Title: Relativistic Quantum (Im)Possibilities

Date: Mar 23, 2011 02:00 PM

URL: http://pirsa.org/11030106

Abstract: Many fundamental results in quantum foundations and quantum information theory can be framed in terms of information-theoretic tasks that are provably (im)possible in quantum mechanics but not in classical mechanics. For example, Bell's theorem, the no-cloning and no-broadcasting theorems, quantum key distribution and quantum teleportation can all naturally be described in this way. More generally, quantum cryptography, quantum communication and quantum computing all rely on intrinsically quantum information-theoretic advantages.
Much less attention has been paid to the information-theoretic power of relativistic quantum theory, although it appears to describe nature better than quantum mechanics. This talk describes some simple information-theoretic tasks that distinguish relativistic quantum theory from quantum mechanics and relativistic classical physics, and a general framework for defining tasks that includes all previously known (im)possibility theorems and raises many open questions. This suggests a new way of thinking about relativistic quantum theory, and a possible new approach to defining non-trivial relativistic quantum theories rigorously. I also describe some simple and surprisingly powerful applications of these ideas to cryptography, including a new secure scheme for simultaneously committing to and encrypting a prediction and ways of securely "tagging" an inaccessible object so as to guarantee its position.

Relativistic Quantum (Im)Possibilities

(arxiv:1101.4612, 1008.2147, 1008.5380, 1101.4620 and 1102.2816)

Adrian Kent

Perimeter Institute

and

Centre for Quantum Information and Foundations, DAMTP, University of Cambridge

talk @PI140023032011

Well-known Quantum (Im)Possibilities

- **Bell's theorem and extensions**: quantum mechanics allows classically impossible output correlations (outcomes) from separated unknown random classical inputs (measurement choices).

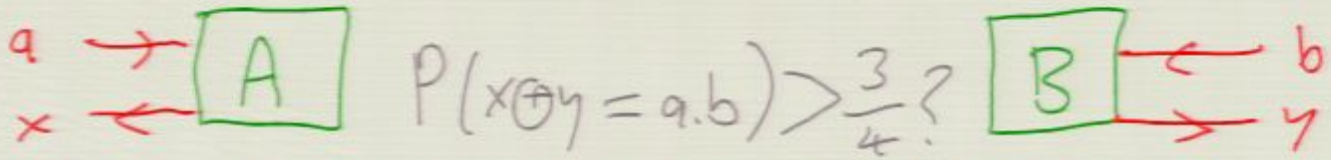

- **quantum no-cloning**: quantum mechanics forbids faithful copying of an unknown random quantum input state, whereas an unknown random classical input state can in principle be copied to arbitrary precision. (Cf. also quantum no-broadcasting, no-deletion, ...)

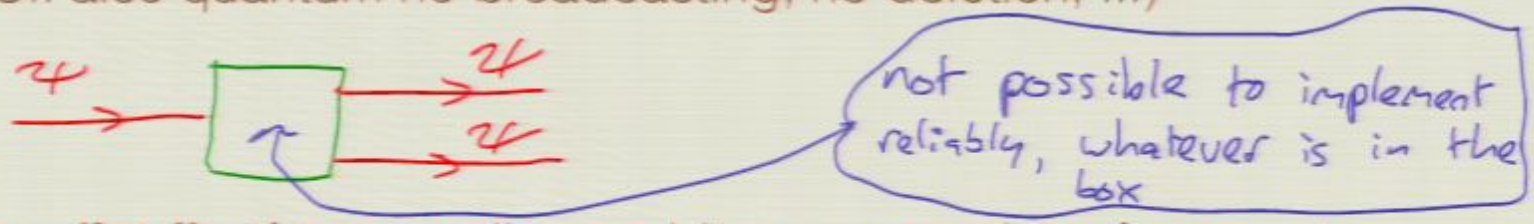

- **Quantum key distribution**: qm allows arbitrary expansion of a secret random string between separated labs, even when eavesdroppers are in between -- no expansion is possible in classical physics.

More Quantum (Im)Possibilities

- **No-signalling**: Whatever state A and B share, if a random bit is input at A, B cannot learn it by actions on B's state -- in classical mechanics, quantum mechanics, or relativistic quantum theory.

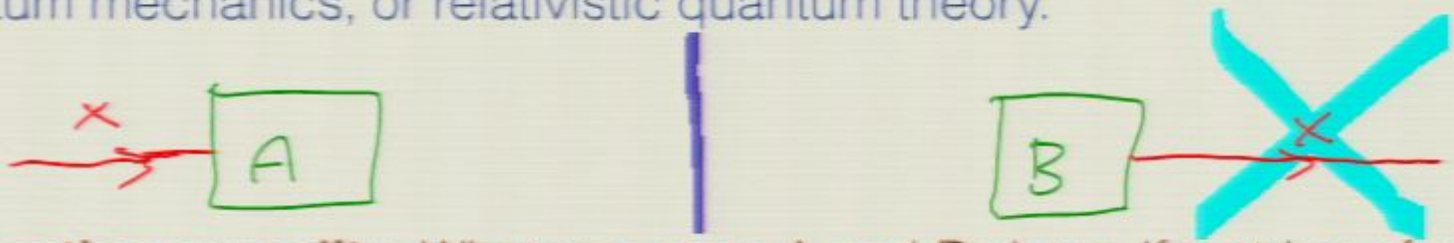

- **Information causality**: Whatever state A and B share, if a string of n random bits are input at A, and A sends m<n classical bits to B, then B can learn no more than m bits of A's string -- in classical mechanics, quantum mechanics, relativistic quantum theory but **not true** in all no-signalling theories. (Pawlowski et al, Nature 2009).

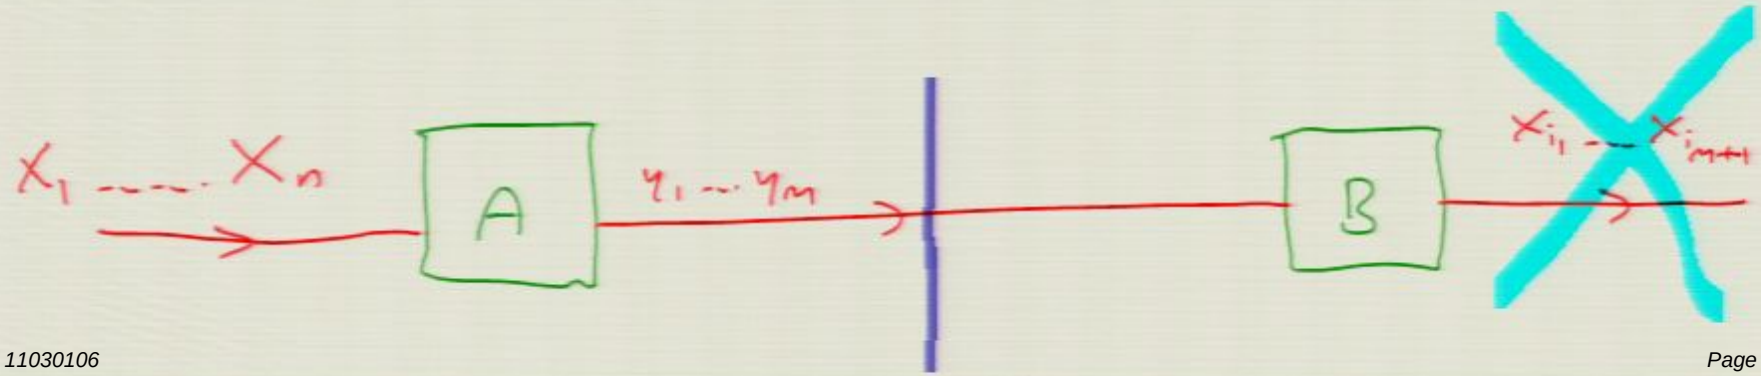

Relativistic quantum (im)possibilities: a general framework

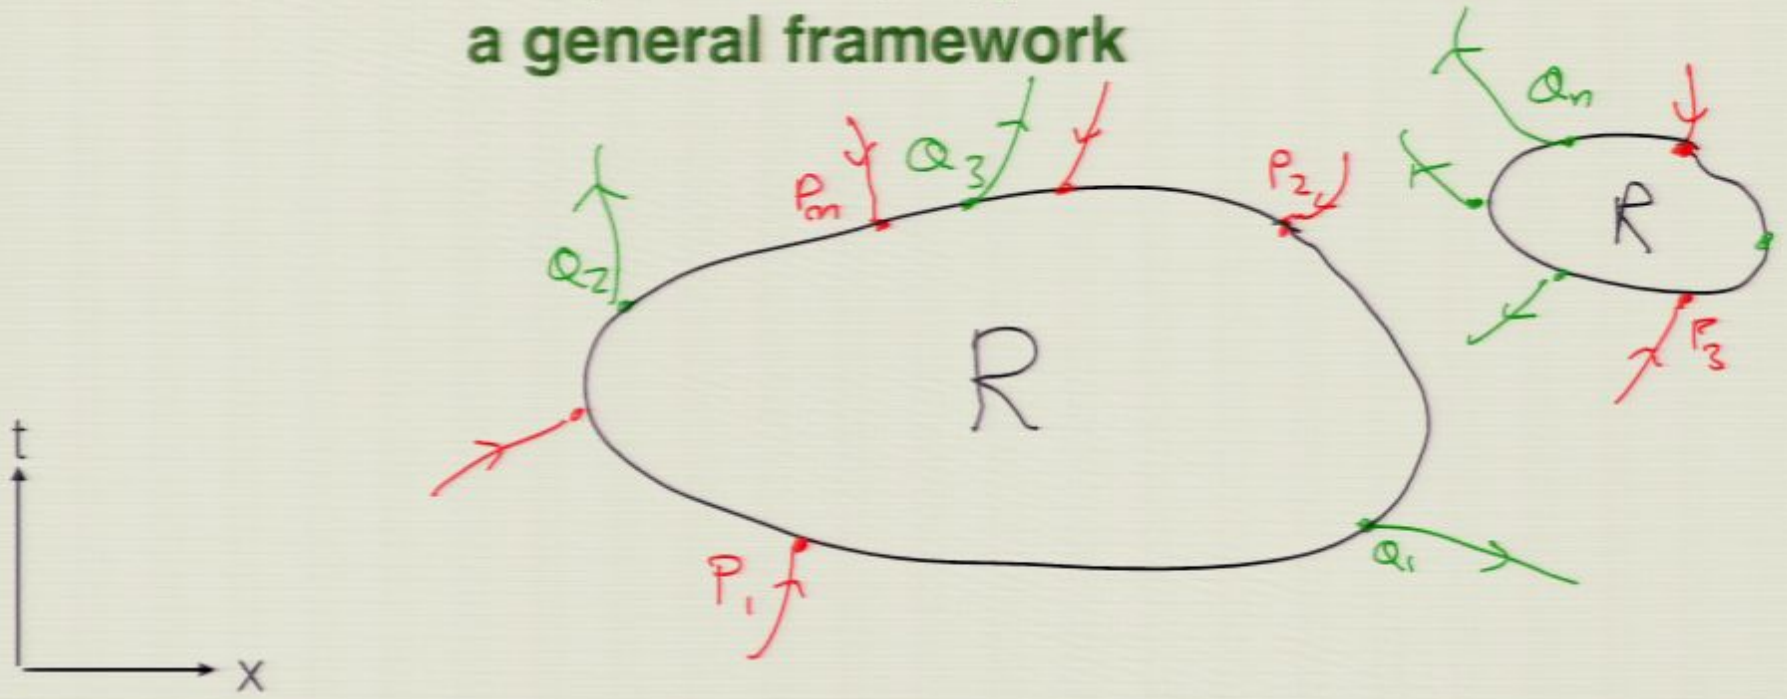

Given inputs at P_i , where both the inputs and the points are drawn from specified distributions and control of the region R (which need not necessarily be connected), and queries at Q_j -- points again drawn from specified distributions -- can one produce outputs there with given relationships to the inputs?

What do we mean by relativistic quantum theory?

- We don't have a rigorous definition of interacting relativistic quantum field theories in Minkowski space.
- But we do have a clear intuitive understanding of non-trivial operations we can carry out on physical quantum states in Minkowski space.
- So we can prove (im)possibility theorems based on standard assumptions -- which (i) are interesting in their own right, (ii) have significant applications for quantum computing and cryptography, and (iii) maybe also suggest a path to axiomatizing relativistic quantum theory.

Assumed properties of relativistic quantum theory

- A global quantum state is defined on any space-like hypersurface, and these states are related by unitary maps, a la Tomonaga-Schwinger.

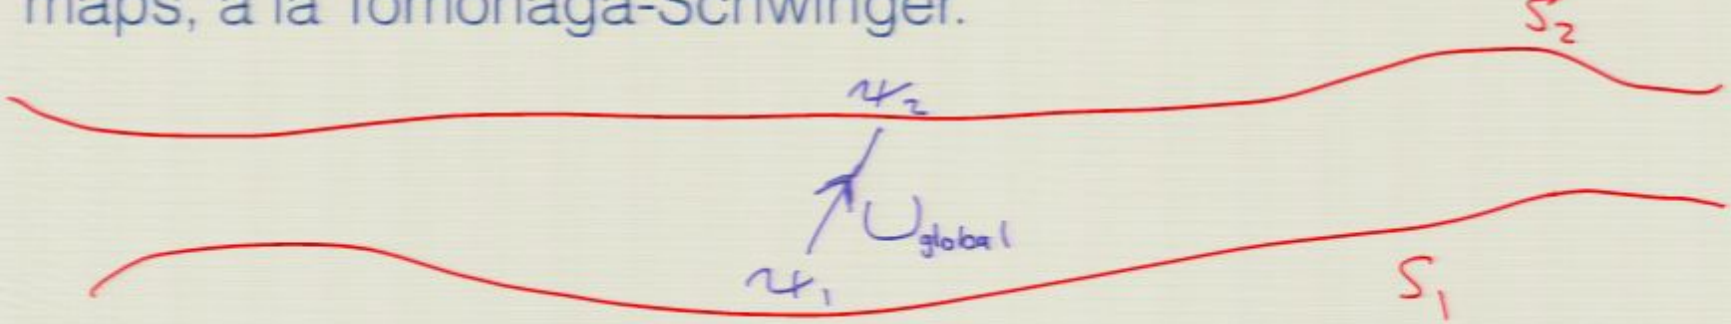

- Any physically allowed local unitary defined on a pointlike subsystem can be implemented with arbitrarily small timelike delay.

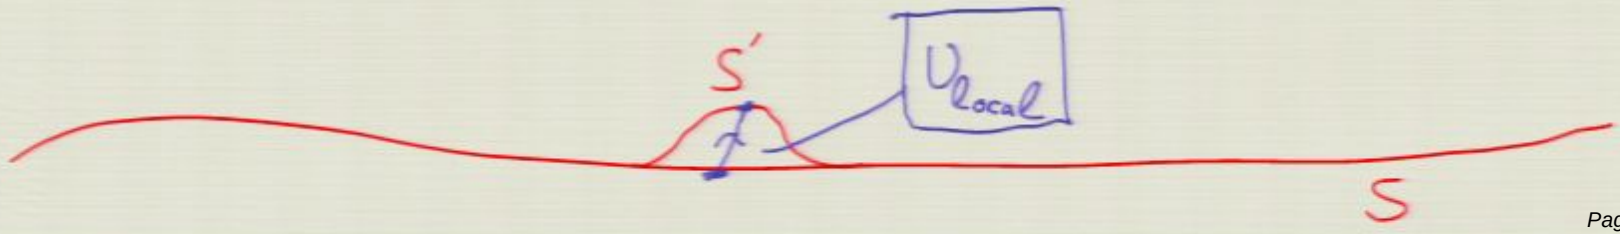

An example of a relativistic quantum impossibility: Summoning a quantum state

Consider two agencies, Alice and Bob, with independent secure networks and (here we idealise for now) representatives everywhere in space-time.

Alice prepares a localised physical state unknown to Bob and gives him it at point P.

At some point Q, in the causal future of P, not known in advance by Bob, Alice summons -- i.e. asks Bob to return -- the state.

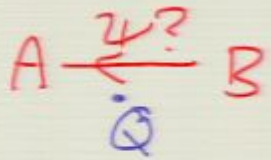

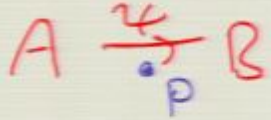

An example of a relativistic quantum impossibility: Summoning a quantum state

Consider two agencies, Alice and Bob, with independent secure networks and (here we idealise for now) representatives everywhere in space-time.

Alice prepares a localised physical state unknown to Bob and gives him it at point P.

At some point Q, in the causal future of P, not known in advance by Bob, Alice summons -- i.e. asks Bob to return -- the state.

In principle, with arbitrarily short delay, Bob can comply if the underlying theory is quantum mechanics in Galilean space-time, or classical mechanics in Minkowski space-time.

However, as we will show, given an unknown quantum state in Minkowski space-time, he cannot comply.

Summoning in classical theories

Given an unknown classical state at point P in Minkowski space, Bob can (in principle) measure it precisely, broadcast the information in all directions, and reconstruct the state at any point Q in the causal future of P -- and so comply with Alice's summons.

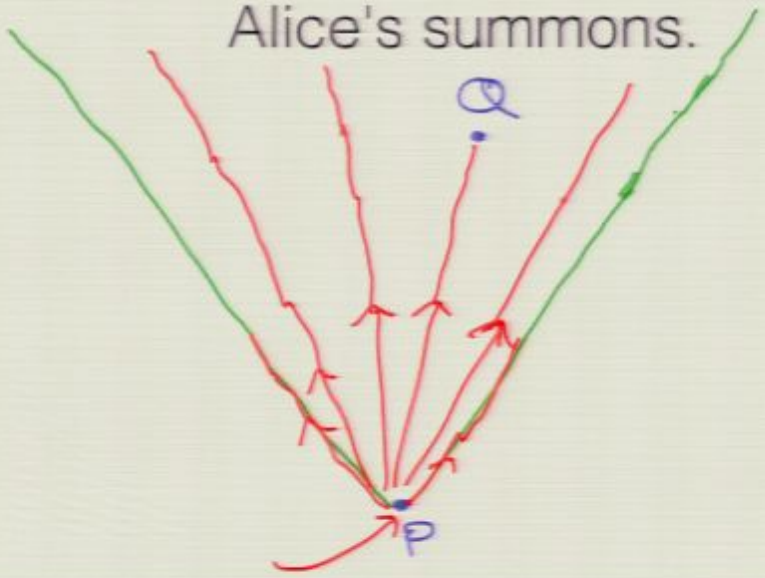

Summoning in non-relativistic quantum mechanics

Given an unknown quantum state rho at a point P=(x,t) in Galilean space-time, Bob can hold the state at position x, wait for a summons at Q=(y,t') (where t'>t), instantaneously send a signal to (x,t') requesting the state, and instantaneously send the state back to Q, and so comply with the summons.

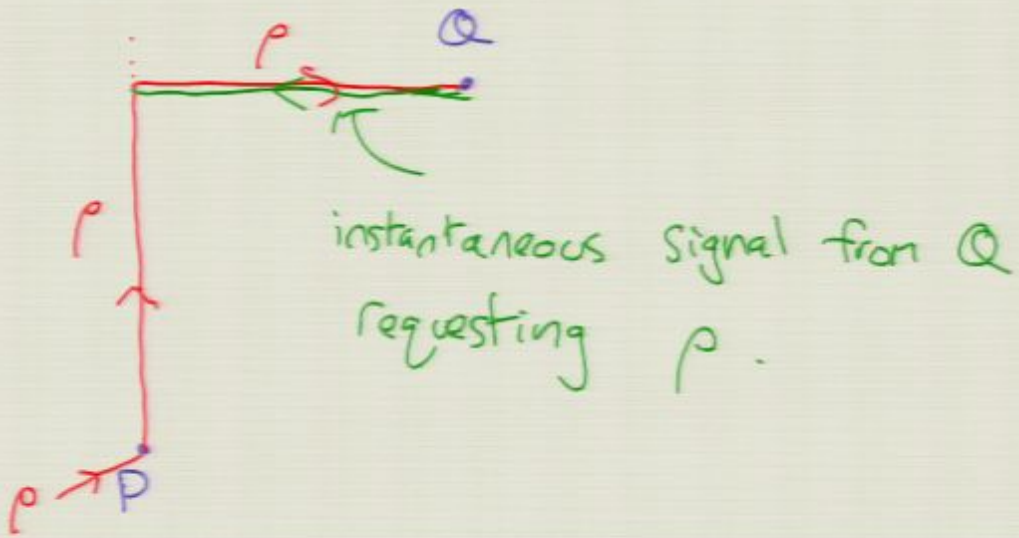

Summoning in classical theories

Given an unknown classical state at point P in Minkowski space, Bob can (in principle) measure it precisely, broadcast the information in all directions, and reconstruct the state at any point Q in the causal future of P -- and so comply with Alice's summons.

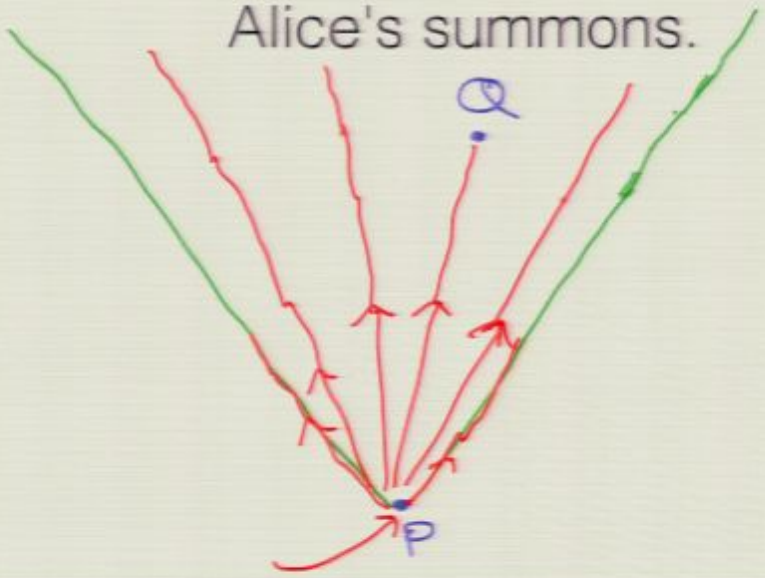

Summoning in non-relativistic quantum mechanics

Given an unknown quantum state rho at a point P=(x,t) in Galilean space-time, Bob can hold the state at position x, wait for a summons at Q=(y,t') (where t'>t), instantaneously send a signal to (x,t') requesting the state, and instantaneously send the state back to Q, and so comply with the summons.

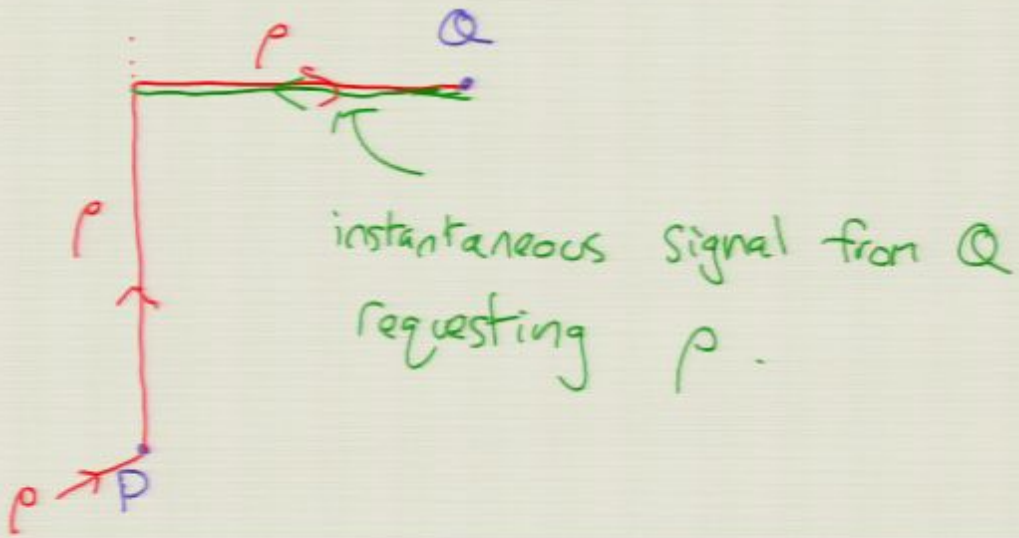

No summoning in relativistic quantum theory

Given an unknown quantum state rho at point P in Minkowski space-time, Bob cannot precisely identify it or copy it (because of the no-cloning theorem).

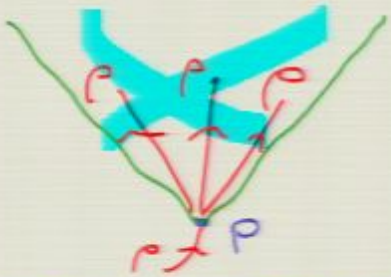

If he holds it at a possible summoning point Q in the causal future of P, he cannot send it to another space like separated possible summoning point Q' (because of the no-superluminal signalling principle).

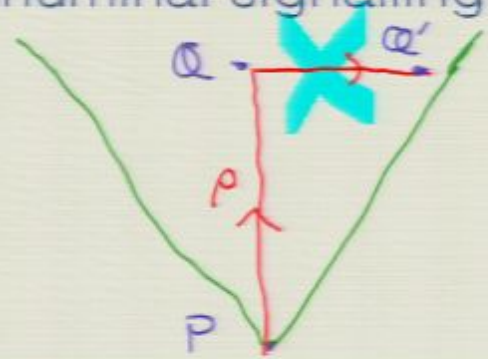

No summoning in relativistic quantum theory

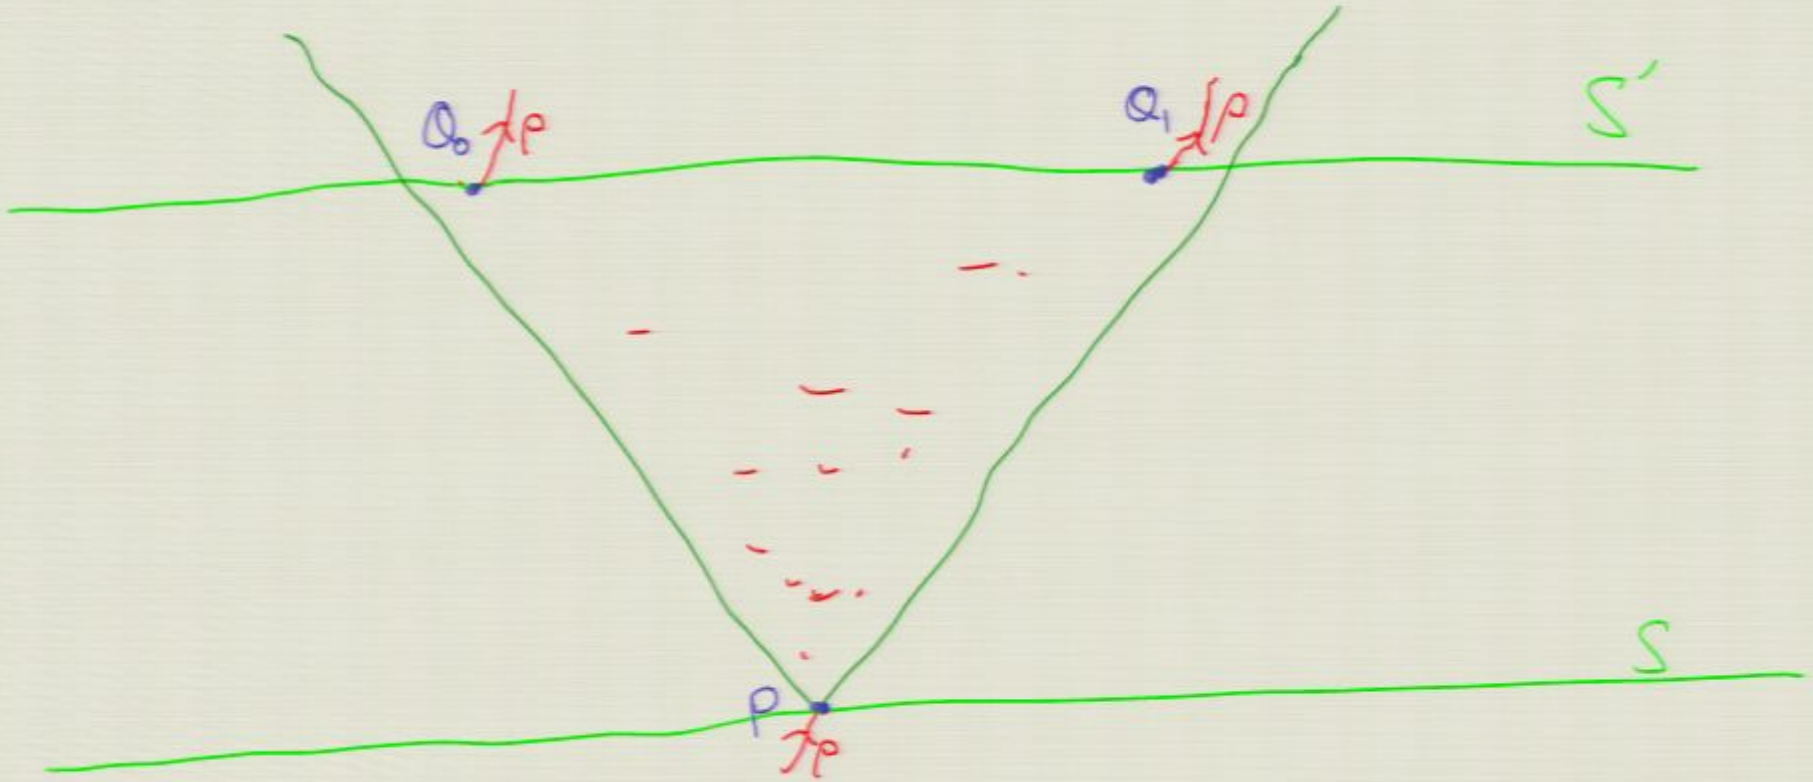

More generally, we can use no-cloning and no-signalling to prove that whatever strategy he follows, Bob cannot generally comply with a summons.

No approximate summoning in relativistic quantum theory

- A more realistic version of the task would allow Bob some time (more precisely, a prescribed space-time region within which) to comply.

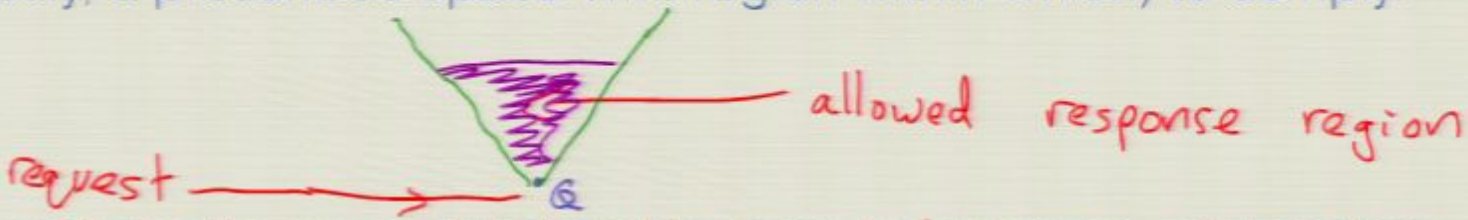

- Also, realistically, we could allow him margin for errors - ok to return approximately the same state (i.e. with fidelity close to 1 to the original)

- Under these definitions, summoning is realistically (not just ideally) possible in non-relativistic quantum mechanics or relativistic classical mechanics.

- But there are non-trivial bounds on the fidelity of approximate cloning. Removing our idealizations doesn't affect the main conclusion. No-approximate-cloning plus no-signalling imply no-approximate-summoning in relativistic quantum theory.

No-summoning and quantum foundations

- The no-summoning theorem follows from the no-cloning theorem and the no-signalling principle, but not from either alone.
- Like the no-cloning theorem, it is mathematically elementary.
- But it says something new about the relationship between quantum theory and relativity: the first (?) example of a simple information-related task that distinguishes relativistic quantum theory from non-relativistic qm and relativistic classical physics.
- Whereas Bell's theorem, no-cloning, no-broadcasting, no-signalling, information causality, ... all apply to non-relativistic qm as well as to relativistic quantum theory.
- And, on the other hand, the impossibility of instantaneous measurement of non-localised states holds in classical relativistic theories as well as in relativistic quantum theory.

Relativistic quantum (im)possibilities: a general framework

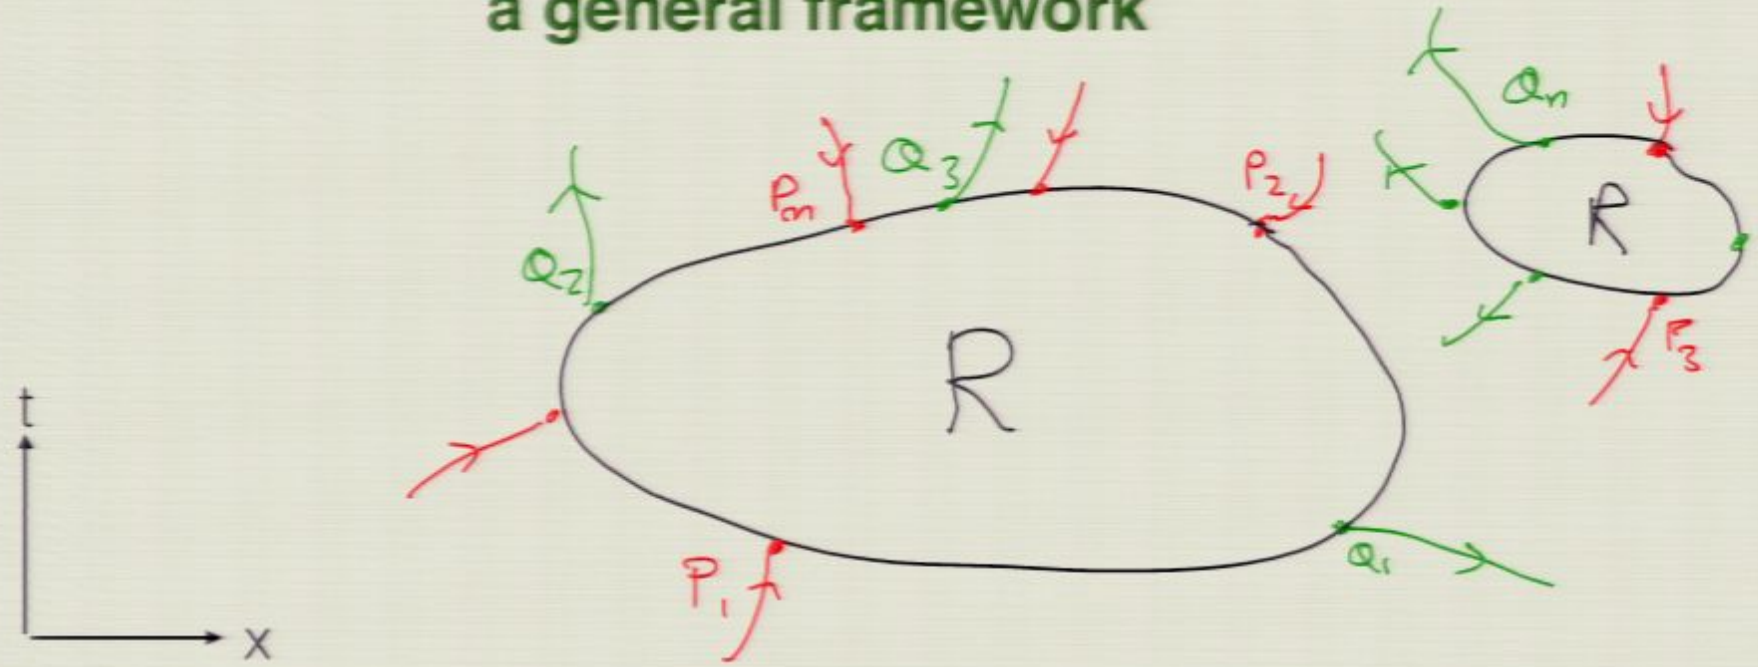

Given inputs at P_i , where both the inputs and the points are drawn from specified distributions and control of the region R (which need not necessarily be connected), and queries at Q_j -- points again drawn from specified distributions -- can one produce outputs there with given relationships to the inputs?

Rewriting quantum (im)possibilities in our general framework

- **Bell's theorem and extensions**:

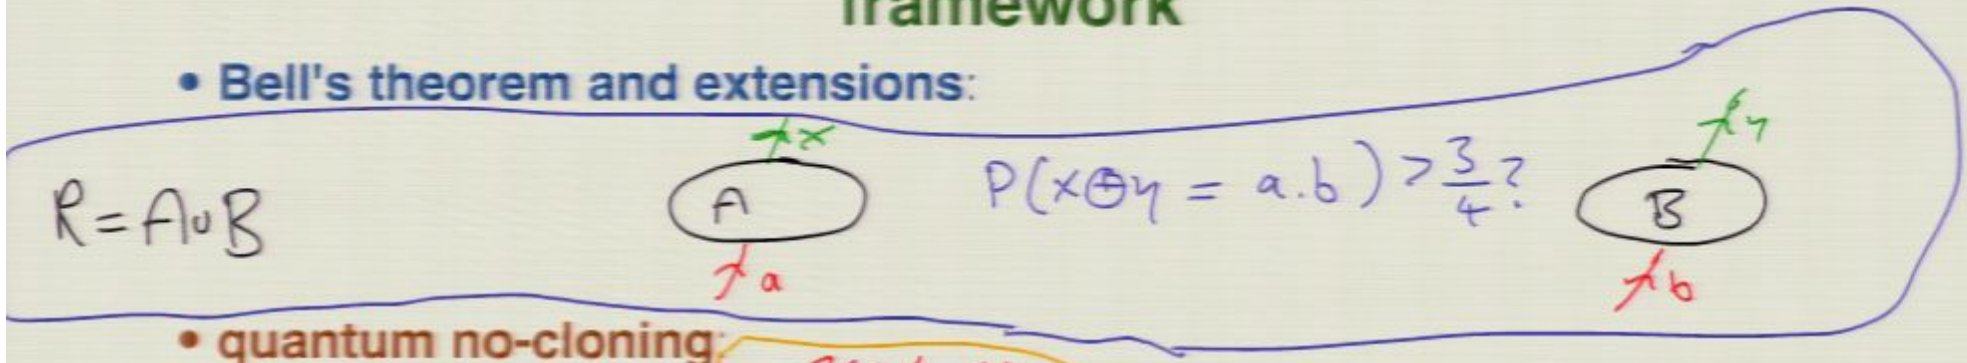

- **quantum no-cloning**:

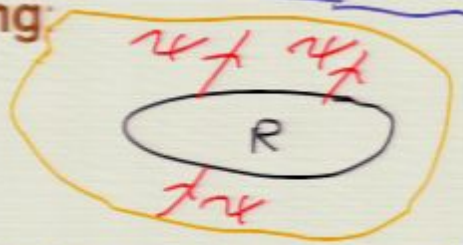

- **No-signalling and information causality**:

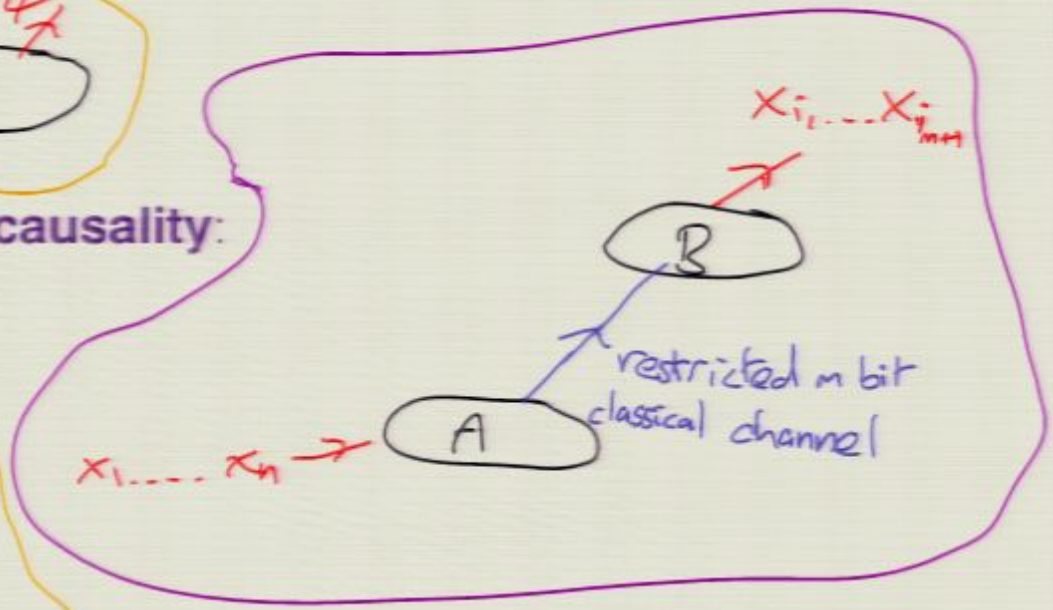

- **No-summoning**:

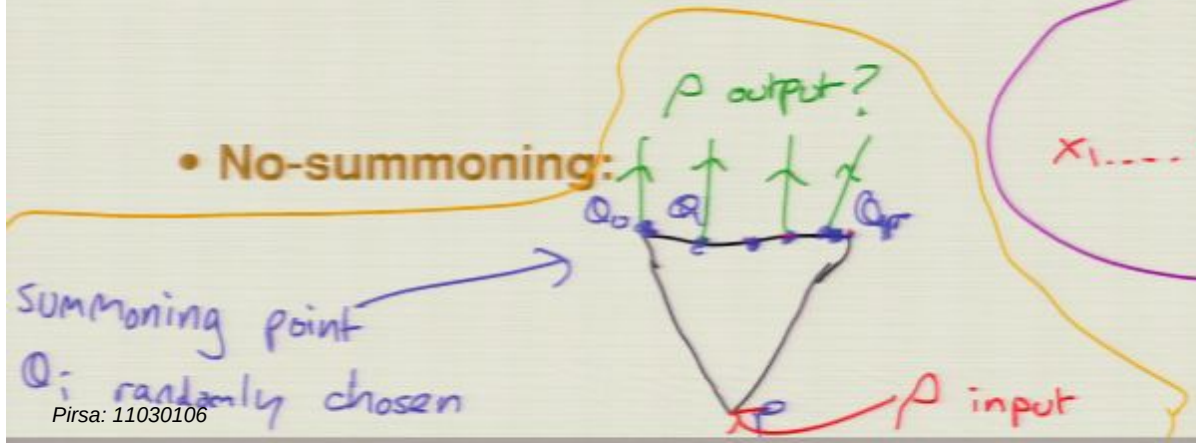

No-summoning and quantum cryptography

(arxiv:1101.4620, see also 1102.2816)

One dramatic example of the power of the no-summoning theorem is a simple and practical solution to the long-standing problem of unconditionally secure quantum bit commitment.

Bit commitment: Alice wants to make an encrypted prediction. She needs a guarantee that the recipient (Bob) cannot decrypt her prediction until she gives him a key - extra data. He needs a guarantee that she is genuinely committed and cannot change her prediction, for instance by having two different keys that will decrypt two different predictions. They both ideally want these guarantees based only on the laws of physics.

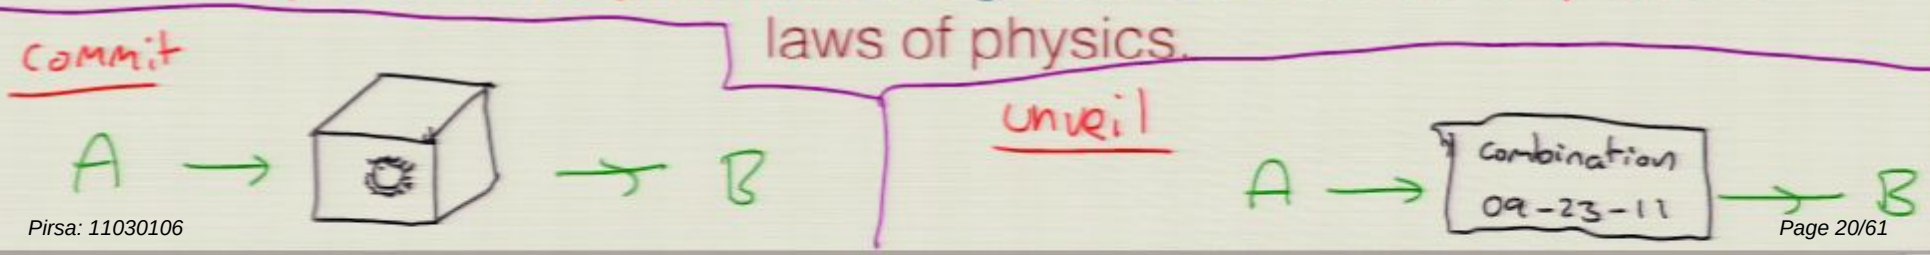

A Brief History of Bit Commitment 1

- Much studied by classical cryptographers, who created protocols provably hard to cheat given computational hardness assumptions. No perfectly secure protocol exists in standard classical crypto models.

- First quantum bit commitment scheme proposed by Bennett-Brassard (1984), who also showed that Alice can cheat it if she can create error-free entangled states and has a perfect quantum memory.

- Much initial optimism that unconditionally secure quantum bit commitment schemes exist; Brassard et al. (BCJL) claimed (1993) to have proven such a scheme secure.

- Mayers and Lo-Chau (1997) showed not only how to break the BCJL scheme but also that all quantum bit commitment schemes of an apparently general type were insecure.

- Mayers conjectured this no-go result also applies to schemes using relativistic signalling constraints.

A Brief History of Bit Commitment 2

- Mayers (1997) conjectured this no-go result also applies to schemes using relativistic signalling constraints.

- Mayers' conjecture turned out to be incorrect (AK, 1999, 2005): schemes exist using relativistic constraints, two sets of separated lab and continuing classical communications that are provably not vulnerable to Mayers-Lo-Chau attacks, and provably secure against all classical attacks. Security against general quantum attacks remains a conjecture.

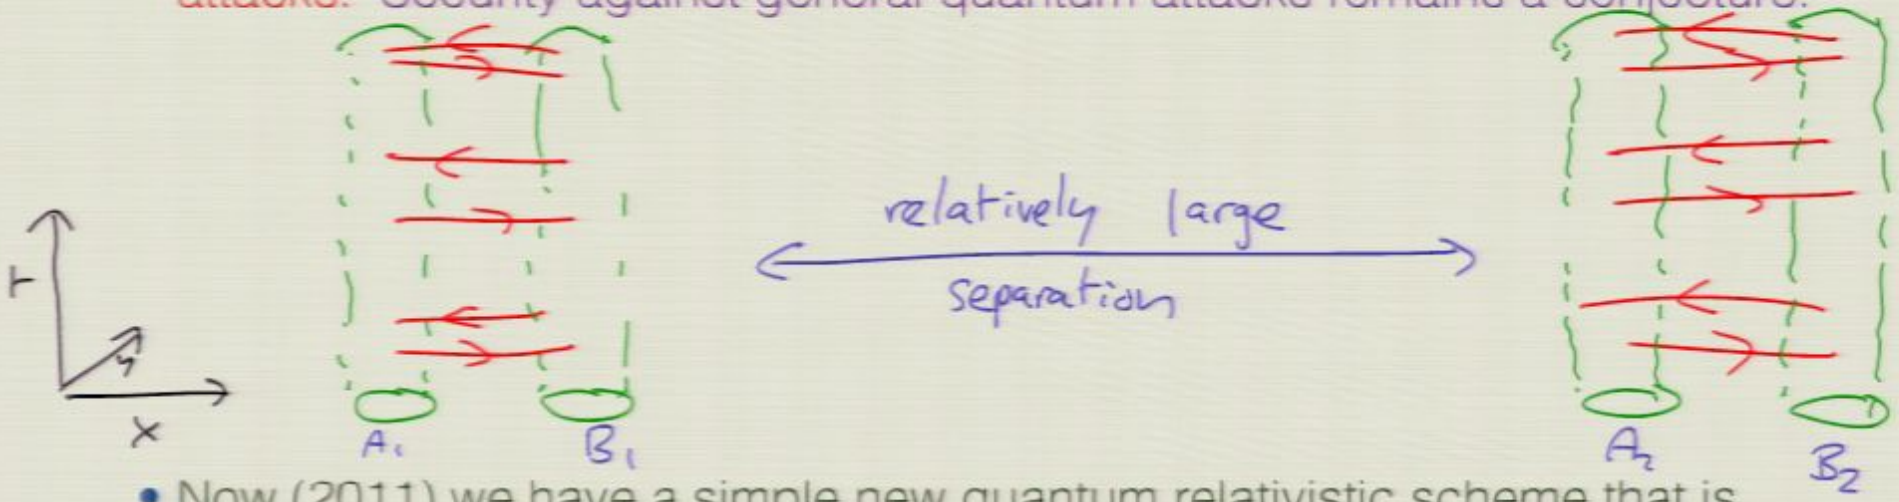

- Now (2011) we have a simple new quantum relativistic scheme that is provably secure, based on the no-summoning theorem.

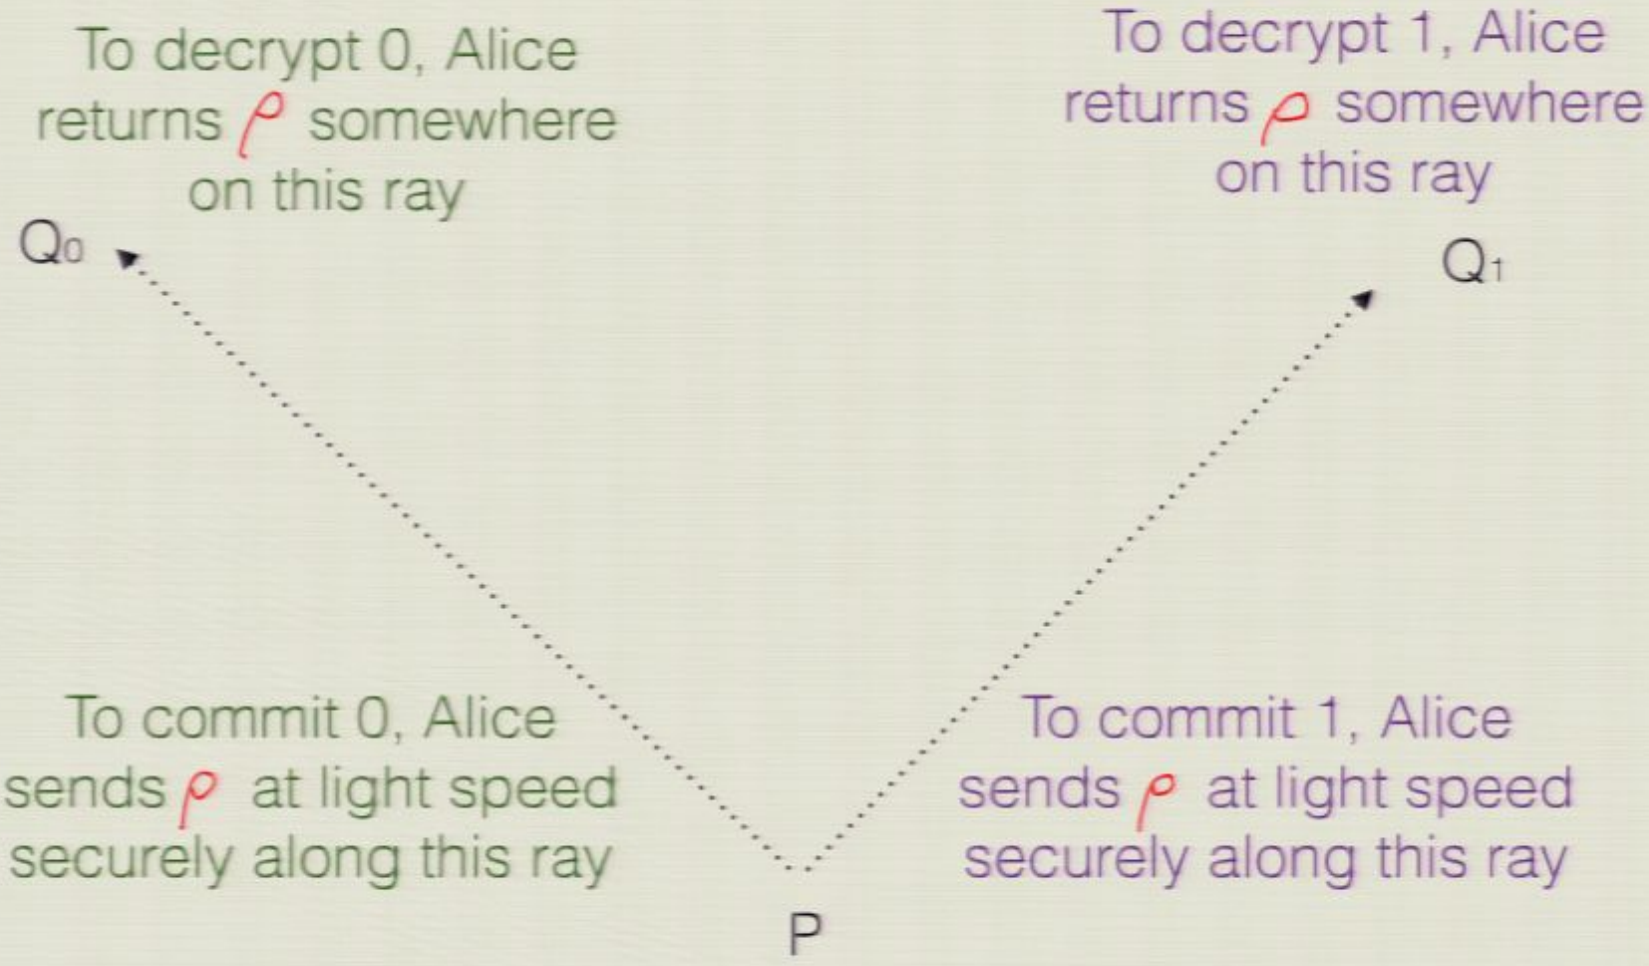

Bob gives Alice state ρ at P

Security against Bob: ensured since Alice sends the state securely (either because she controls a region around the relevant light rays, or e.g. via teleportation)

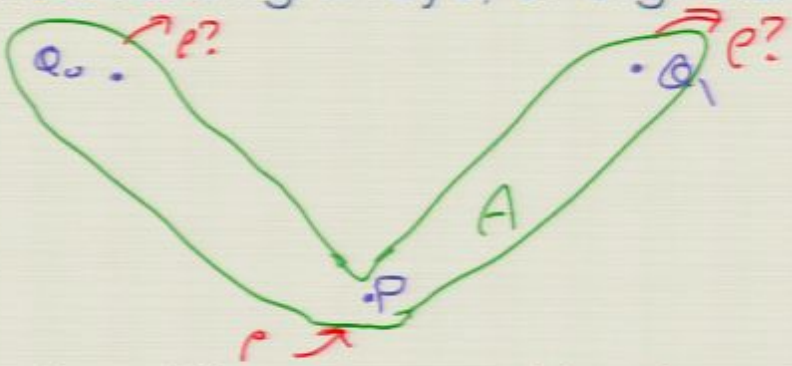

Security against Alice: ensured by the no-summoning theorem -- she cannot return ρ independently at points on both light rays.

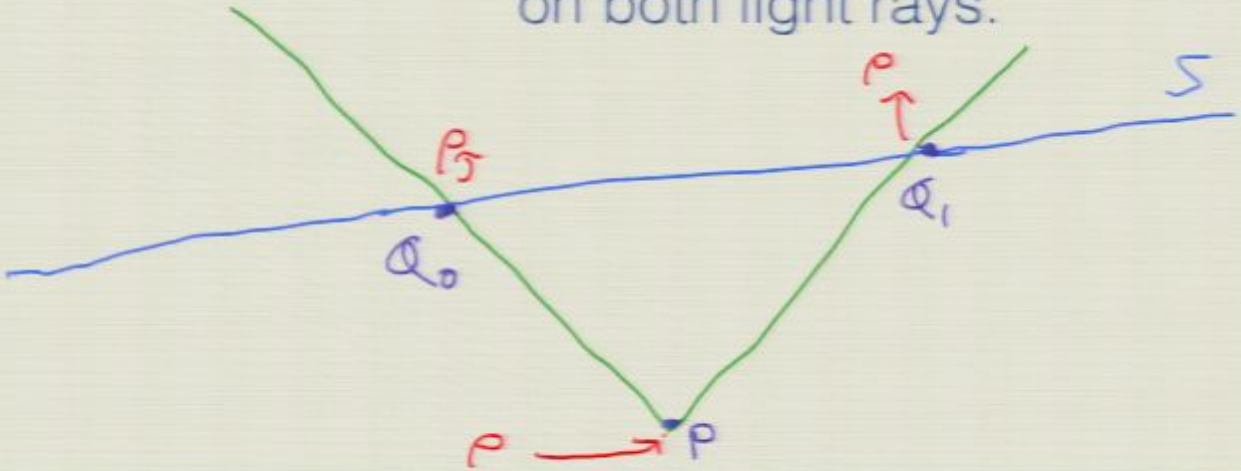

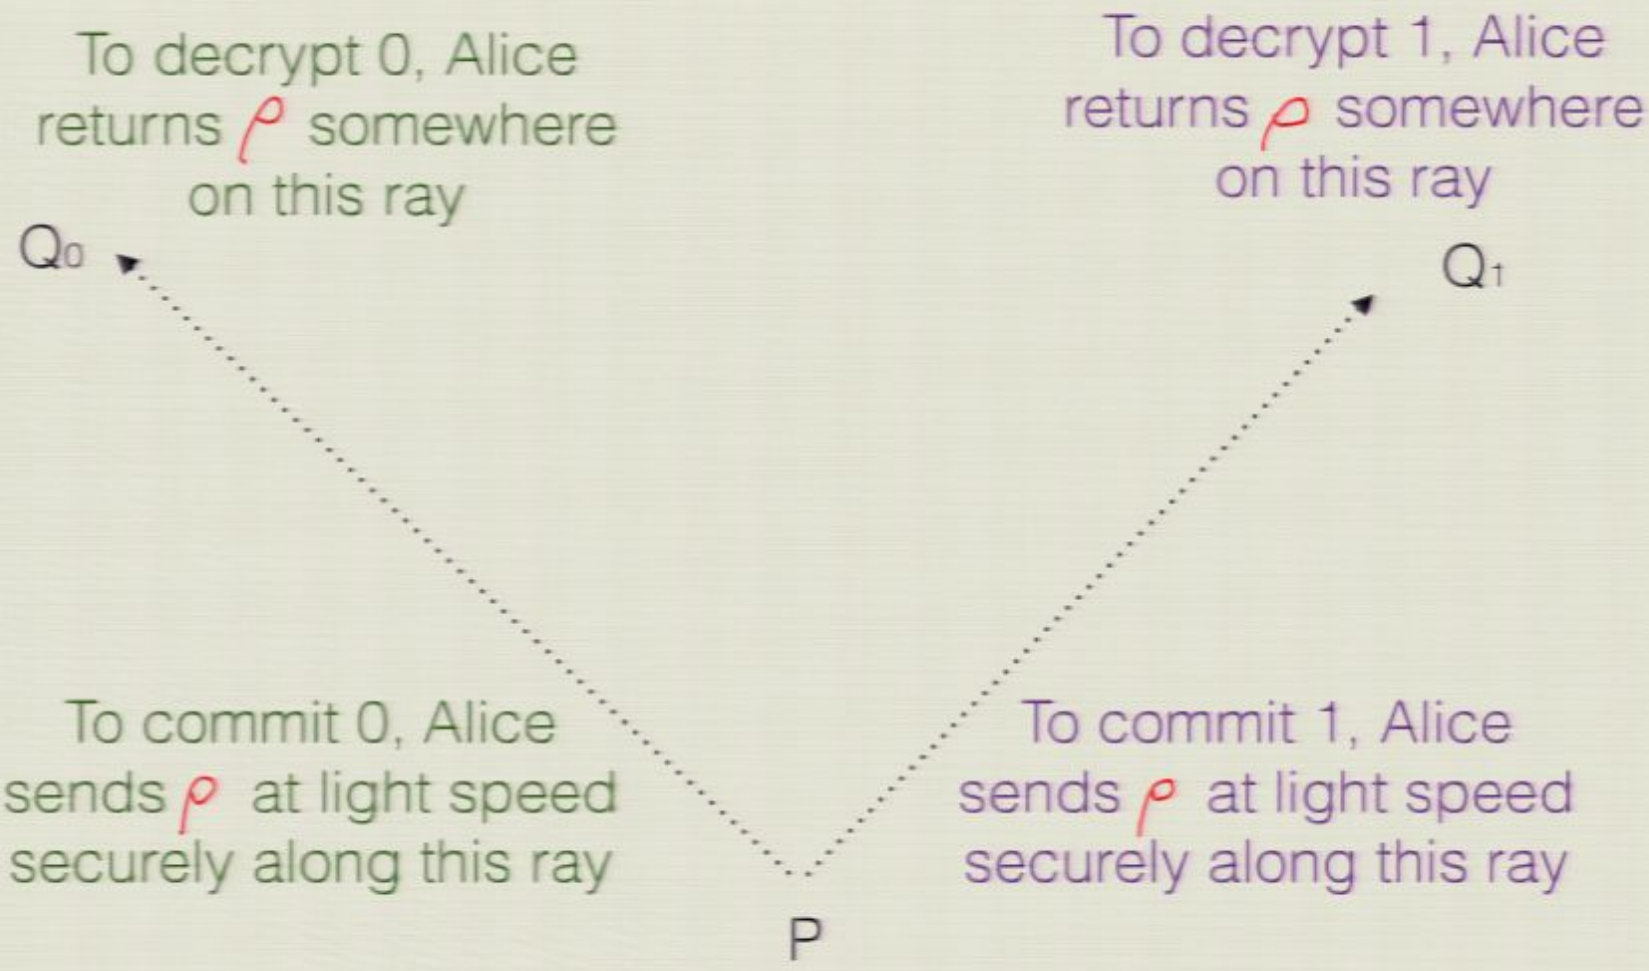

Bob gives Alice state ρ at P

A Brief History of Bit Commitment 2

- Mayers (1997) conjectured this no-go result also applies to schemes using relativistic signalling constraints.

- Mayers' conjecture turned out to be incorrect (AK, 1999, 2005): schemes exist using relativistic constraints, two sets of separated lab and continuing classical communications that are provably not vulnerable to Mayers-Lo-Chau attacks, and provably secure against all classical attacks. Security against general quantum attacks remains a conjecture.

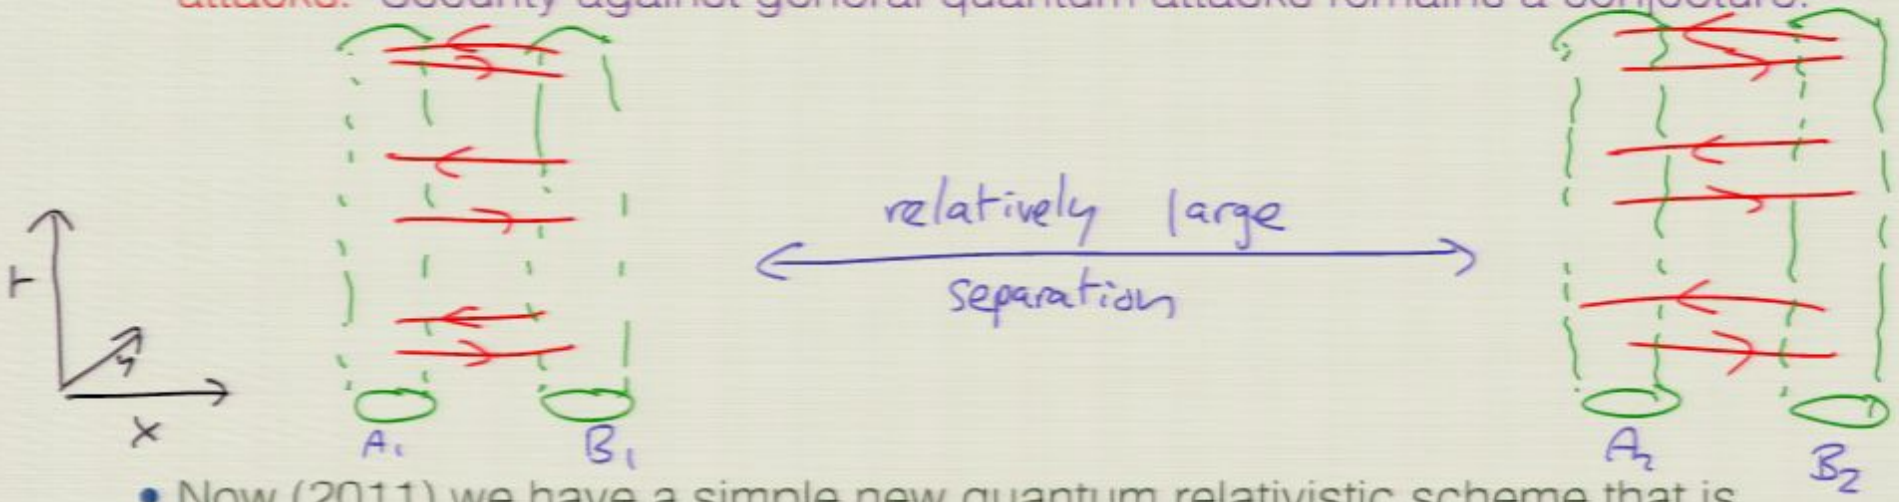

- Now (2011) we have a simple new quantum relativistic scheme that is provably secure, based on the no-summoning theorem.

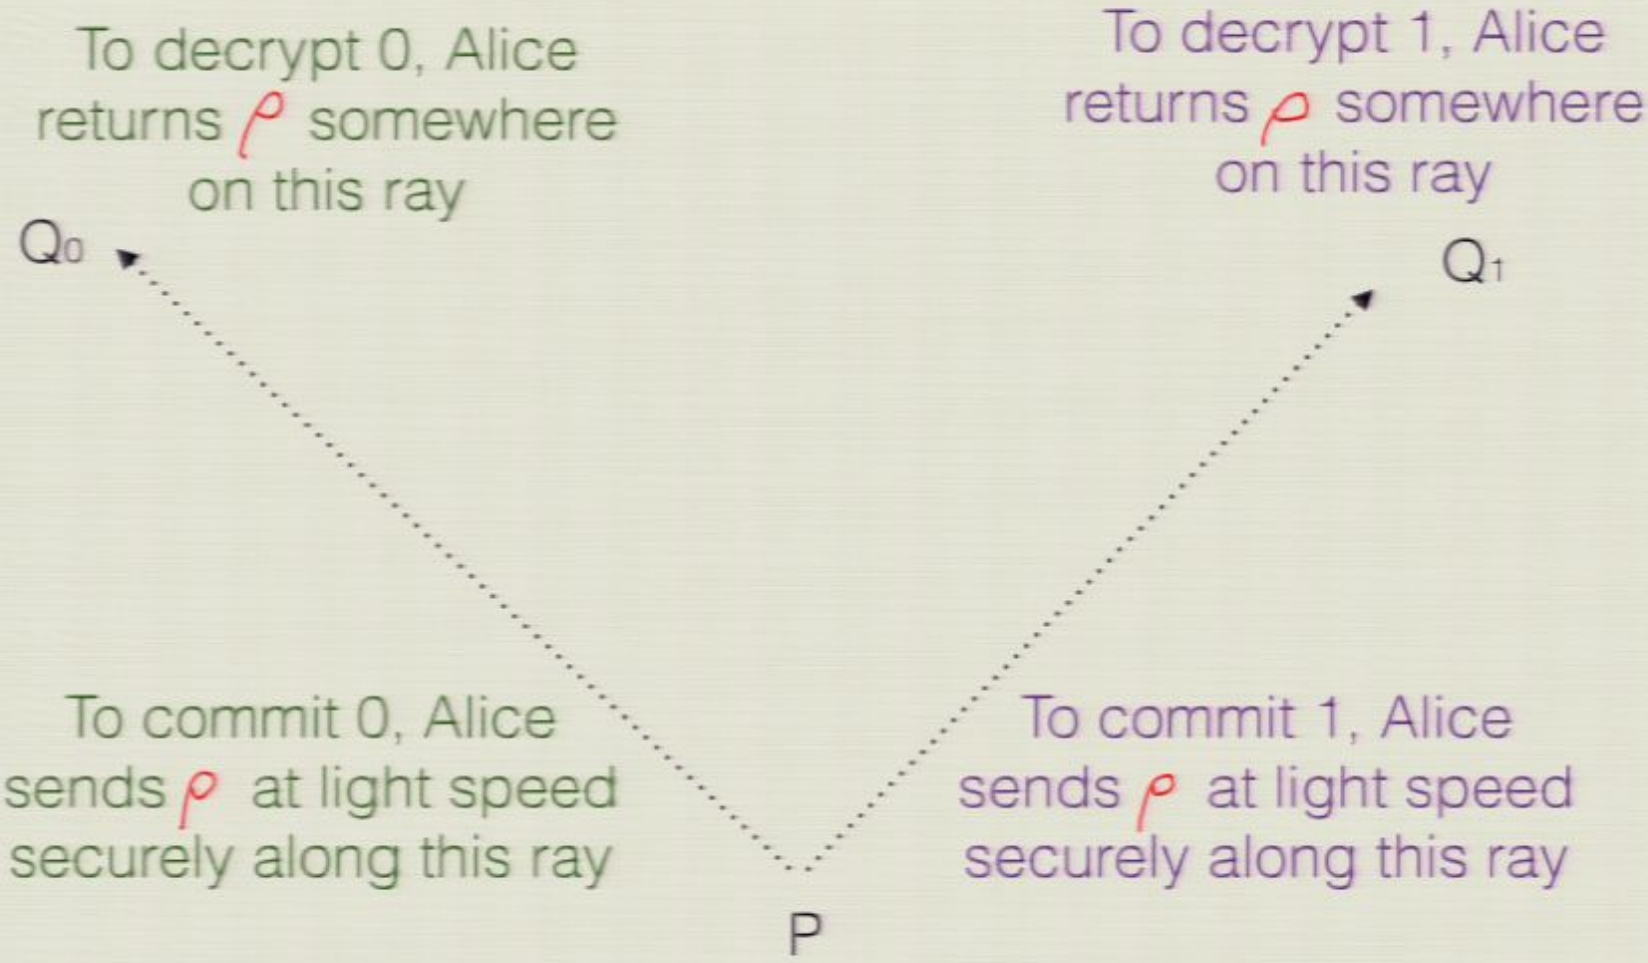

Bob gives Alice state ρ at P

Security against Bob: ensured since Alice sends the state securely (either because she controls a region around the relevant light rays, or e.g. via teleportation)

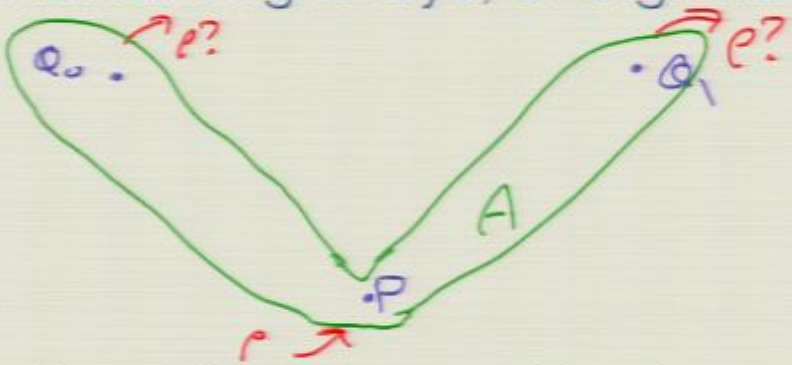

Security against Alice: ensured by the no-summoning theorem -- she cannot return ρ independently at points on both light rays.

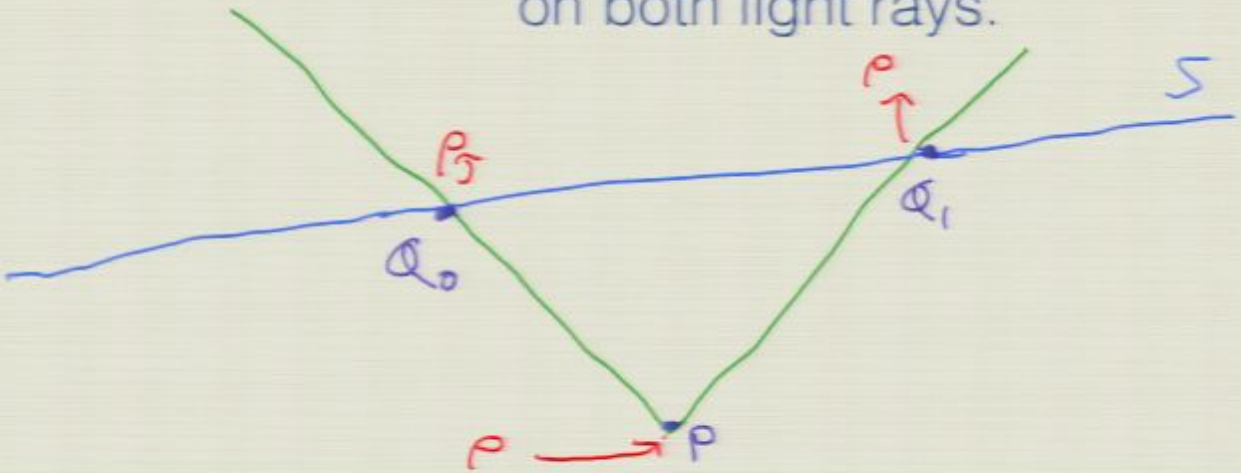

More precisely, we can quantify the security in terms of the dimension d of the space of the unknown state: Alice's cheating probability is bounded by O(1/d).

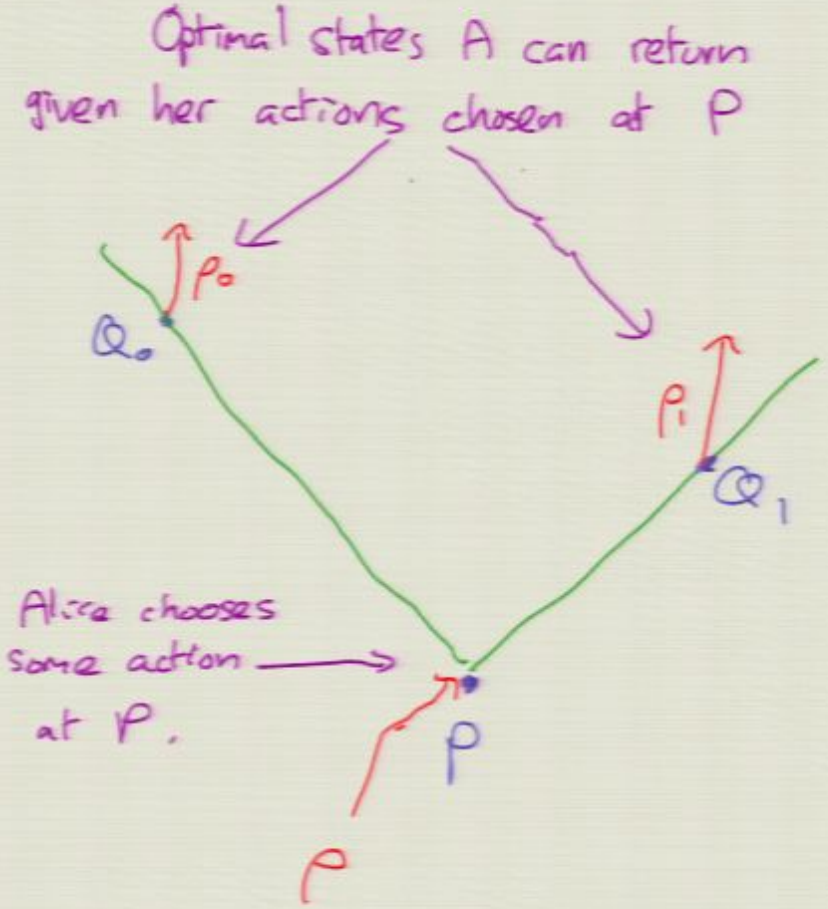

$$P(\text{Bob accepts unveiling at } Q_0) + P(\text{Bob accepts unveiling at } Q_1)$$

$$= Tr(\rho P_0) + Tr(\rho P_1)$$

$$\leq 1 + \frac{2}{d+1}$$

Alice's "wiggle room" decays exponentially in # qubits $= \log_2(d)$.

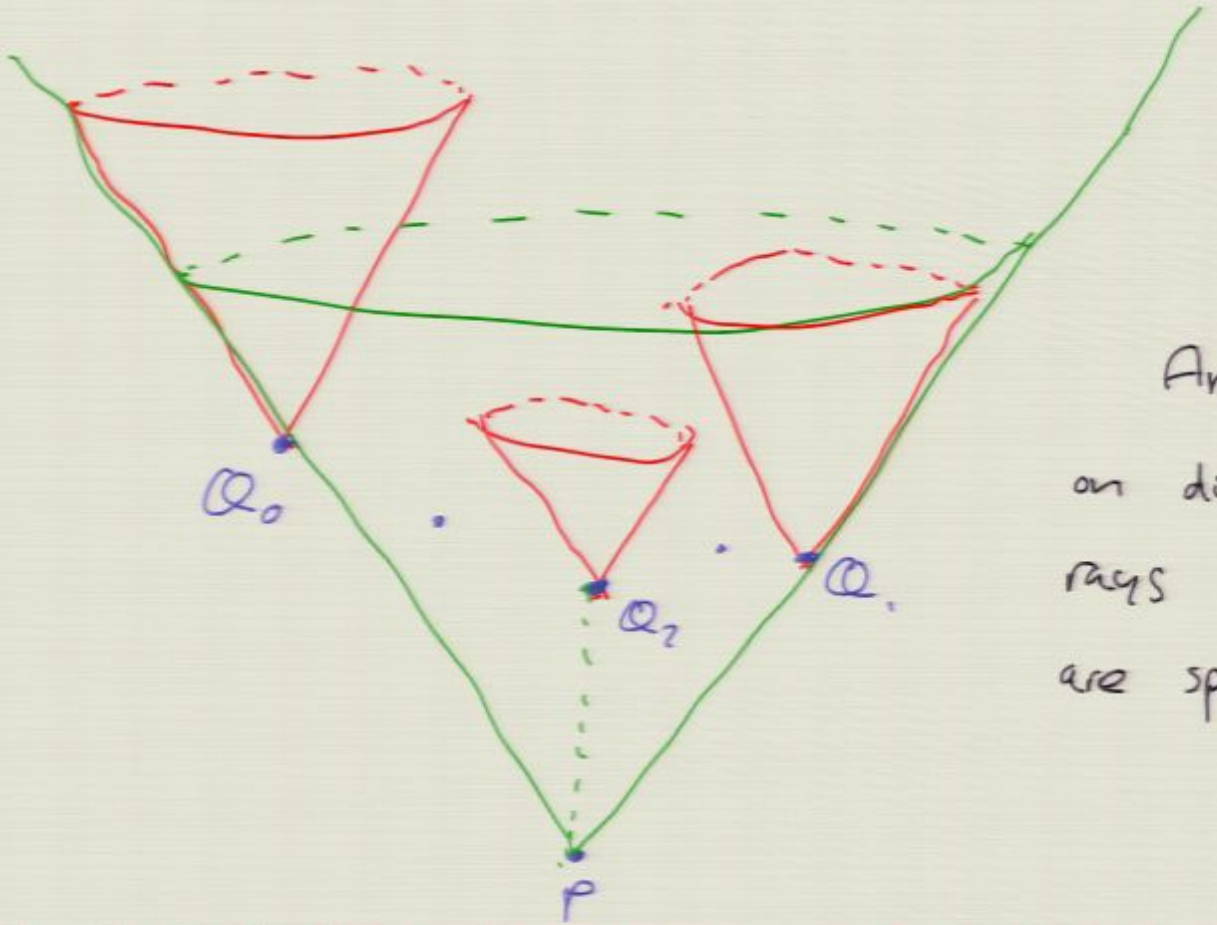

Any 2 points on distinct light rays through P are spacelike separated

This works in 3+1 dimensions also -- and now each possible light like direction can code for a different data value, so the amount of data committed is bounded only by the precision of Alice's transmission and Bob's measurement.

More precisely, we can quantify the security in terms of the dimension d of the space of the unknown state: Alice's cheating probability is bounded by O(1/d).

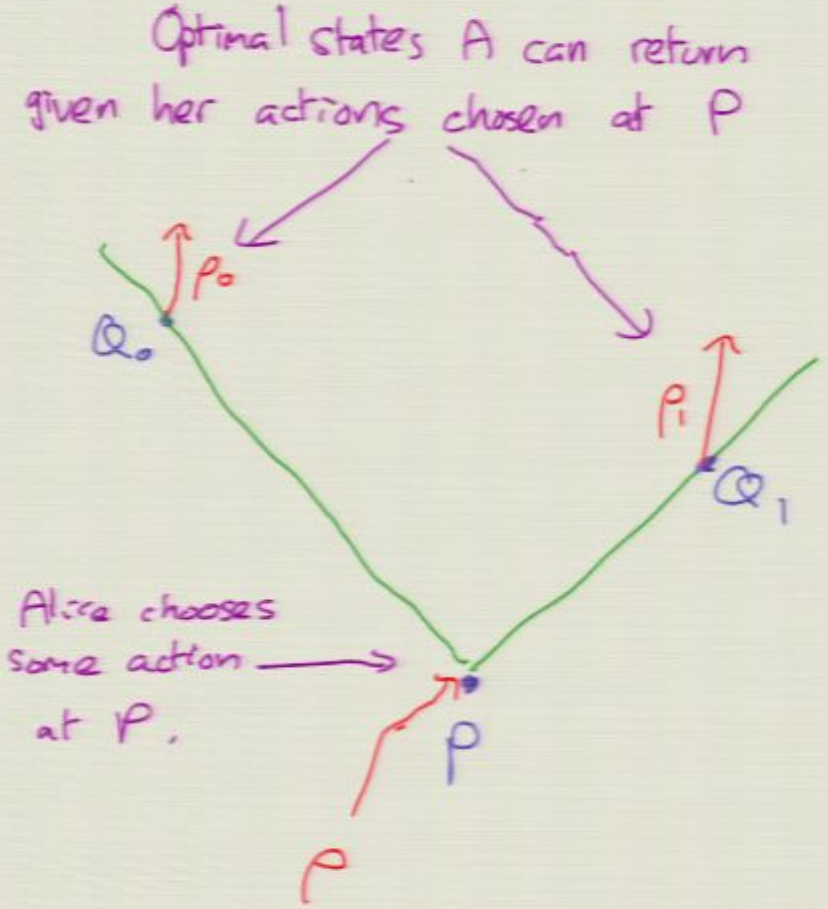

$$P(\text{Bob accepts unveiling at } Q_0) + P(\text{Bob accepts unveiling at } Q_1)$$

$$= Tr(\rho P_0) + Tr(\rho P_1)$$

$$\leq 1 + \frac{2}{d+1}$$

Alice's "wiggle room" decays exponentially in # qubits $= \log_2(d)$.

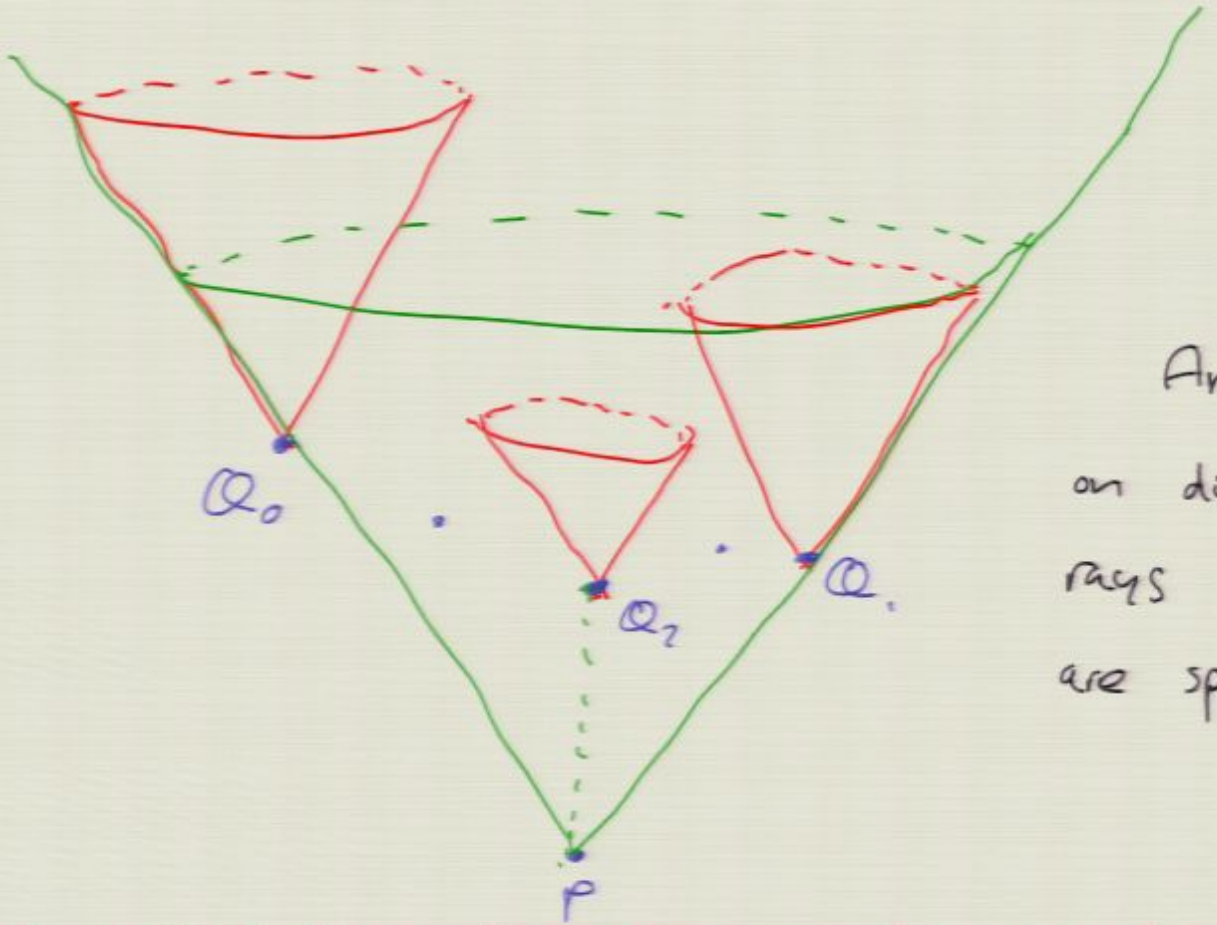

Any 2 points on distinct light rays through P are spacelike separated

This works in 3+1 dimensions also -- and now each possible light like direction can code for a different data value, so the amount of data committed is bounded only by the precision of Alice's transmission and Bob's measurement.

No contradiction with the Mayers-Lo-Chau no-go theorem

Mayers and Lo-Chau's celebrated result shows that unconditionally secure bit commitment is impossible for a large class of quantum protocols -- but the proof makes some tacit assumptions.

In particular, it assumes that, if there is a unitary map taking a 0 commitment to a 1 commitment, known to Alice, she can implement it physically -- and so cheat by altering her commitments.

In our protocol Alice does know the relevant unitary -- which takes a qudit going along one light ray to the same qudit going along another.

But this unitary cannot be implemented physically, as it would violate causality. So the Mayers-Lo-Chau cheating strategy doesn't apply.

More relativistic quantum (im)possibilities: Tagging an inaccessible object

Alice has constructed a tagging device, which Eve can't break into or detach, and wants to authenticate its position by exchanging quantum signals from remote sites.

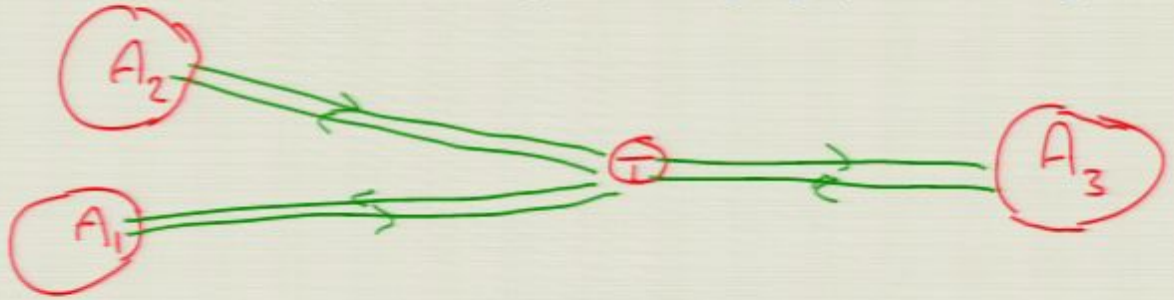

Eve wants to spoof the tagging device by intercepting Alice's signals elsewhere and responding to them -- she can then spirit away the tagged object without Alice's knowledge.

A Brief History of Quantum Tagging

- Independently invented by KMSB (2002, patent 2006), CFGGO (2010) (who used the name quantum position-verification, and extended to more general position-based quantum cryptography), Malaney (2009).

- Various tagging schemes proposed: CFGGO and Malaney schemes claimed proven secure, but broken by teleportation attacks (KMS 2010). New schemes proposed by KMS 2010 (security left open) and LL 2010 (security conjectured).

- (Im)possibility of security turns out to depend crucially on subtleties in the properties assumed for the tag: in particular, whether Eve can read information from within it. Secure quantum tagging is possible if the tag can keep secret data shared with Alice (K 2010).

- For tags that cannot hold secrets, a large class of tagging schemes including KMS 2010 and LL 2010 are provably insecure (BCFGGOS, 2010) -- a beautiful result that relies on earlier work by Vaidman (2003) on non-local quantum measurements.

Quantum Tagging References

- KMSB (2006): AK, Munro, Spiller, Beausoleil, "Tagging Systems", US patent 2006/0022832
- Malaney (2009): Phys Rev A 81 042319 and arxiv 1004.2689
- CFGGO (2010): Chandran, Fehr, Gelles, Goyal, Ostrovsky arxiv: 1005.1750
- KMS (2010): AK, Munro, Spiller, arxiv:1008.2147
- K (2010): AK, arxiv:1008.5380
- LL (2010): Lau, Lo: arxiv:1009.2256
- BCFGGOS (2010): Buhrman,CFGGO,Schaffner, arxiv: 1009.2490
- (Relies on Vaidman (2003): Phys. Rev. Lett 90 010402.)

Quantum Tagging in our general framework

A tagging protocol requires the tag to produce outputs reaching the A_i at specified space-time points, in response to inputs.

This must be possible given control of the region occupied by T. To be secure, it should be impossible given control of the complement of T.

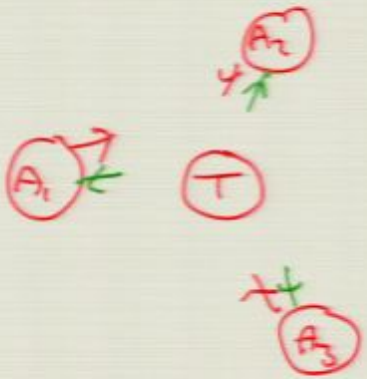

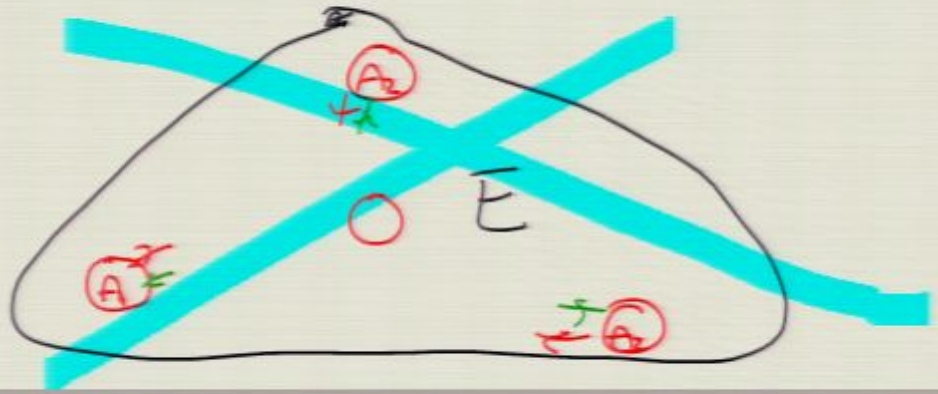

The CFGGO tagging scheme and how to break it

States from this list are sent from left, a bit telling T to measure in $\{|0\rangle, |1\rangle\}$ or $\{|+\rangle, |-\rangle\}$ basis is sent from right. They arrive simultaneously : T broadcasts the outcome.

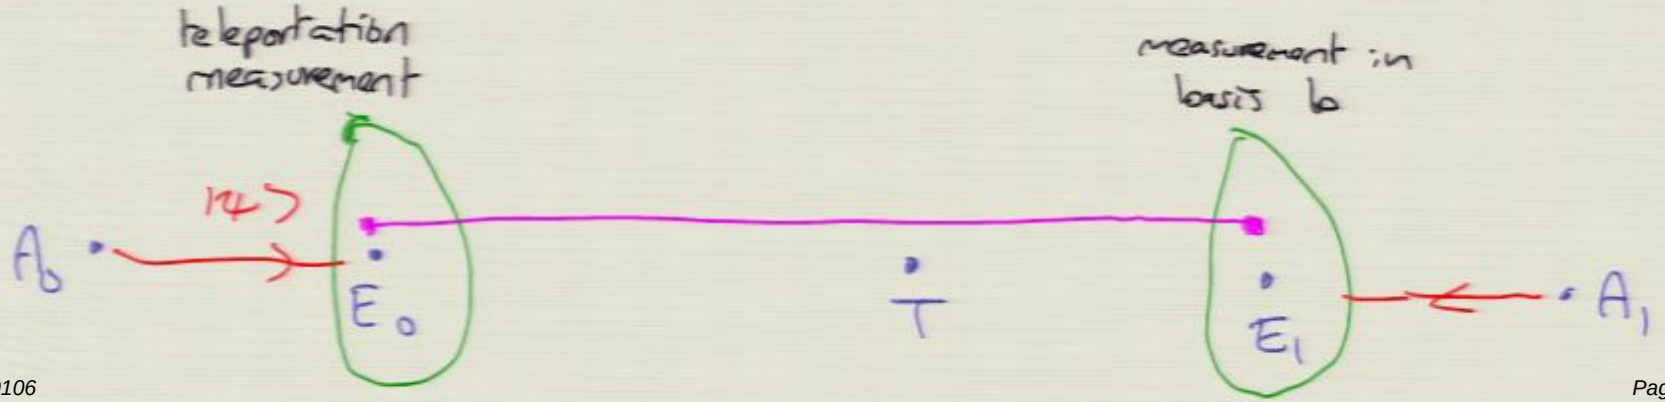

The CFGGO tagging scheme and how to break it

(1D Case)

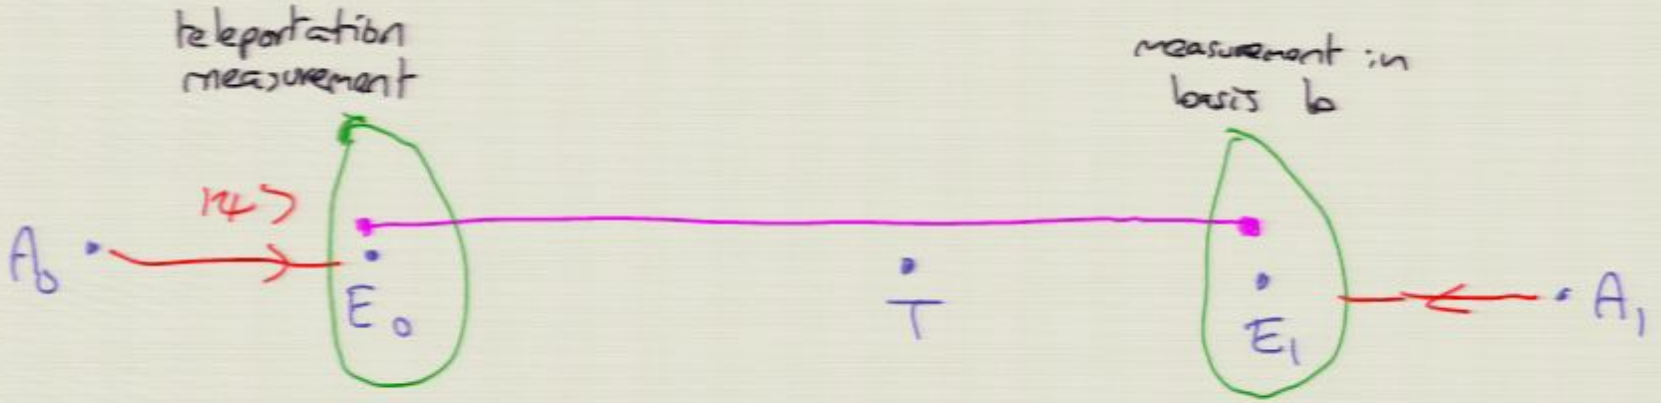

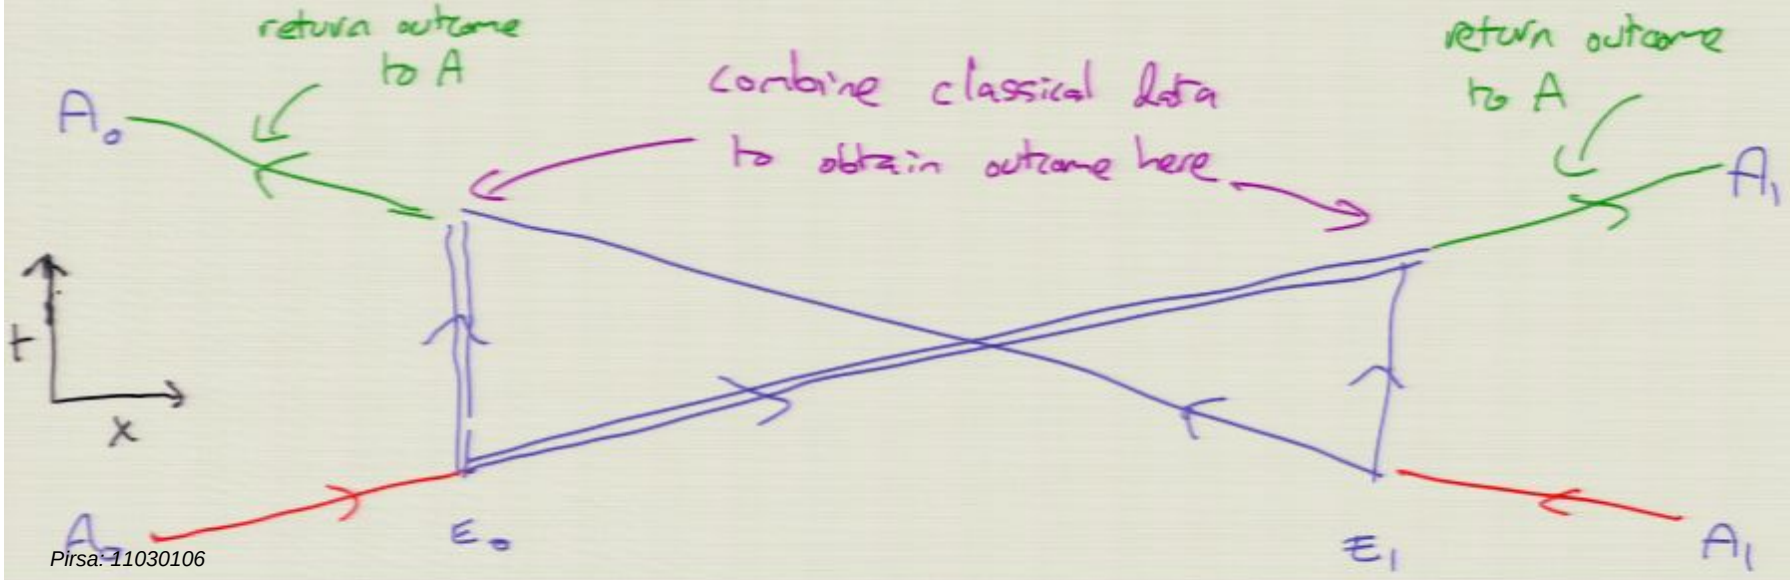

Generalising teleportation attacks (BCFGGOS 2010)

A large class of schemes effectively reduce to requiring T to apply a unitary U on inputs and return result as outputs

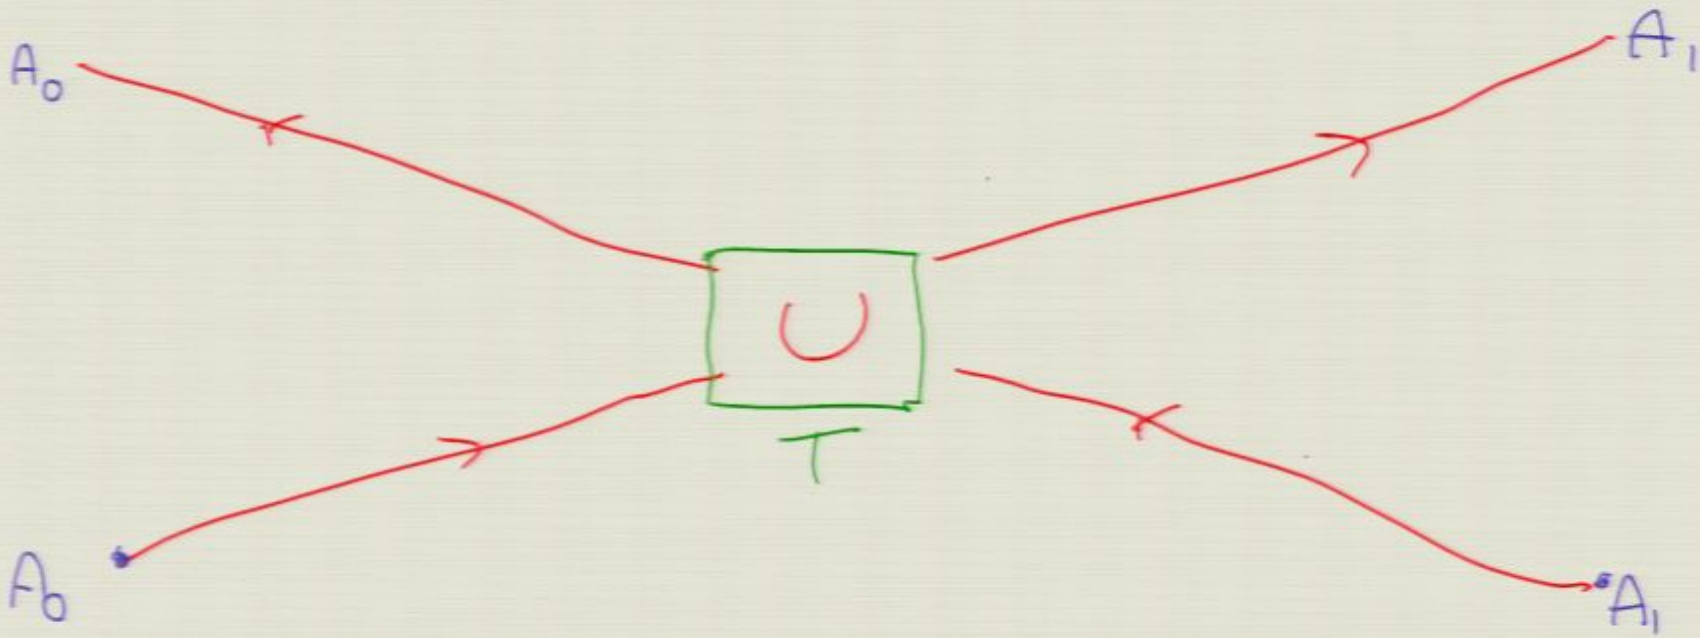

Buhrman et al. (BCFGGOS) show that this:

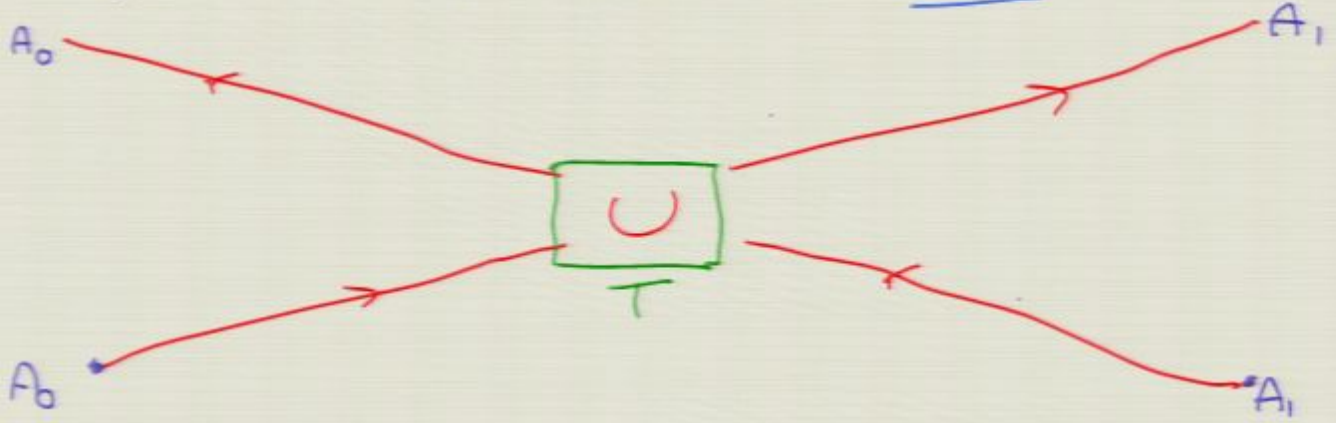

Can be spoofed by this:

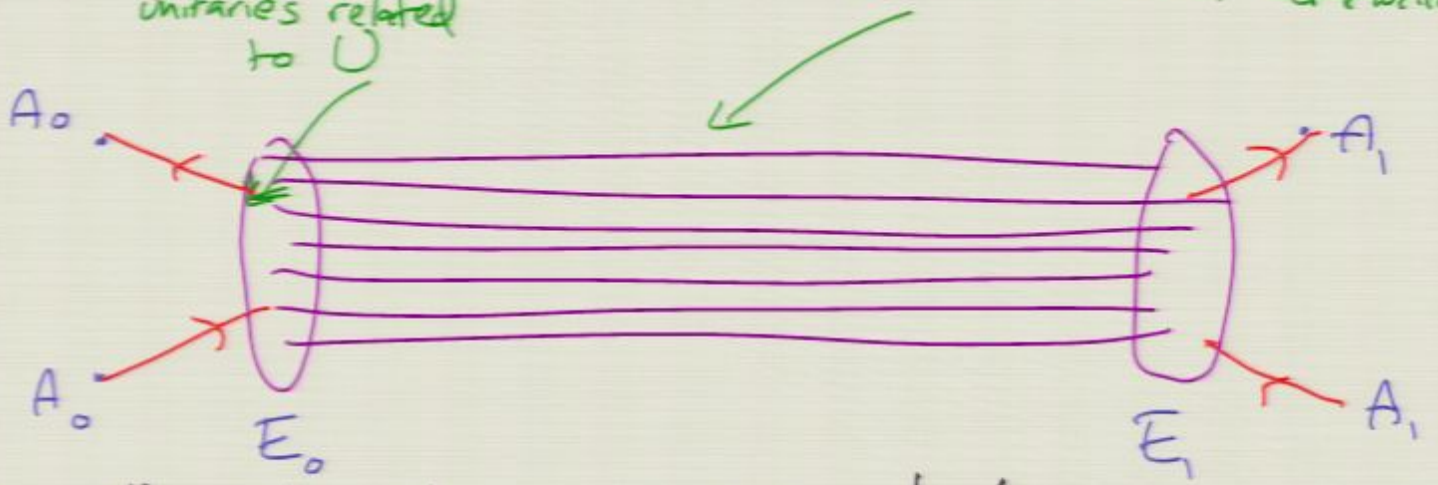

making all such schemes insecore! (generalizes to 3D also)

Generalising teleportation attacks (BCFGGOS 2010)

A large class of schemes effectively reduce to requiring T to apply a unitary U on inputs and return result as outputs

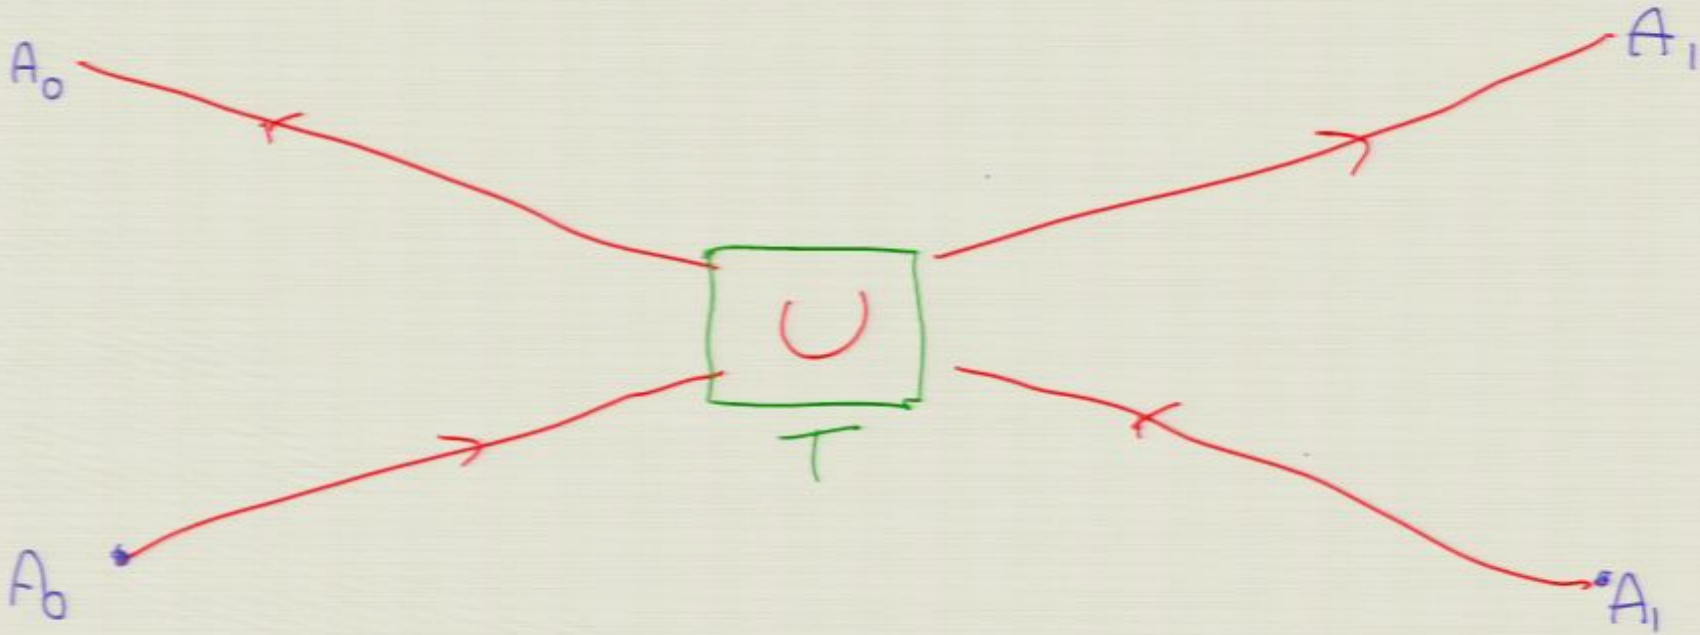

Buhrman et al. (BCFGGOS) show that this:

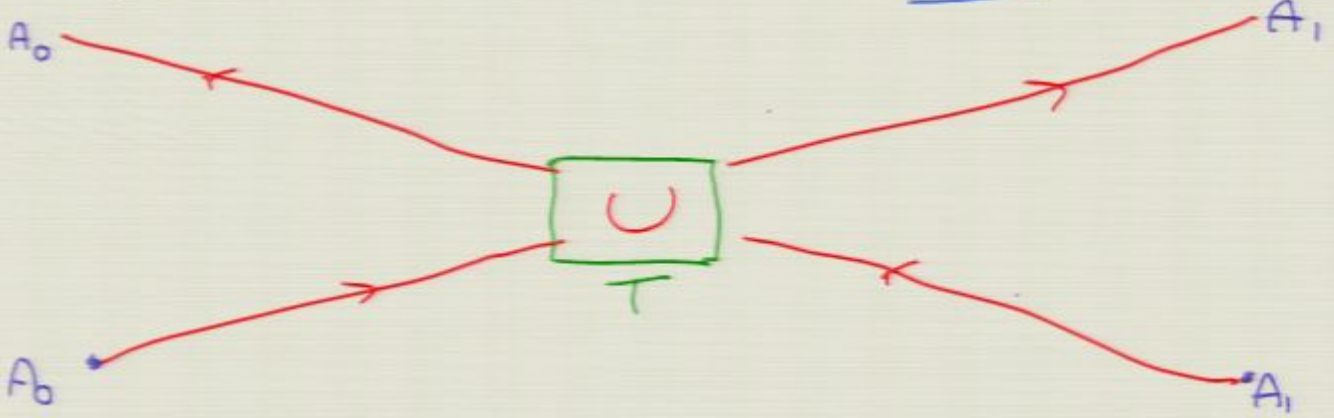

Can be spoofed by this:

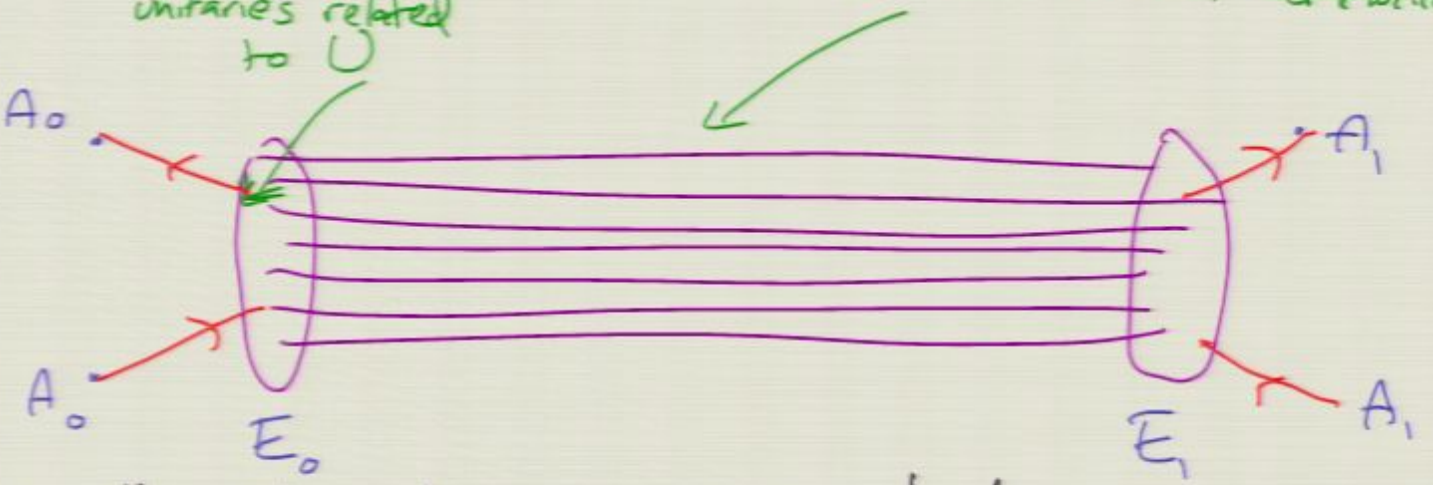

making all such schemes insecore! (generalizes to 3D also)

Secure quantum tagging using tags containing secret data (K 2010)

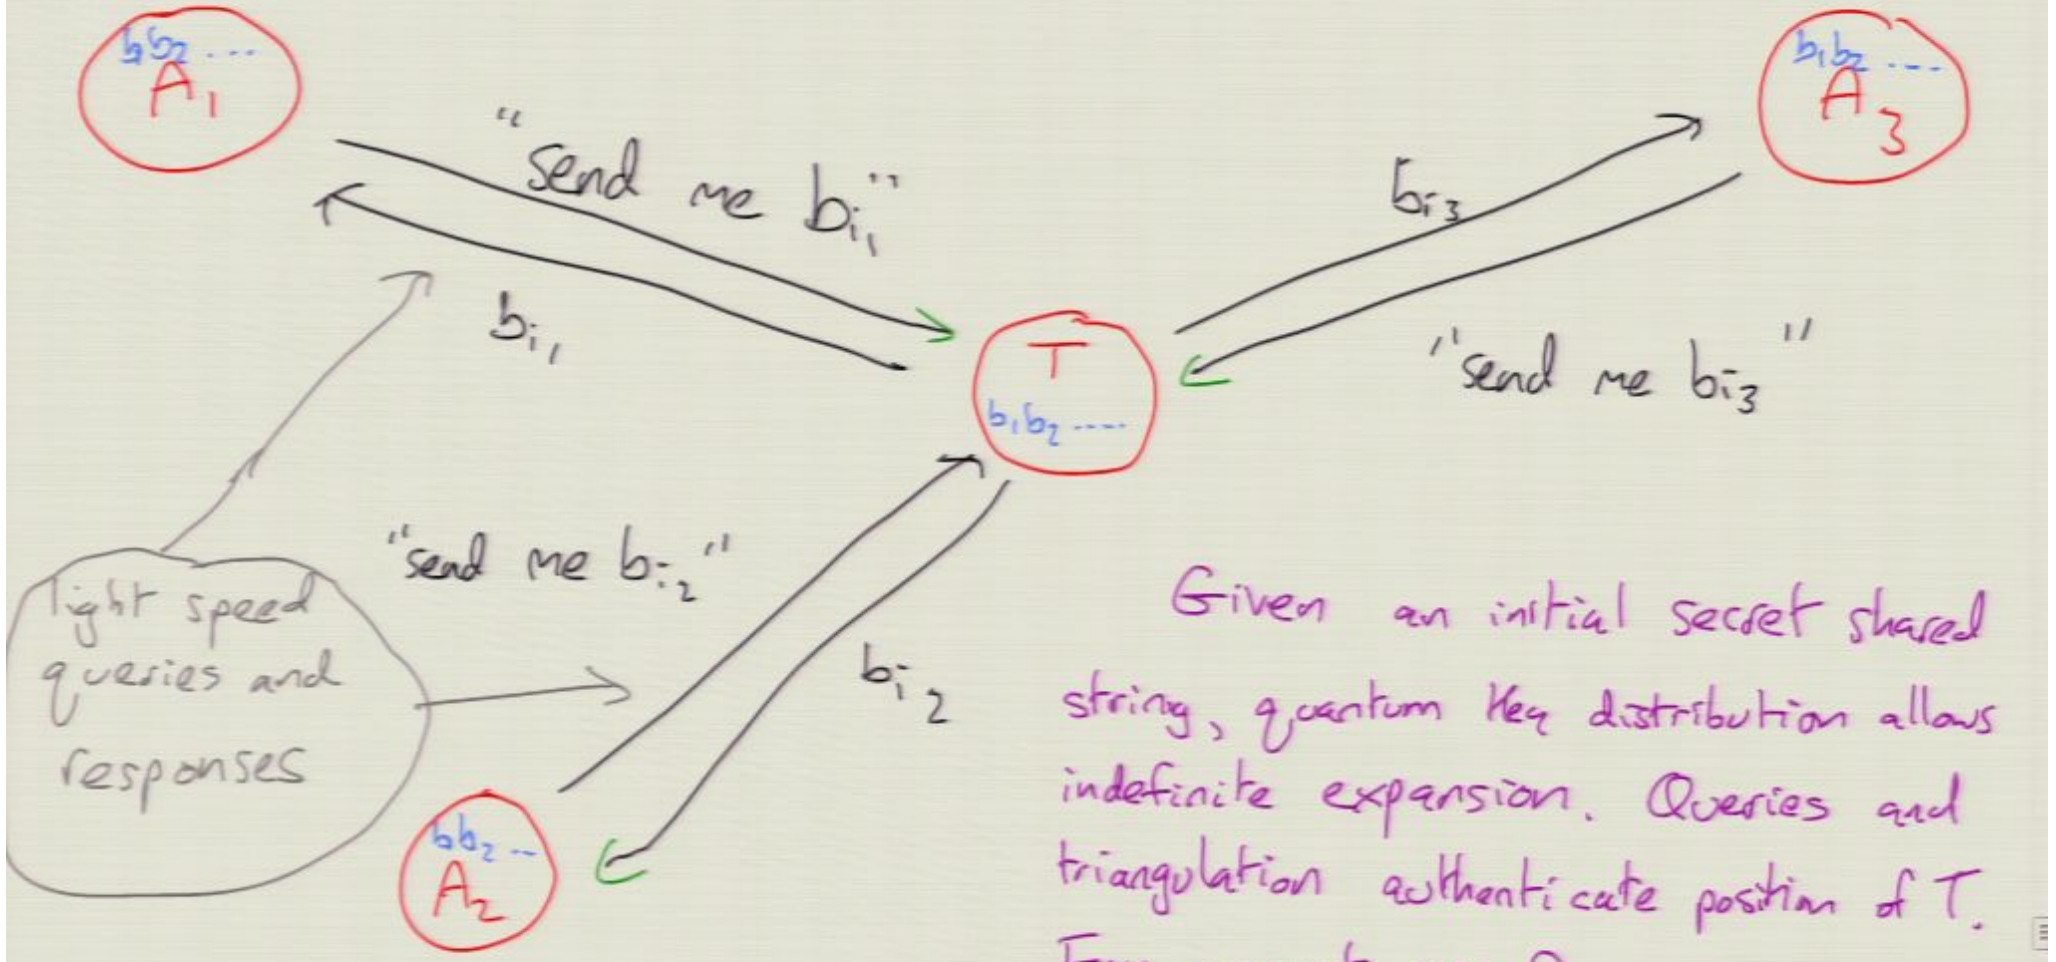

Given an initial secret shared string, quantum key distribution allows indefinite expansion. Queries and triangulation authenticate position of T. Eve cannot spoof.

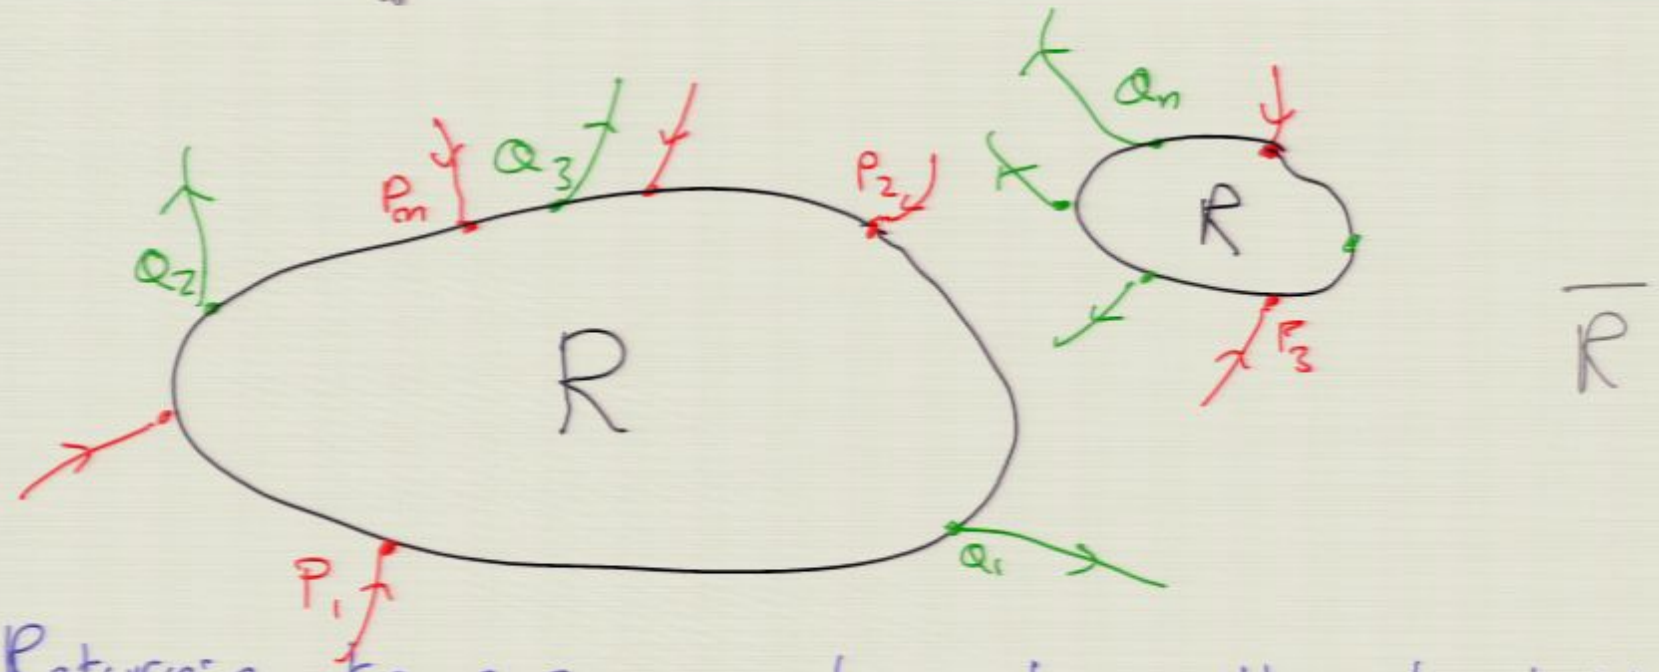

Returning to our general question, this highlights an interesting subtlety: we might wish to allow Eve control of R and nothing more, or control of R and the ability to communicate through $\bar{R}$, or also to read classical or quantum data in $\bar{R}$. All of these are interesting questions, each potentially practically relevant in the right context.

Summary

- We can learn new features of relativistic quantum theory by considering intrinsically relativistic and quantum tasks.
- Summoning is a simple example, which singles out relativistic quantum theory from NRQM or relativistic classical mechanics.
- It's cryptographically powerful, with a direct application to quantum bit commitment; it also allows other relativistic cryptographic tasks to be implemented securely.
- Quantum tagging and position-based quantum cryptography are further natural applications, with intriguing (and practically relevant) possibilities and impossibilities.
- There are surely many other interesting tasks, many open questions, and many new quantum cryptographic and computational applications.

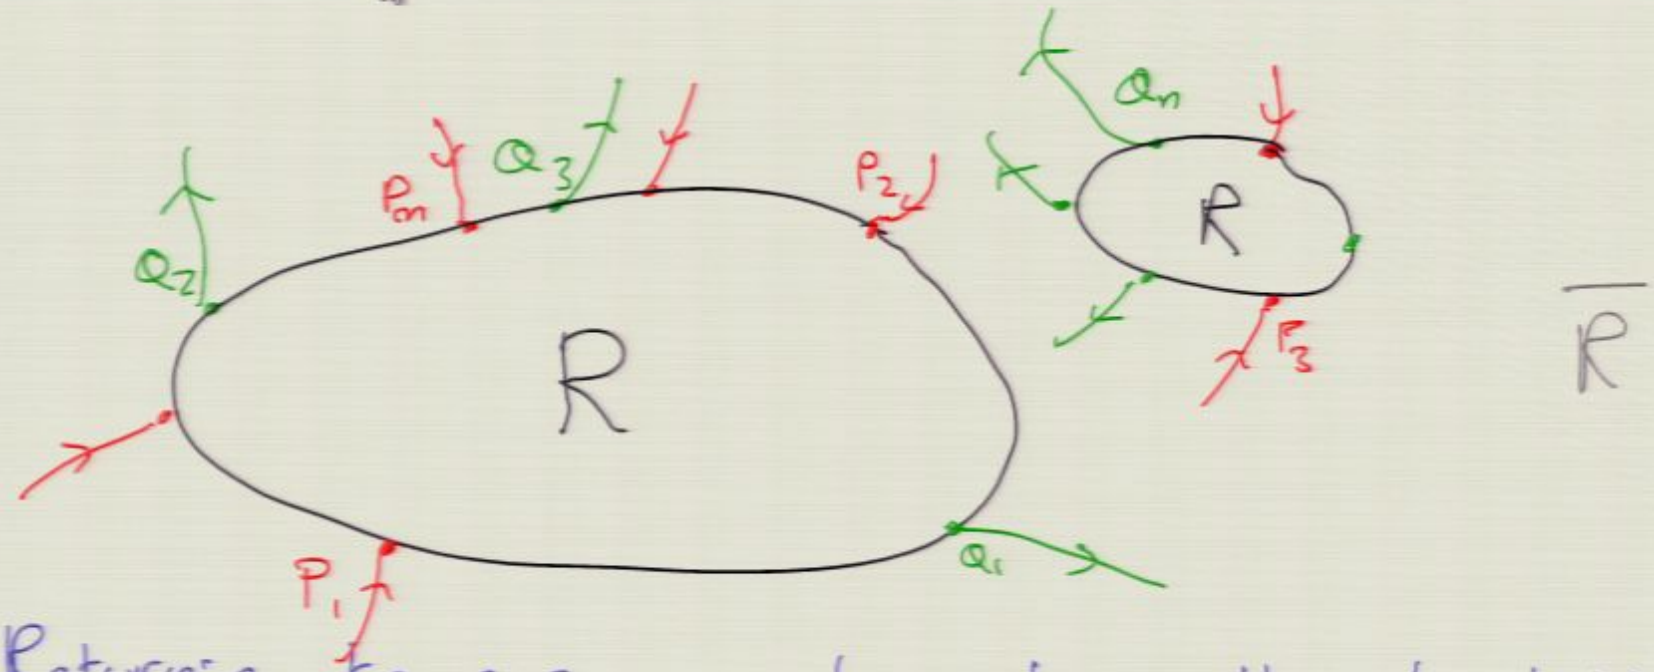

Returning to our general question, this highlights an interesting subtlety: we might wish to allow Eve control of R and nothing more, or control of R and the ability to communicate through $\bar{R}$, or also to read classical or quantum data in $\bar{R}$. All of these are interesting questions, each potentially practically relevant in the right context.

Buhrman et al. (BCFGGOS) show that this:

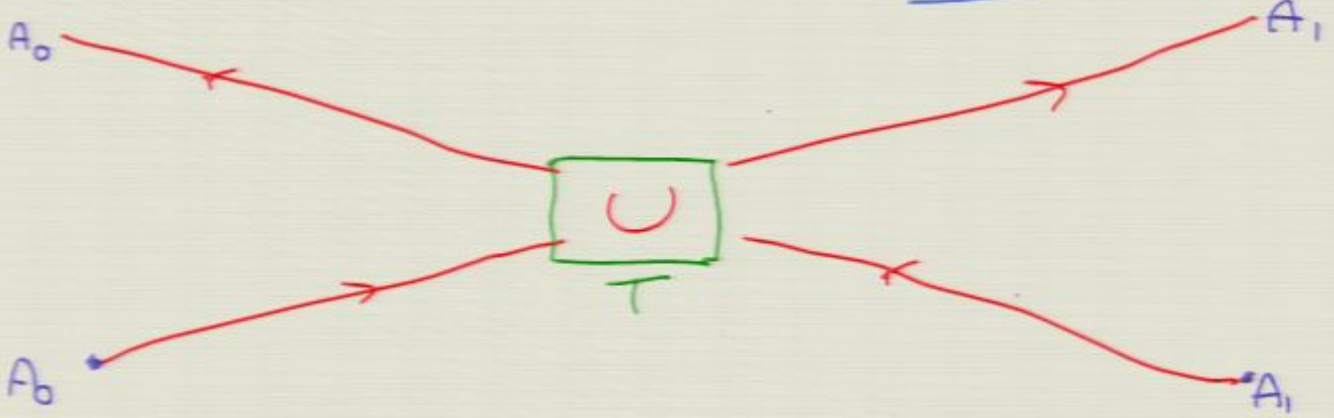

Can be spoofed by this:

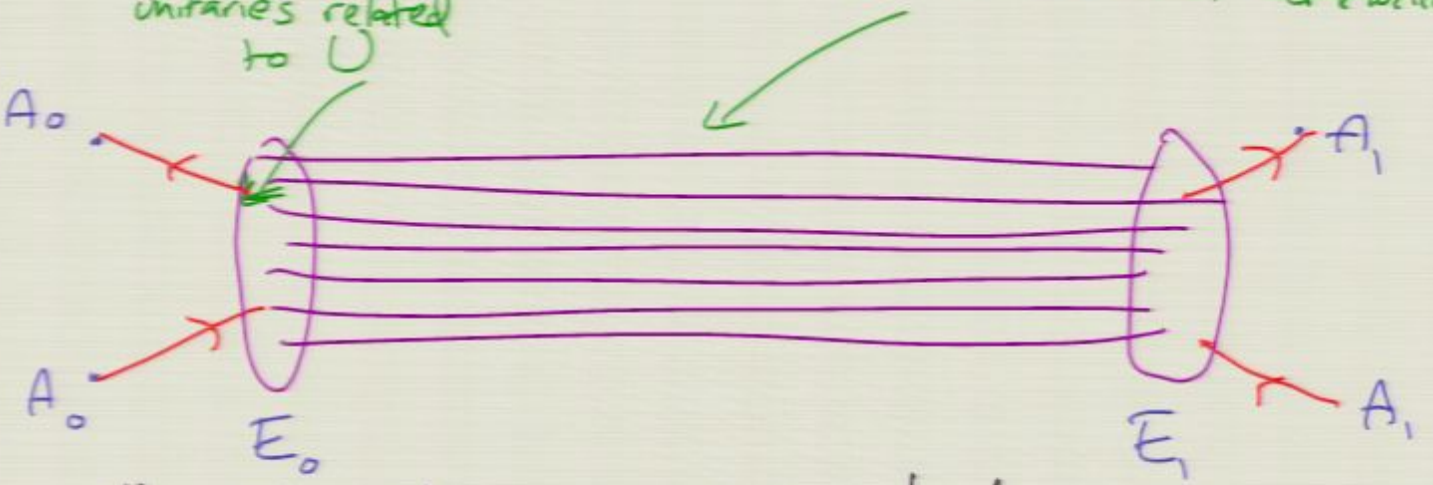

making all such schemes insecure! (generalizes to 3D also)

Quantum Tagging in our general framework

A tagging protocol requires the tag to produce outputs reaching the A_i at specified space-time points, in response to inputs.

This must be possible given control of the region occupied by T. To be secure, it should be impossible given control of the complement of T.

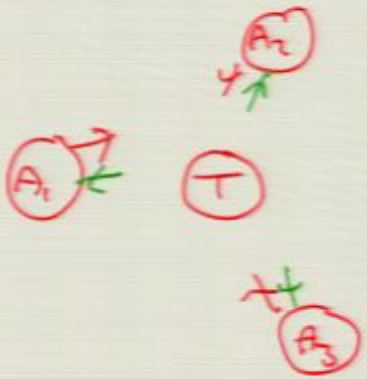

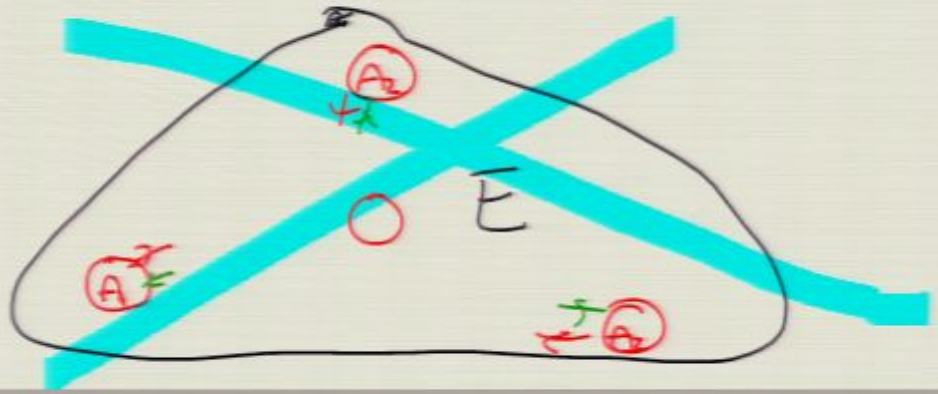

Quantum Tagging References

- KMSB (2006): AK, Munro, Spiller, Beausoleil, "Tagging Systems", US patent 2006/0022832
- Malaney (2009): Phys Rev A 81 042319 and arxiv 1004.2689
- CFGGO (2010): Chandran, Fehr, Gelles, Goyal, Ostrovsky arxiv: 1005.1750
- KMS (2010): AK, Munro, Spiller, arxiv:1008.2147
- K (2010): AK, arxiv:1008.5380
- LL (2010): Lau, Lo: arxiv:1009.2256
- BCFGGOS (2010): Buhrman,CFGGO,Schaffner, arxiv: 1009.2490
- (Relies on Vaidman (2003): Phys. Rev. Lett 90 010402.)

A Brief History of Quantum Tagging

- Independently invented by KMSB (2002, patent 2006), CFGGO (2010) (who used the name quantum position-verification, and extended to more general position-based quantum cryptography), Malaney (2009).

- Various tagging schemes proposed: CFGGO and Malaney schemes claimed proven secure, but broken by teleportation attacks (KMS 2010). New schemes proposed by KMS 2010 (security left open) and LL 2010 (security conjectured).

- (Im)possibility of security turns out to depend crucially on subtleties in the properties assumed for the tag: in particular, whether Eve can read information from within it. Secure quantum tagging is possible if the tag can keep secret data shared with Alice (K 2010).

- For tags that cannot hold secrets, a large class of tagging schemes including KMS 2010 and LL 2010 are provably insecure (BCFGGOS, 2010) -- a beautiful result that relies on earlier work by Vaidman (2003) on non-local quantum measurements.

No contradiction with the Mayers-Lo-Chau no-go theorem

Mayers and Lo-Chau's celebrated result shows that unconditionally secure bit commitment is impossible for a large class of quantum protocols -- but the proof makes some tacit assumptions.

In particular, it assumes that, if there is a unitary map taking a 0 commitment to a 1 commitment, known to Alice, she can implement it physically -- and so cheat by altering her commitments.

In our protocol Alice does know the relevant unitary -- which takes a qudit going along one light ray to the same qudit going along another.

But this unitary cannot be implemented physically, as it would violate causality. So the Mayers-Lo-Chau cheating strategy doesn't apply.

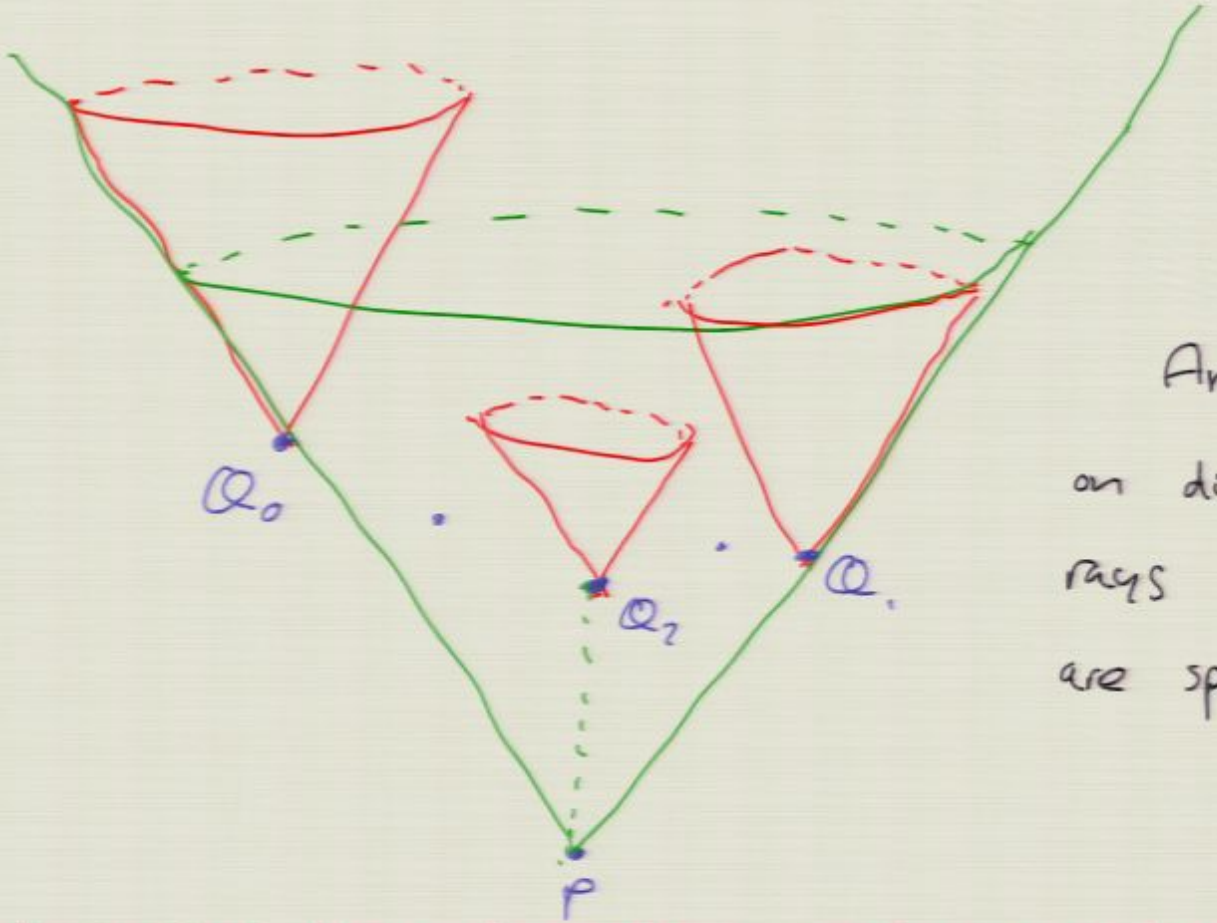

Any 2 points on distinct light rays through P are spacelike separated

This works in 3+1 dimensions also -- and now each possible light like direction can code for a different data value, so the amount of data committed is bounded only by the precision of Alice's transmission and Bob's measurement.

More precisely, we can quantify the security in terms of the dimension d of the space of the unknown state: Alice's cheating probability is bounded by O(1/d).

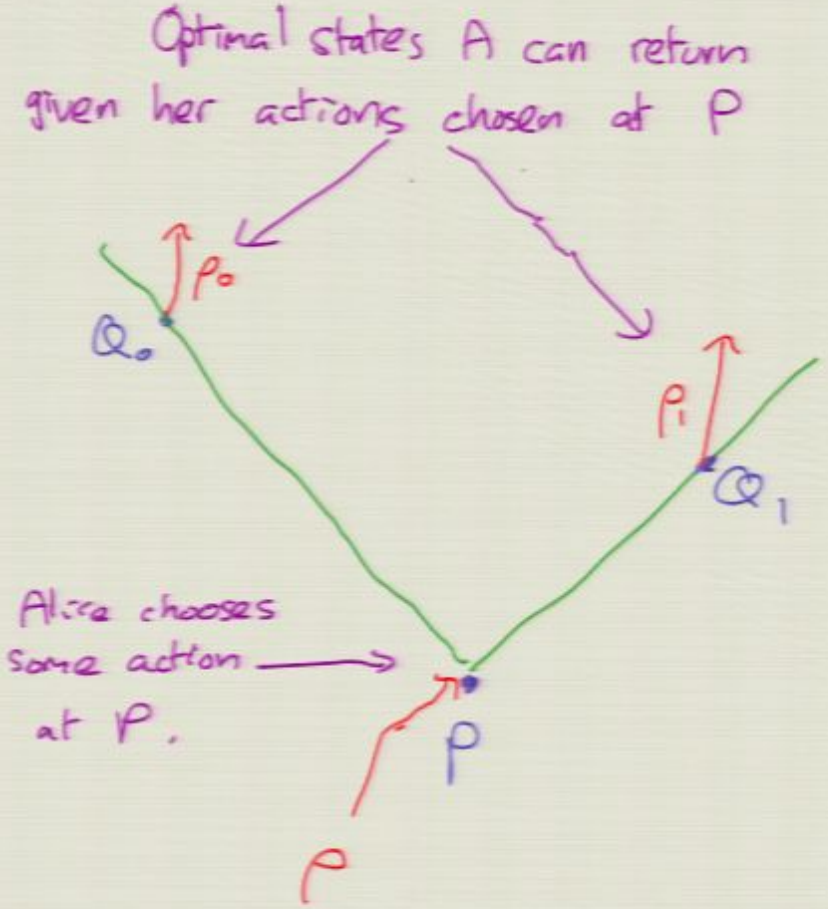

$$P(\text{Bob accepts unveiling at } Q_0) + P(\text{Bob accepts unveiling at } Q_1)$$

$$= Tr(\rho P_0) + Tr(\rho P_1)$$

$$\leq 1 + \frac{2}{d+1}$$

Alice's "wiggle room" decays exponentially in # qubits $= \log_2(d)$.

Security against Bob: ensured since Alice sends the state securely (either because she controls a region around the relevant light rays, or e.g. via teleportation)

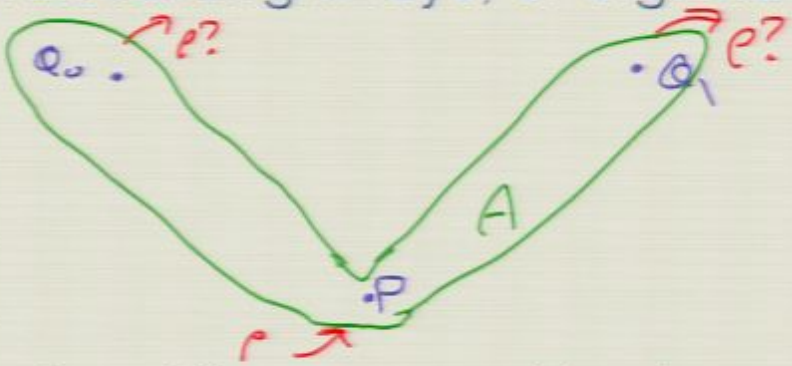

Security against Alice: ensured by the no-summoning theorem -- she cannot return ρ independently at points on both light rays.

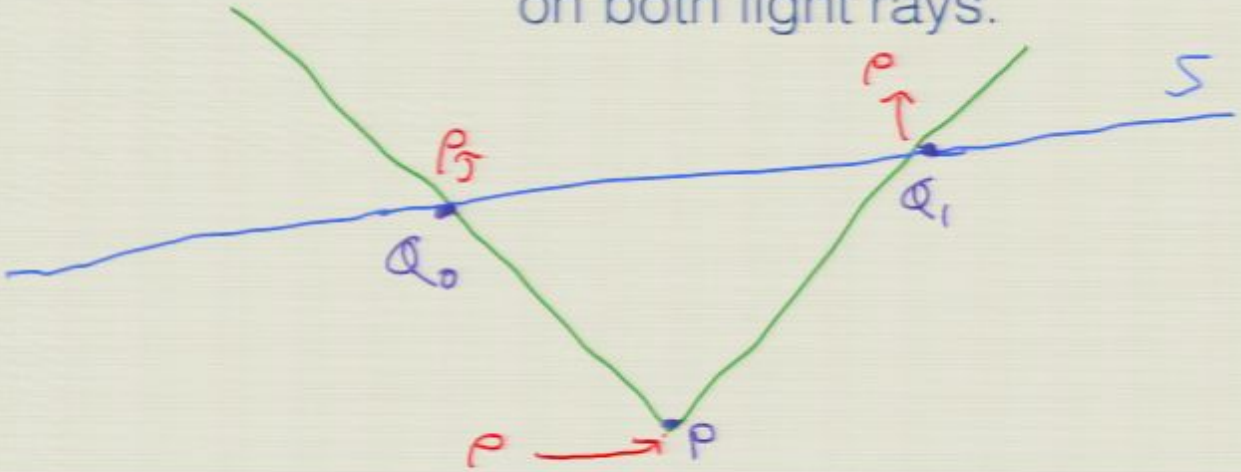

More precisely, we can quantify the security in terms of the dimension d of the space of the unknown state: Alice's cheating probability is bounded by O(1/d).

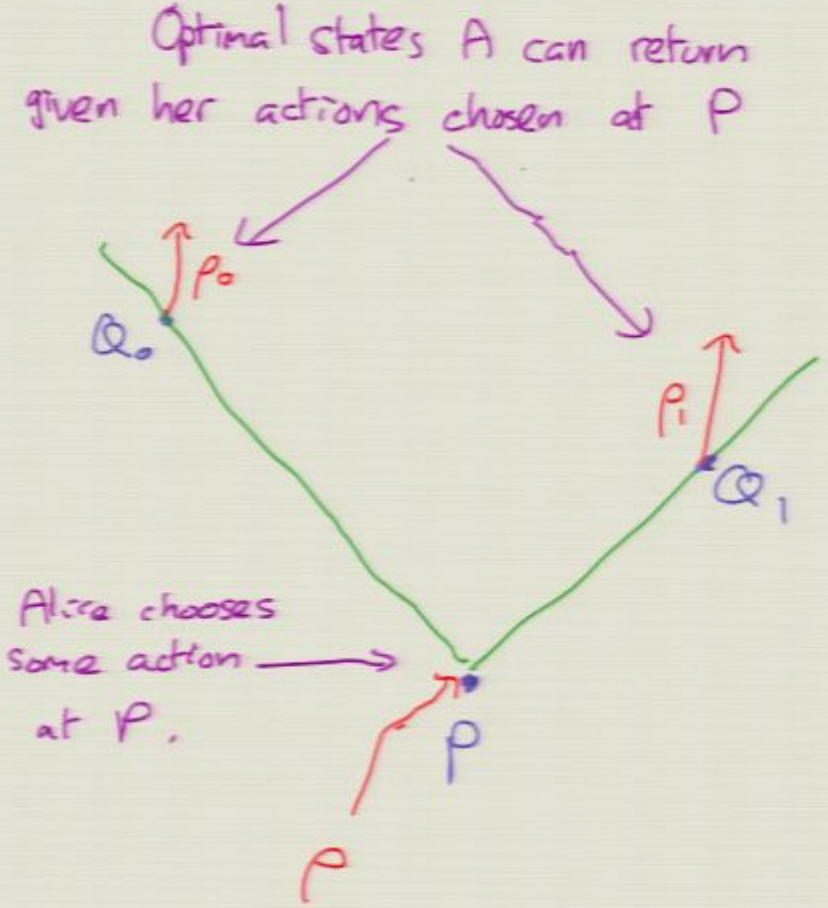

$$P(\text{Bob accepts unveiling at } Q_0) + P(\text{Bob accepts unveiling at } Q_1)$$

$$= Tr(\rho P_0) + Tr(\rho P_1)$$

$$\leq 1 + \frac{2}{d+1}$$

Alice's "wiggle room" decays exponentially in # qubits $= \log_2(d)$.

Security against Bob: ensured since Alice sends the state securely (either because she controls a region around the relevant light rays, or e.g. via teleportation)

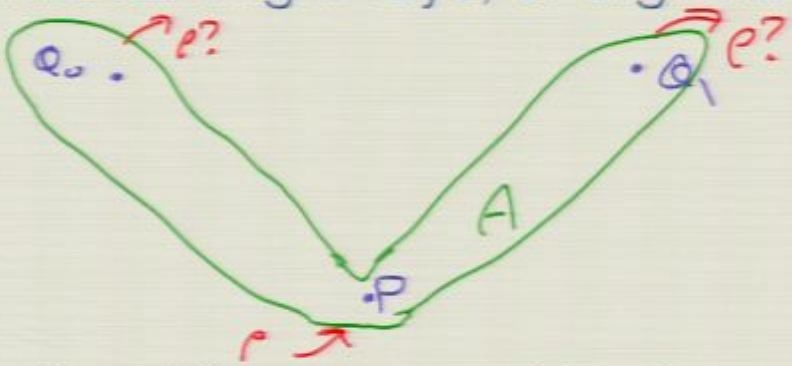

Security against Alice: ensured by the no-summoning theorem -- she cannot return ρ independently at points on both light rays.

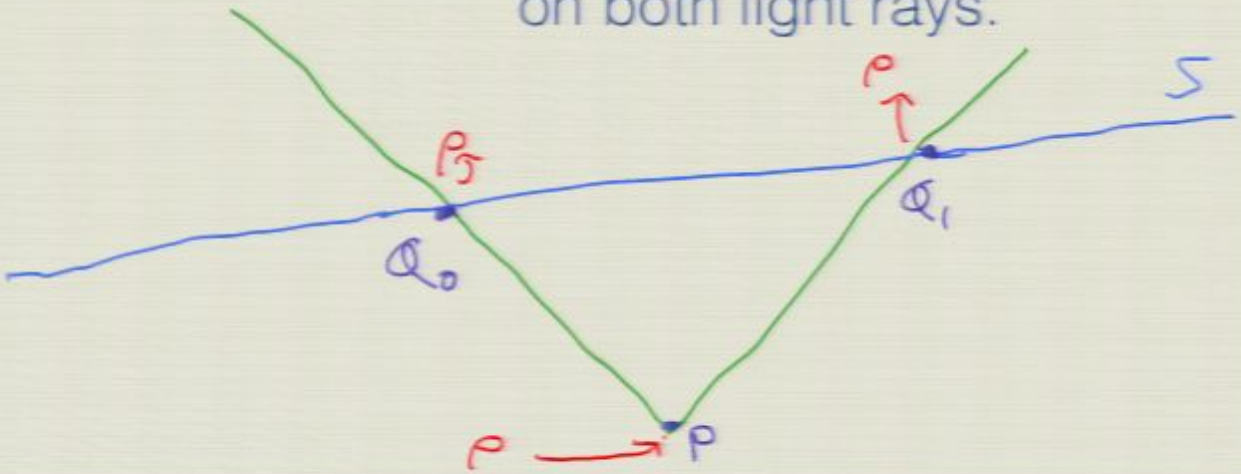

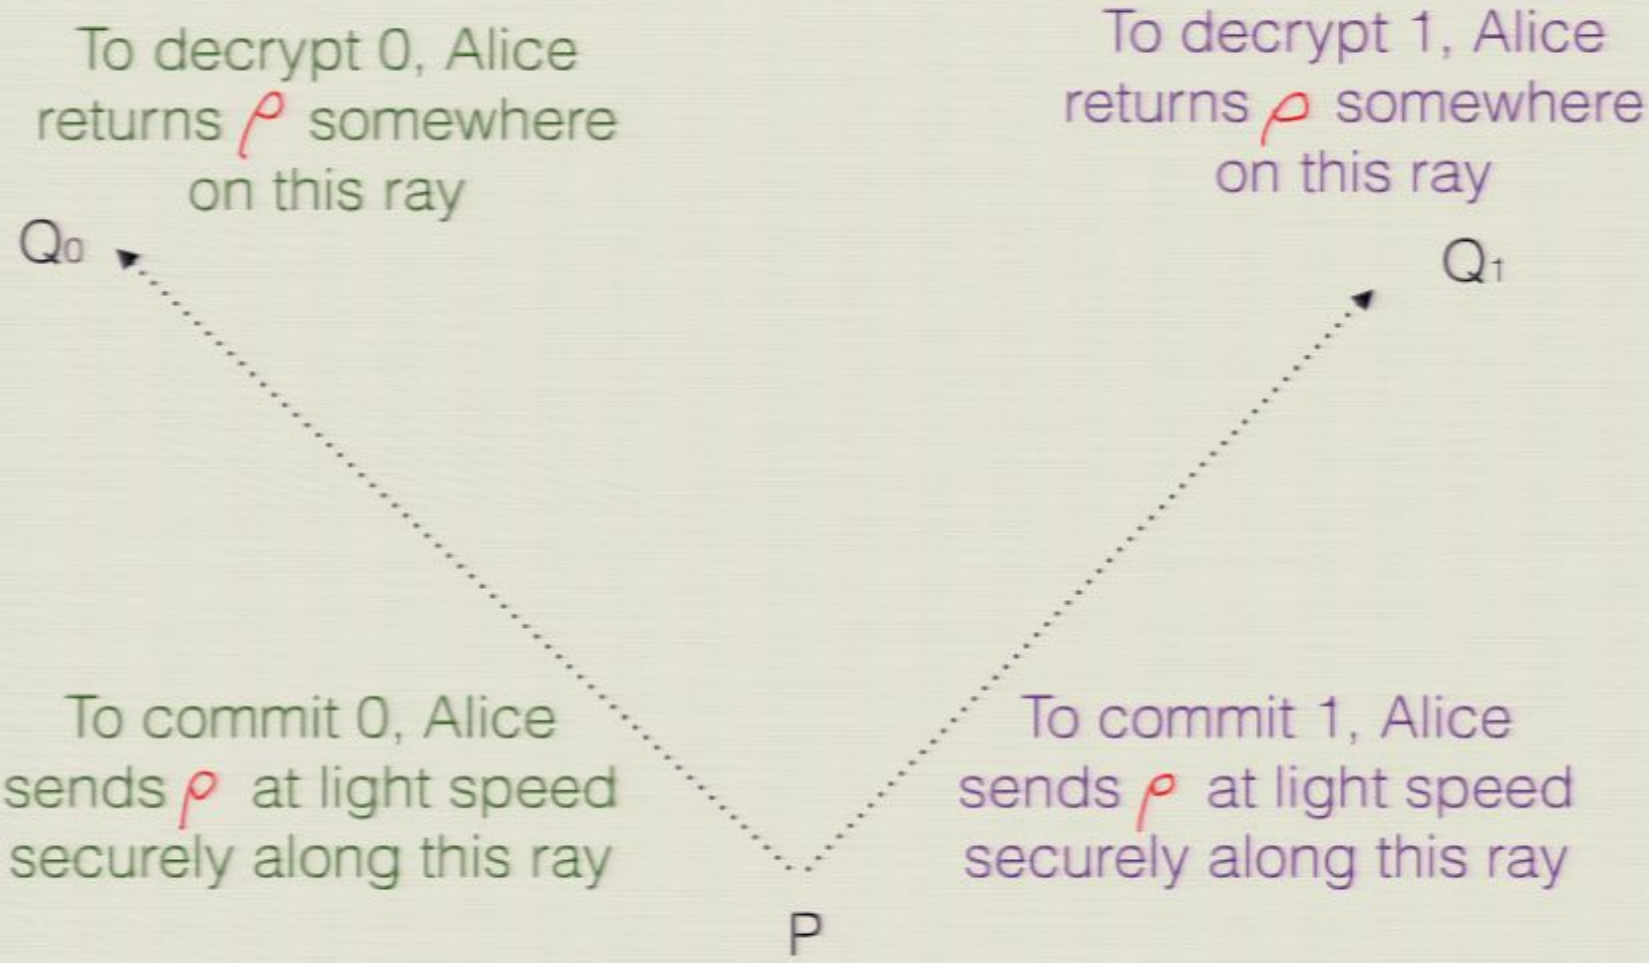

Bob gives Alice state ρ at P

Security against Bob: ensured since Alice sends the state securely (either because she controls a region around the relevant light rays, or e.g. via teleportation)

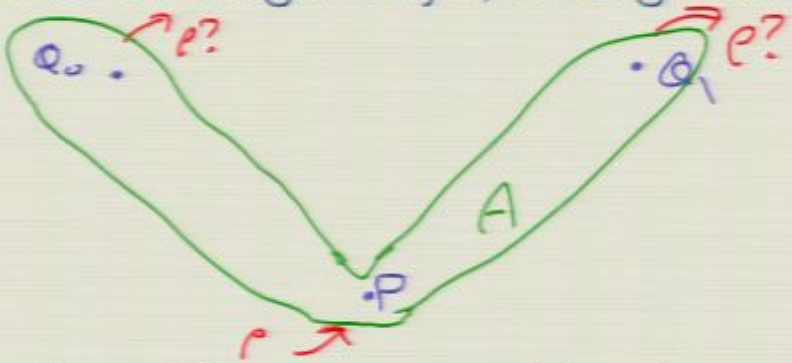

Security against Alice: ensured by the no-summoning theorem -- she cannot return ρ independently at points on both light rays.

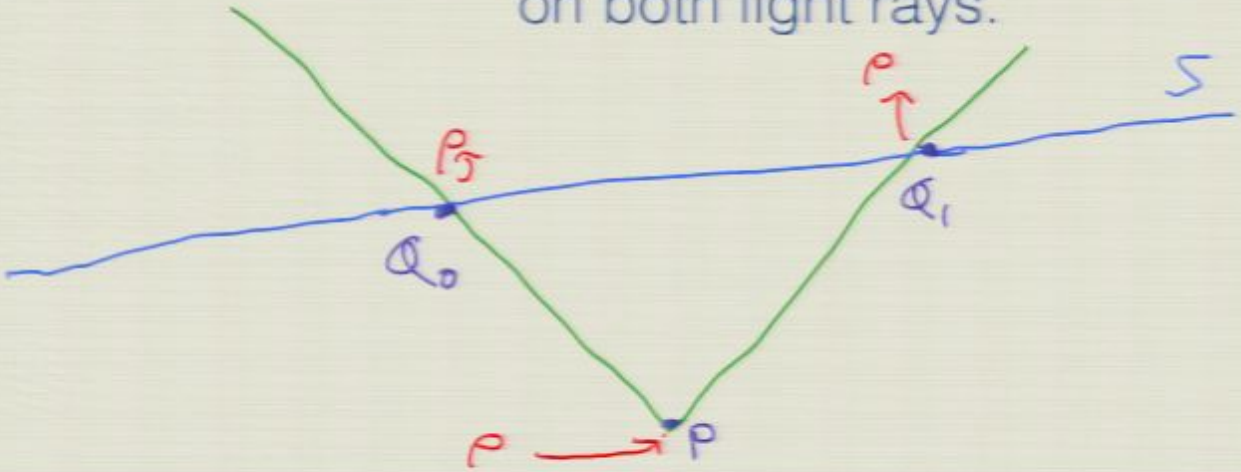

More precisely, we can quantify the security in terms of the dimension d of the space of the unknown state: Alice's cheating probability is bounded by O(1/d).

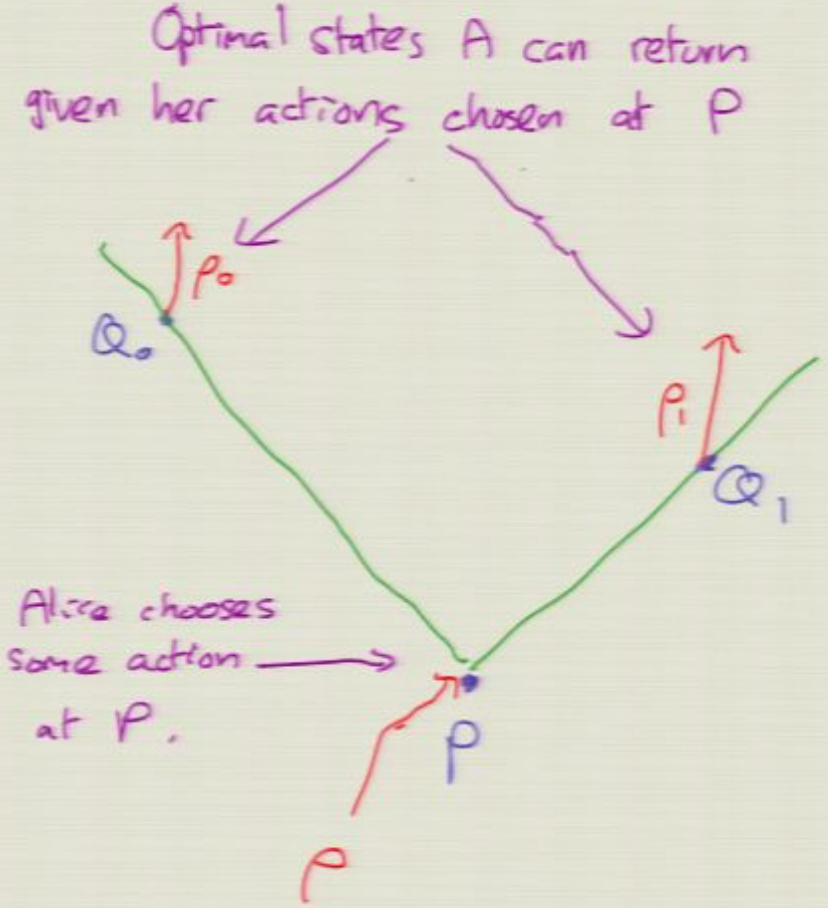

$$P(\text{Bob accepts unveiling at } Q_0) + P(\text{Bob accepts unveiling at } Q_1)$$

$$= Tr(\rho P_0) + Tr(\rho P_1)$$

$$\leq 1 + \frac{2}{d+1}$$

Alice's "wiggle room" decays exponentially in # qubits $= \log_2(d)$.

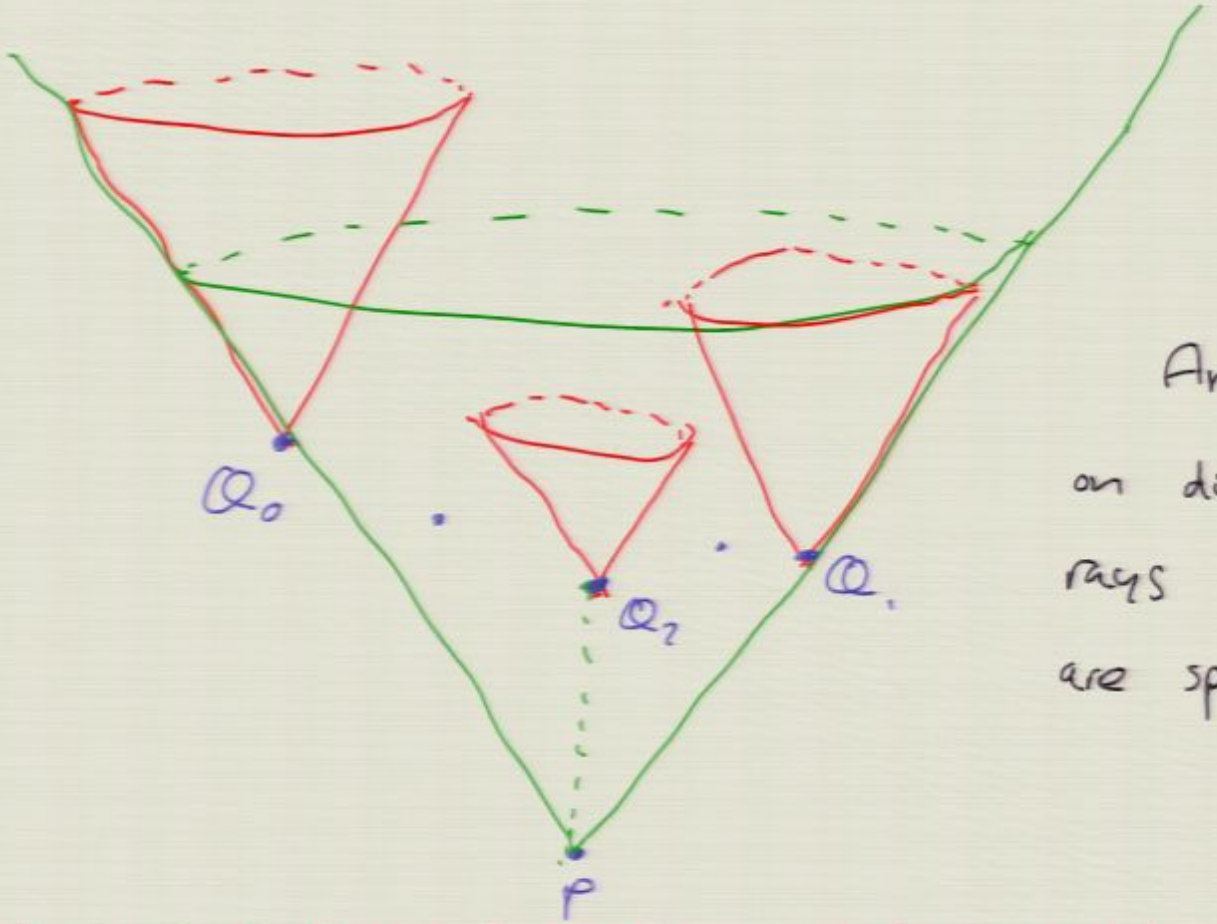

Any 2 points on distinct light rays through P are spacelike separated

This works in 3+1 dimensions also -- and now each possible light like direction can code for a different data value, so the amount of data committed is bounded only by the precision of Alice's transmission and Bob's measurement.