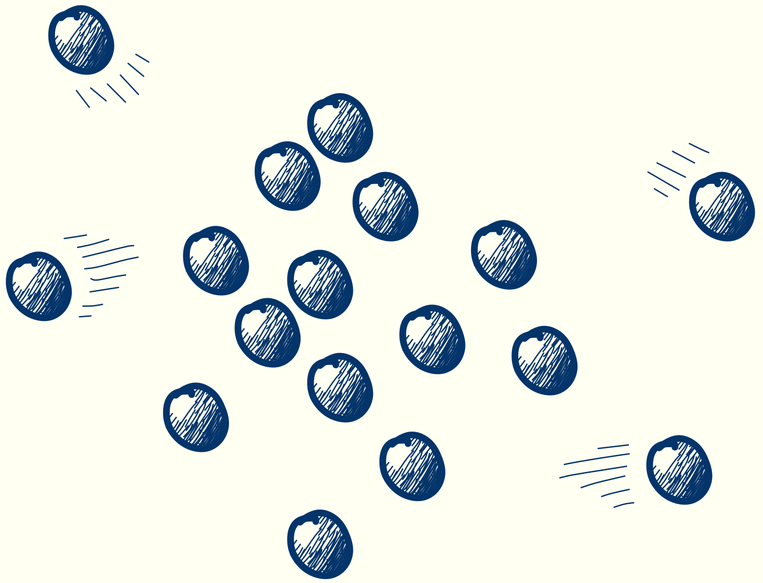

The bosons are not always paired: the sampler might ``double bounce'', i.e., add momentum to a boson that already was moving.

We can show this almost never happens → The sampler implies a contradiction!

Takeaways and future directions

Some problems are only meant to be proven once!

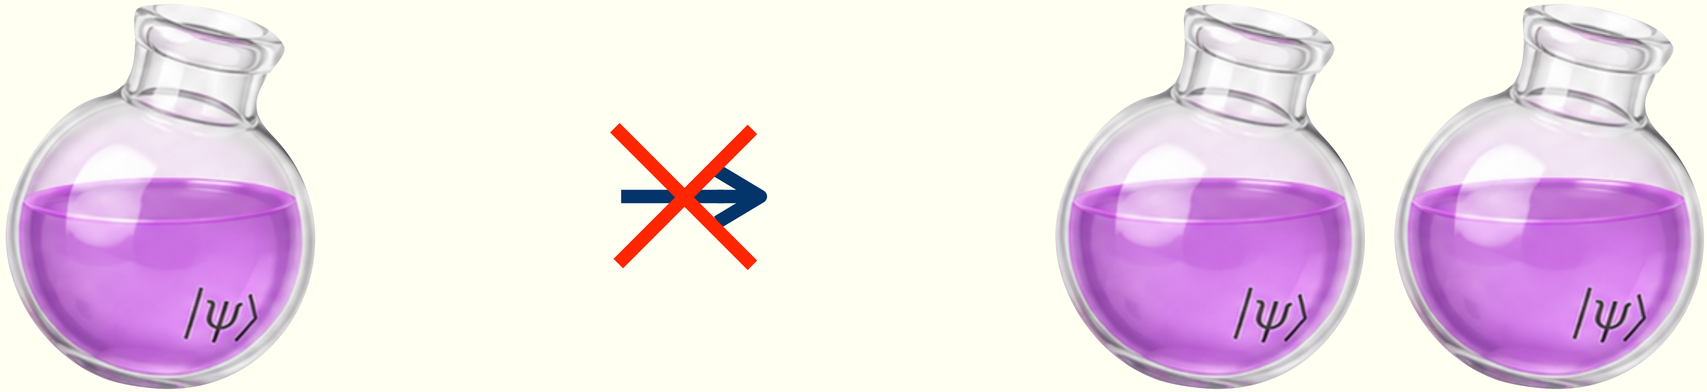

Takeaways and future directions

Some problems are only meant to be proven once!
Our work gives some evidence that quantum proofs might be more powerful than classical states for solving certain problems.

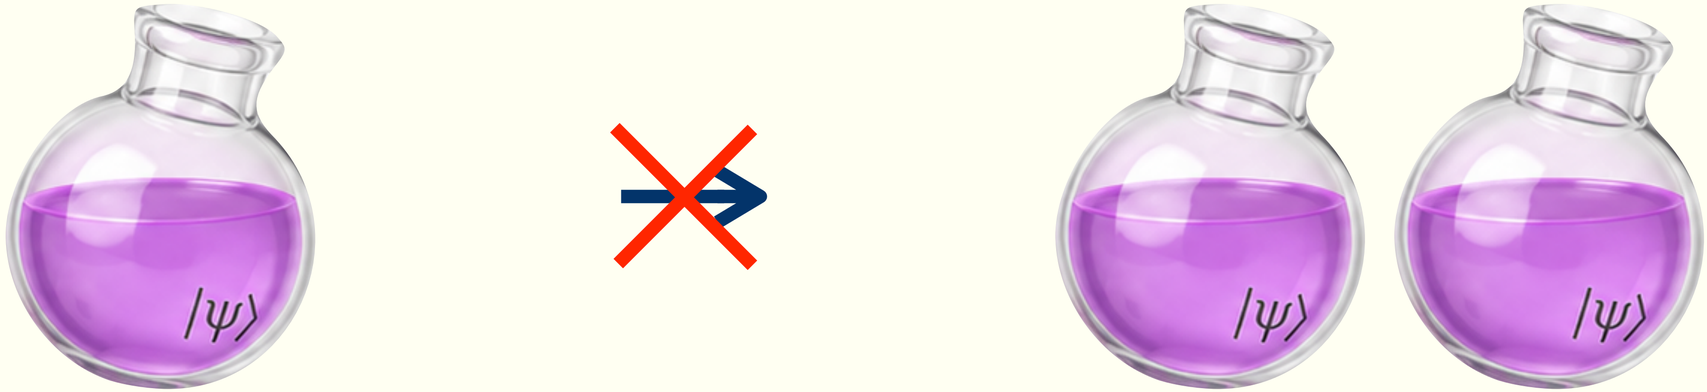

Takeaways and future directions

Some problems are only meant to be proven once!
Our work gives some evidence that quantum proofs might be more powerful than classical states for solving certain problems.

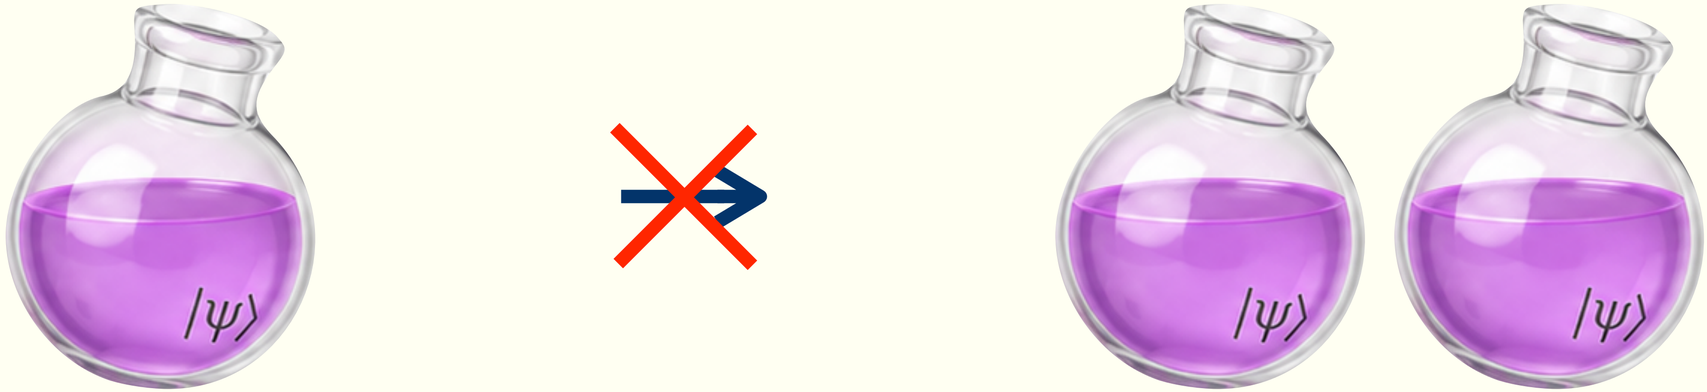

Unfortunately, this power actually prevents us from making them clonable.

If we could copy any proof for spectral Forrelation, we would arrive at a contradiction!

Takeaways and future directions

But, sometimes that is a good thing!

Takeaways and future directions

But, sometimes that is a good thing!

1. We would like to have money that is publicly verifiable, but impossible to counterfeit (clone). Quantum money is the idea that we might be able to use quantum proof to achieve this.

Takeaways and future directions

But, sometimes that is a good thing!

1. We would like to have money that is publicly verifiable, but impossible to counterfeit (clone). Quantum money is the idea that we might be able to use quantum proof to achieve this.

2. We want to be able to publicly verify quantum computation without having to own a quantum computer ourselves. This usually requires getting a quantum computer to verify problems that can't be solved with a classical proof.

Takeaways and future directions

But, sometimes that is a good thing!

1. We would like to have money that is publicly verifiable, but impossible to counterfeit (clone). Quantum money is the idea that we might be able to use quantum proof to achieve this.

2. We want to be able to publicly verify quantum computation without having to own a quantum computer ourselves. This usually requires getting a quantum computer to verify problems that can't be solved with a classical proof.

In these settings, we actually want there to not be a clonable proof! Maybe our ideas can give us insight into how to build quantum money and verify quantum computation!

Takeaways and future directions

But, sometimes that is a good thing!

1. We would like to have money that is publicly verifiable, but impossible to counterfeit (clone). Quantum money is the idea that we might be able to use quantum proof to achieve this.

2. We want to be able to publicly verify quantum computation without having to own a quantum computer ourselves. This usually requires getting a quantum computer to verify problems that can't be solved with a classical proof.

In these settings, we actually want there to not be a clonable proof! Maybe our ideas can give us insight into how to build quantum money and verify quantum computation!

These problems are still really hard though…

Thank you for listening!