Title: MIP* = RE

Speakers: Henry Yuen

Series: Colloquium

Date: May 13, 2020 - 2:00 PM

URL: http://pirsa.org/20050011

Abstract: MIP* denotes the class of problems that admit interactive proofs with quantum entangled provers. It has been an outstanding question to characterize the complexity of this class. Most notably, there was no known computable upper bound on MIP*.

We show that MIP* is equal to the class RE, the set of recursively enumerable languages. In particular, this shows that MIP* contains uncomputable problems. Through a series of known connections, this also yields a negative answer to Connesâ€™ Embedding Problem from the theory of operator algebras. In this talk, I will explain the connection between Connes' Embedding Problem, quantum information theory, and complexity theory. I will then give an overview of our approach, which involves reducing the Halting Problem to the problem of approximating the entangled value of nonlocal games.

Joint work with Zhengfeng Ji, Anand Natarajan, Thomas Vidick, and John Wright.

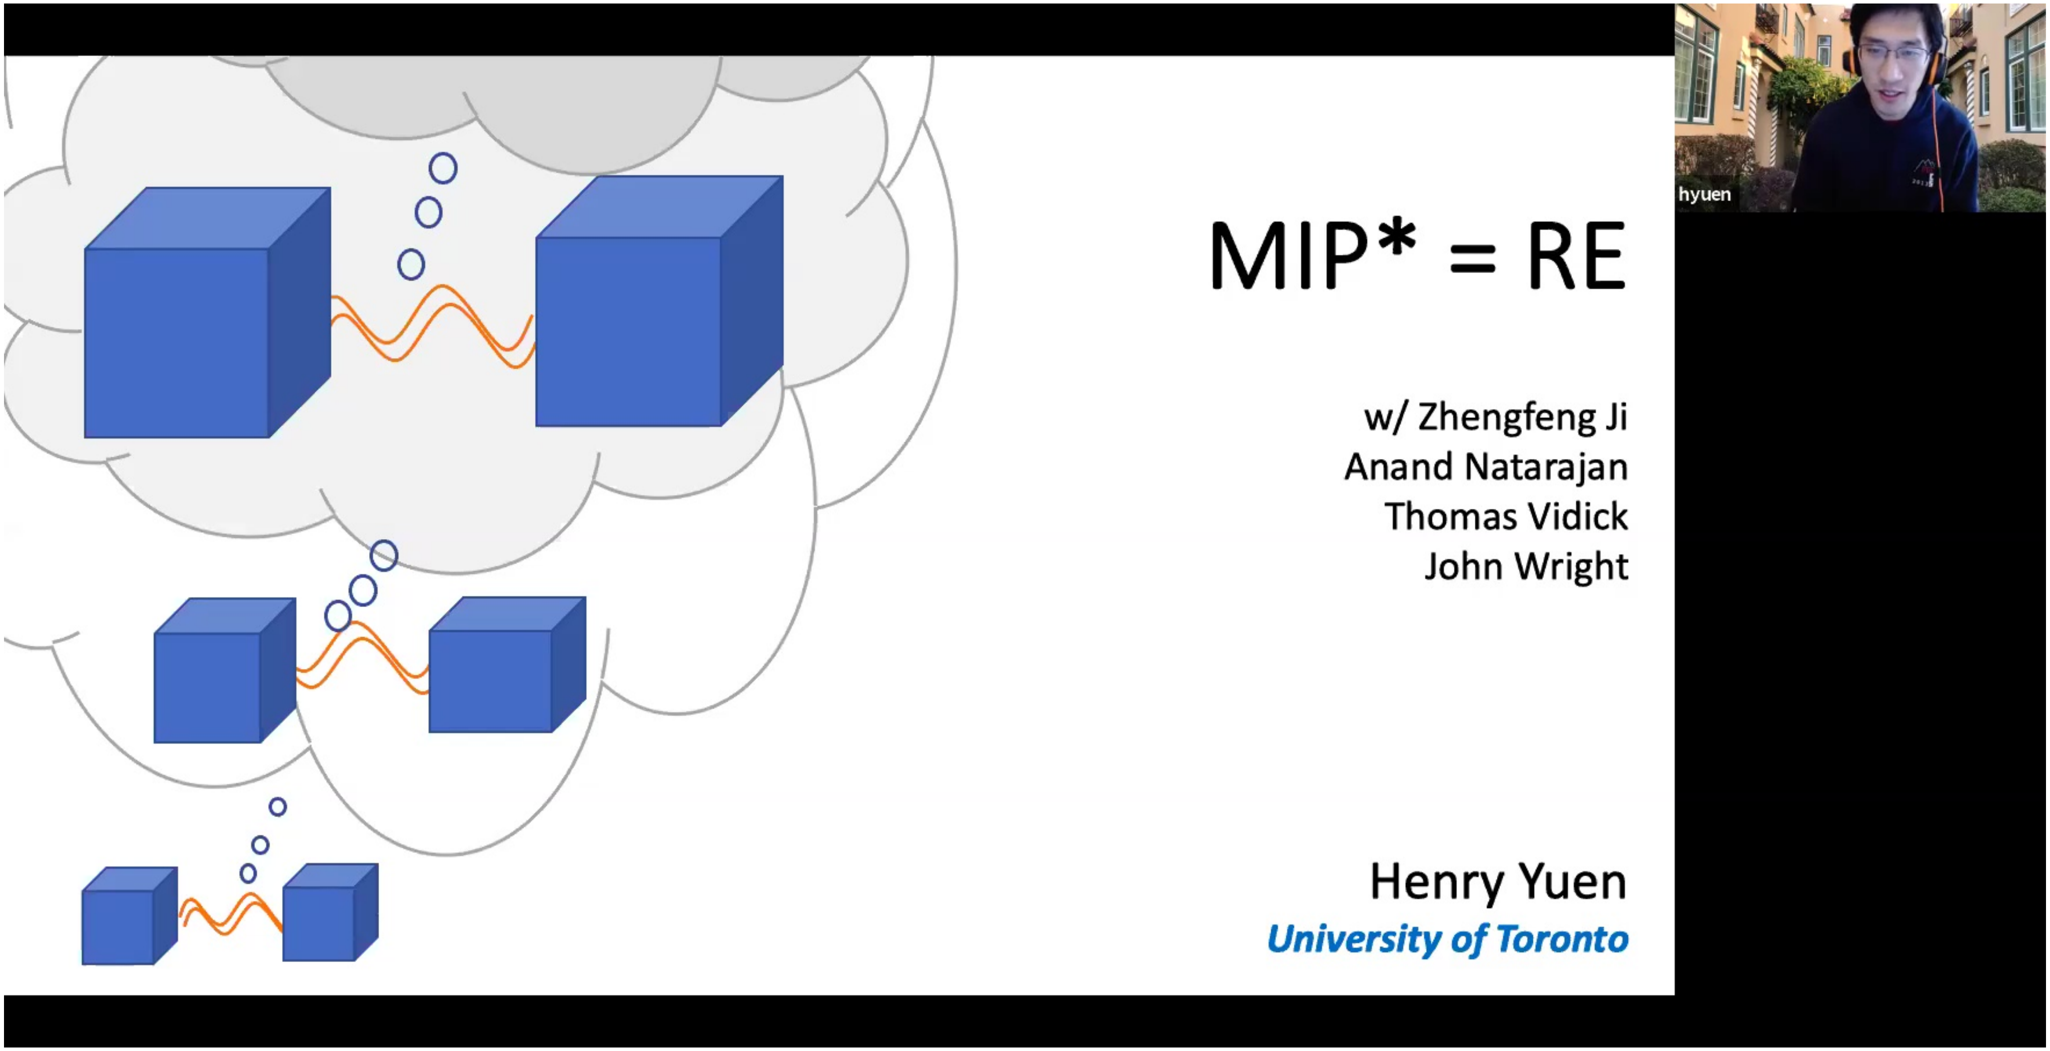

MIP* = RE

w/ Zhengfeng Ji
Anand Natarajan
Thomas Vidick
John Wright

Henry Yuen
University of Toronto

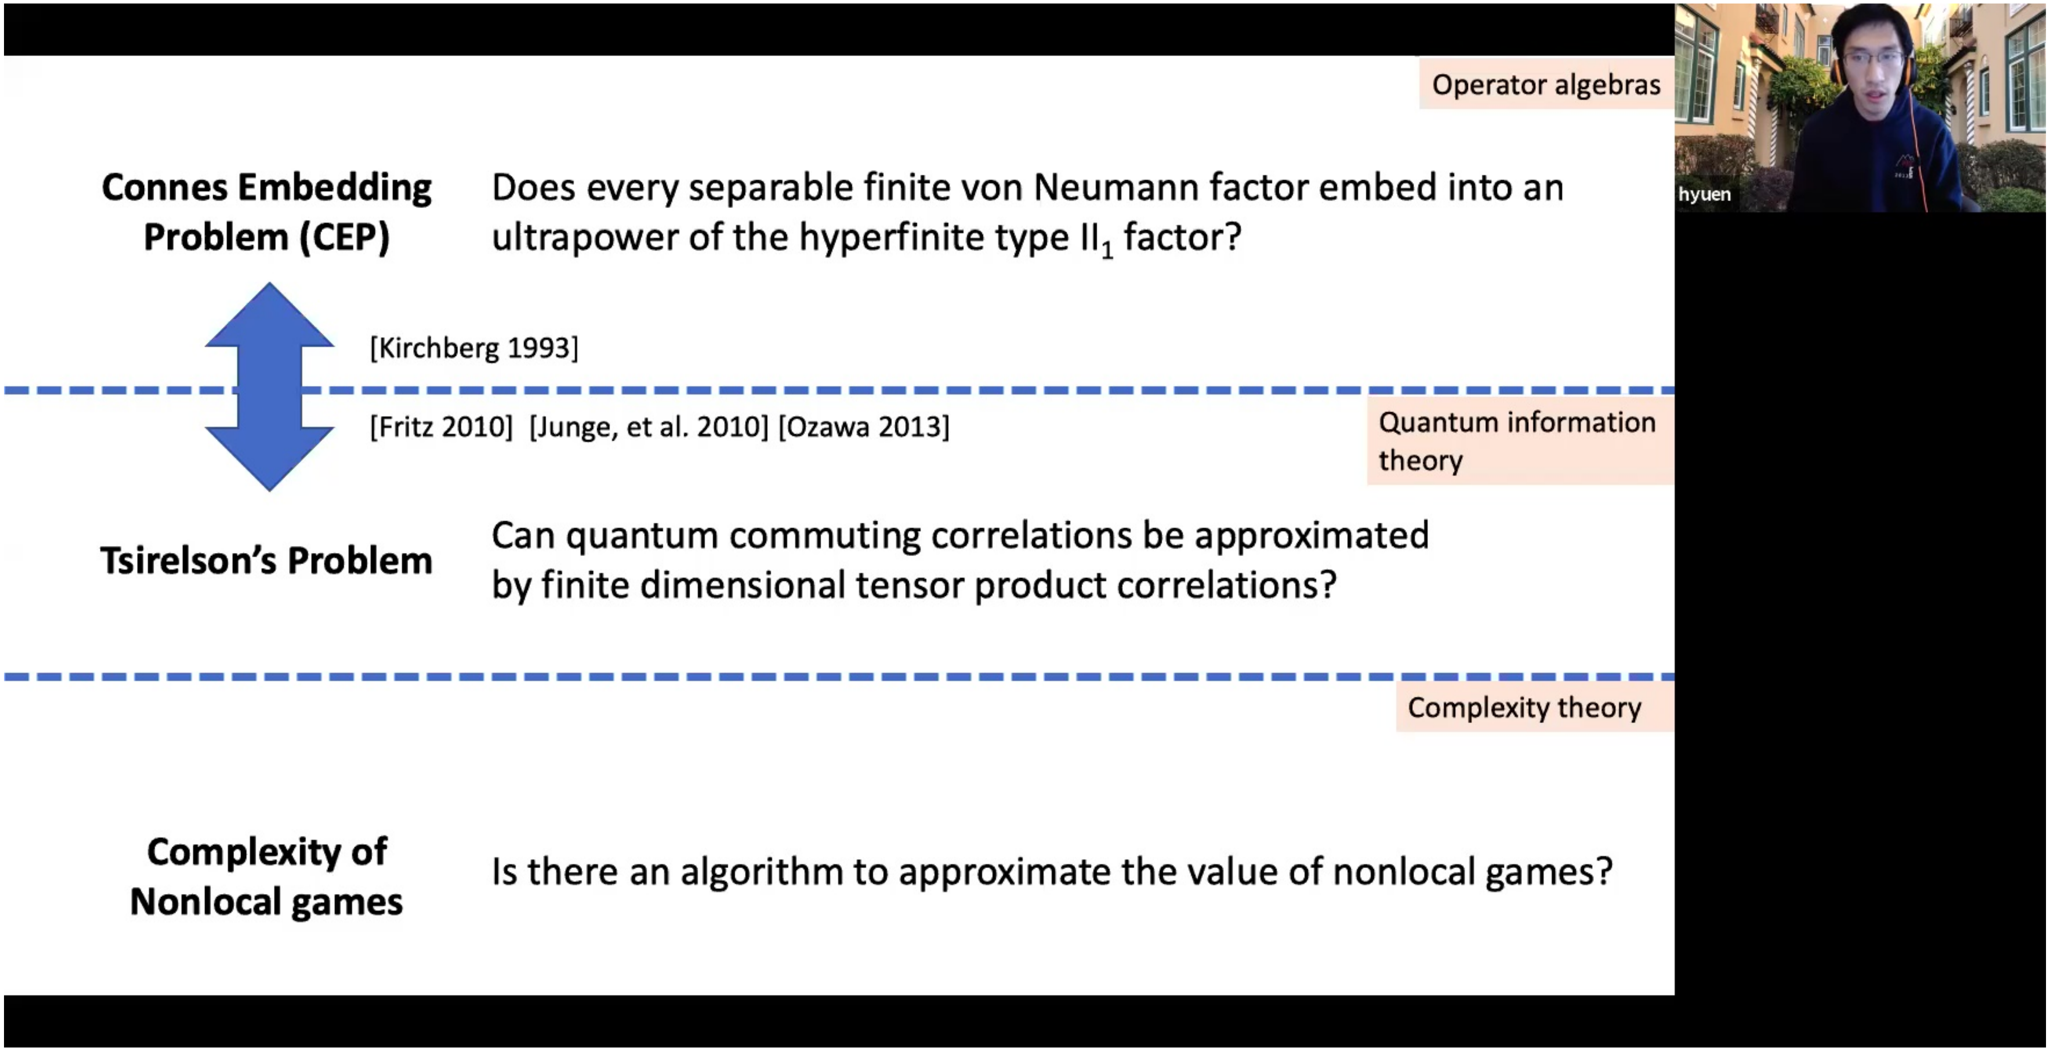

Operator algebras

Connes Embedding Problem (CEP)

Does every separable finite von Neumann factor embed into an ultrapower of the hyperfinite type II_1 factor?

[Kirchberg 1993]

[Fritz 2010] [Junge, et al. 2010] [Ozawa 2013]

Quantum information theory

Tsirelson's Problem

Can quantum commuting correlations be approximated by finite dimensional tensor product correlations?

Complexity theory

Complexity of Nonlocal games

Is there an algorithm to approximate the value of nonlocal games?

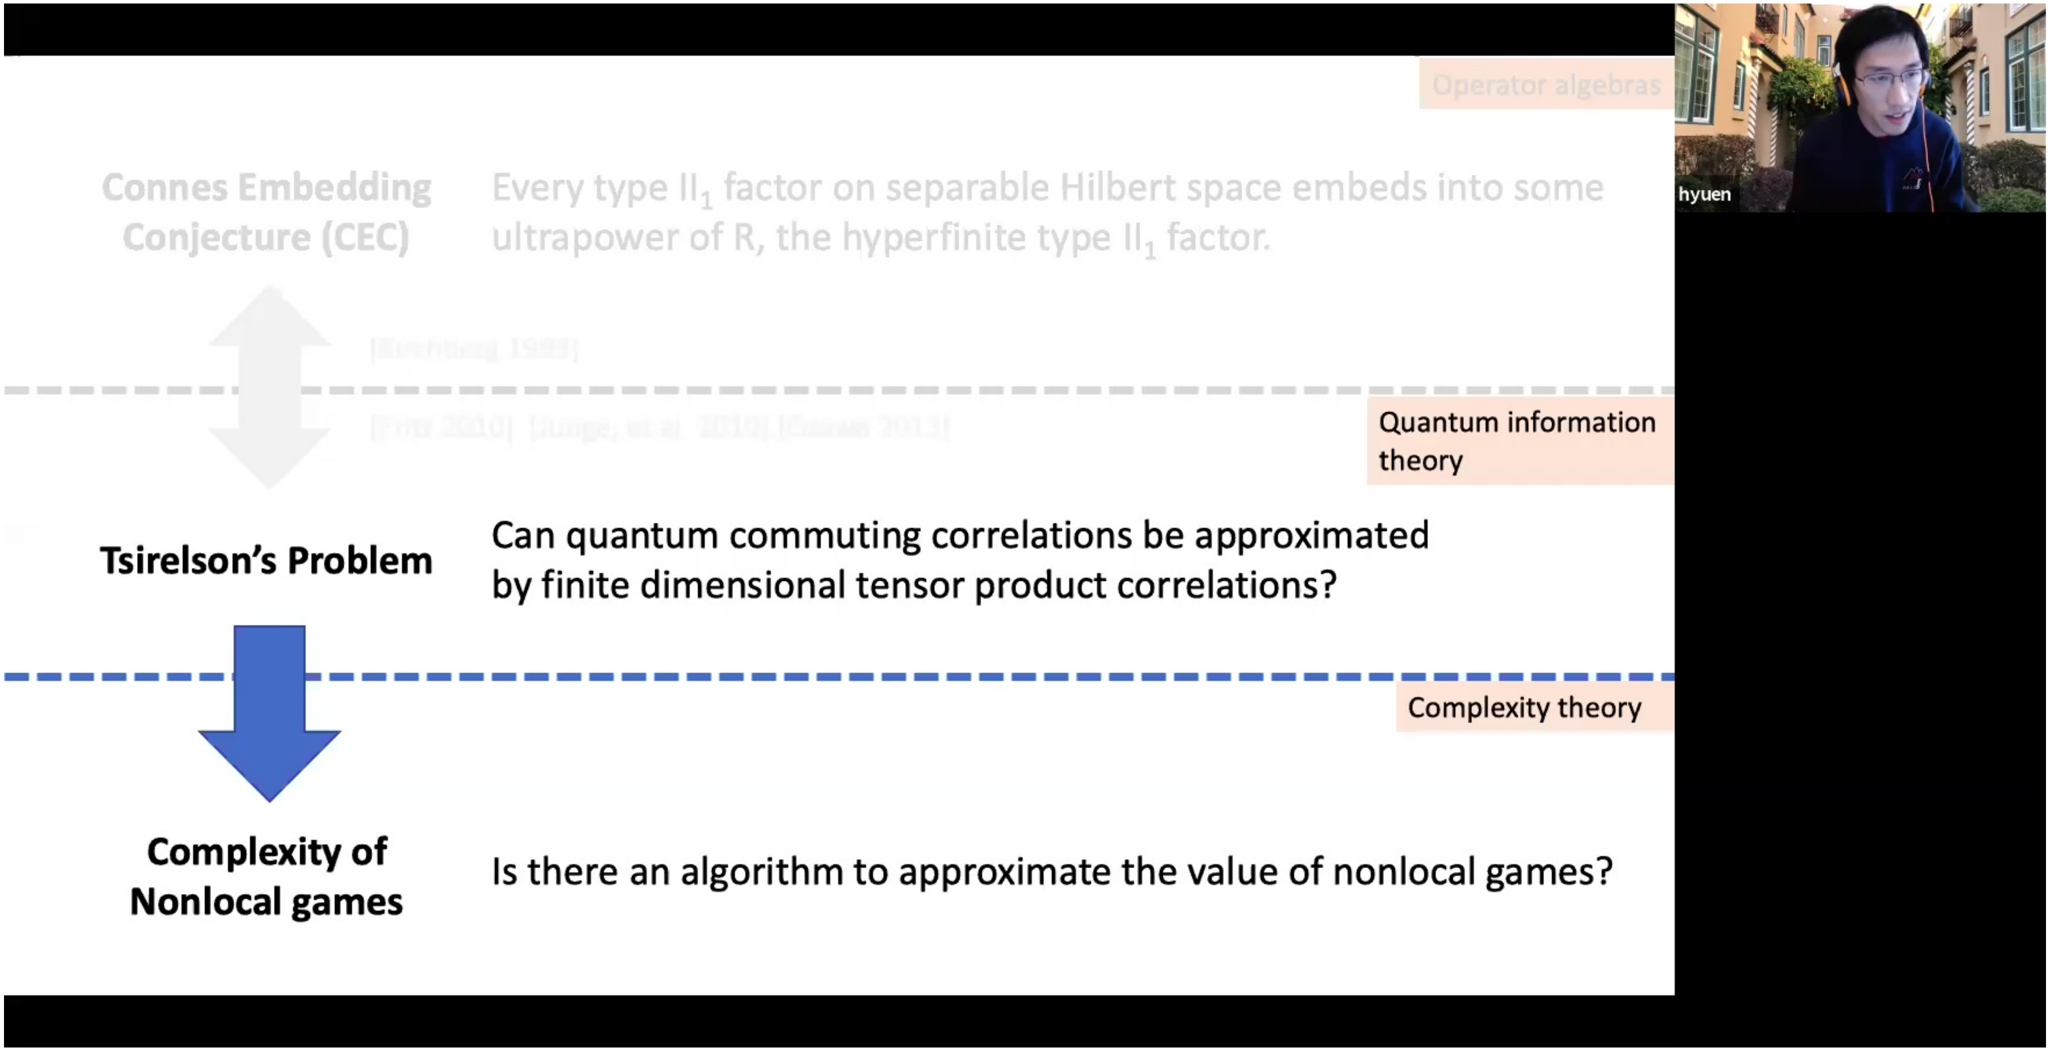

Operator algebras

Connes Embedding Conjecture (CEC)

Every type II_1 factor on separable Hilbert space embeds into some ultrapower of R, the hyperfinite type II_1 factor.

[Kirchberg 1993]

[Fritz 2010] [Junge, et al. 2010] [Ozawa 2013]

Quantum information theory

Tsirelson's Problem

Can quantum commuting correlations be approximated by finite dimensional tensor product correlations?

Complexity theory

Complexity of Nonlocal games

Is there an algorithm to approximate the value of nonlocal games?

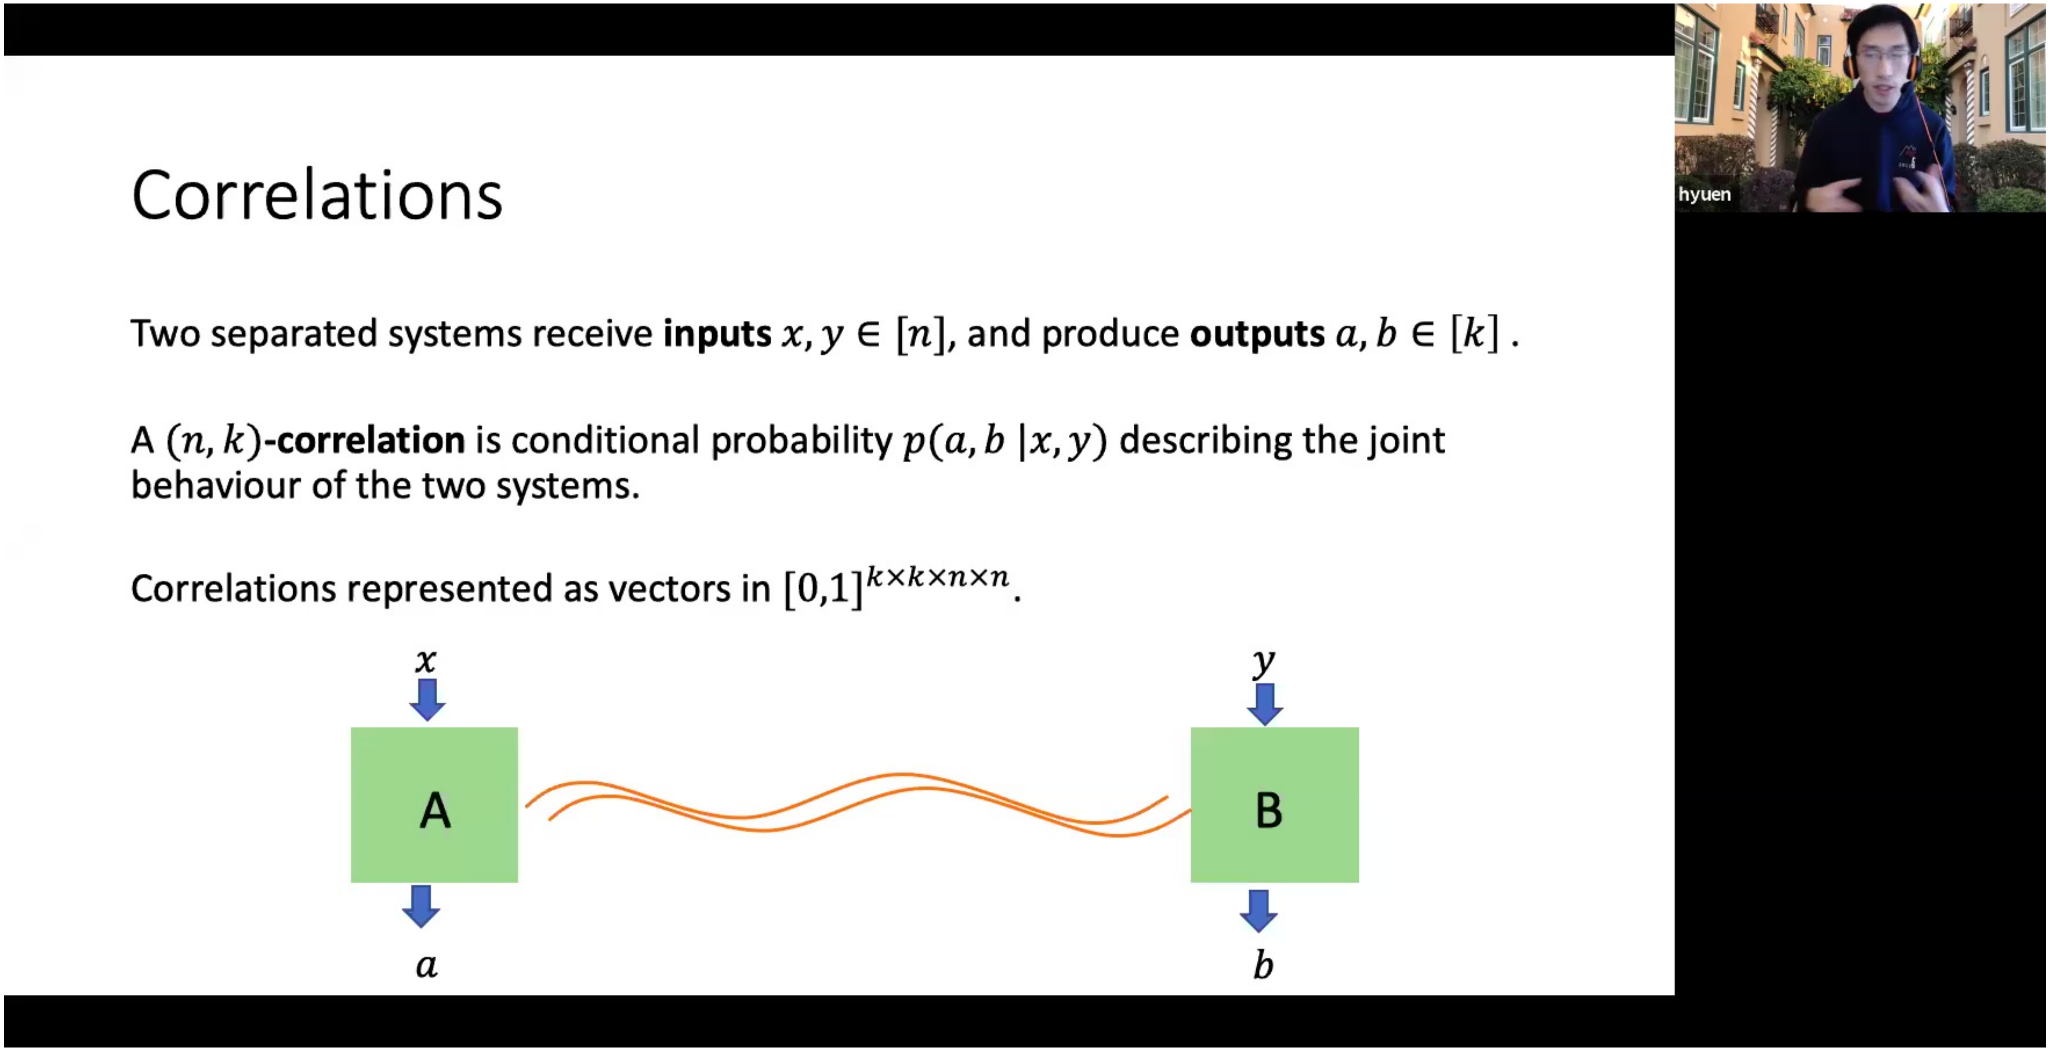

Correlations

Two separated systems receive **inputs** $x, y \in [n]$, and produce **outputs** $a, b \in [k]$.

A (n, k)**-correlation** is conditional probability $p(a, b\ |x, y)$ describing the joint behaviour of the two systems.

Correlations represented as vectors in $[0,1]^{k\times k\times n\times n}$.

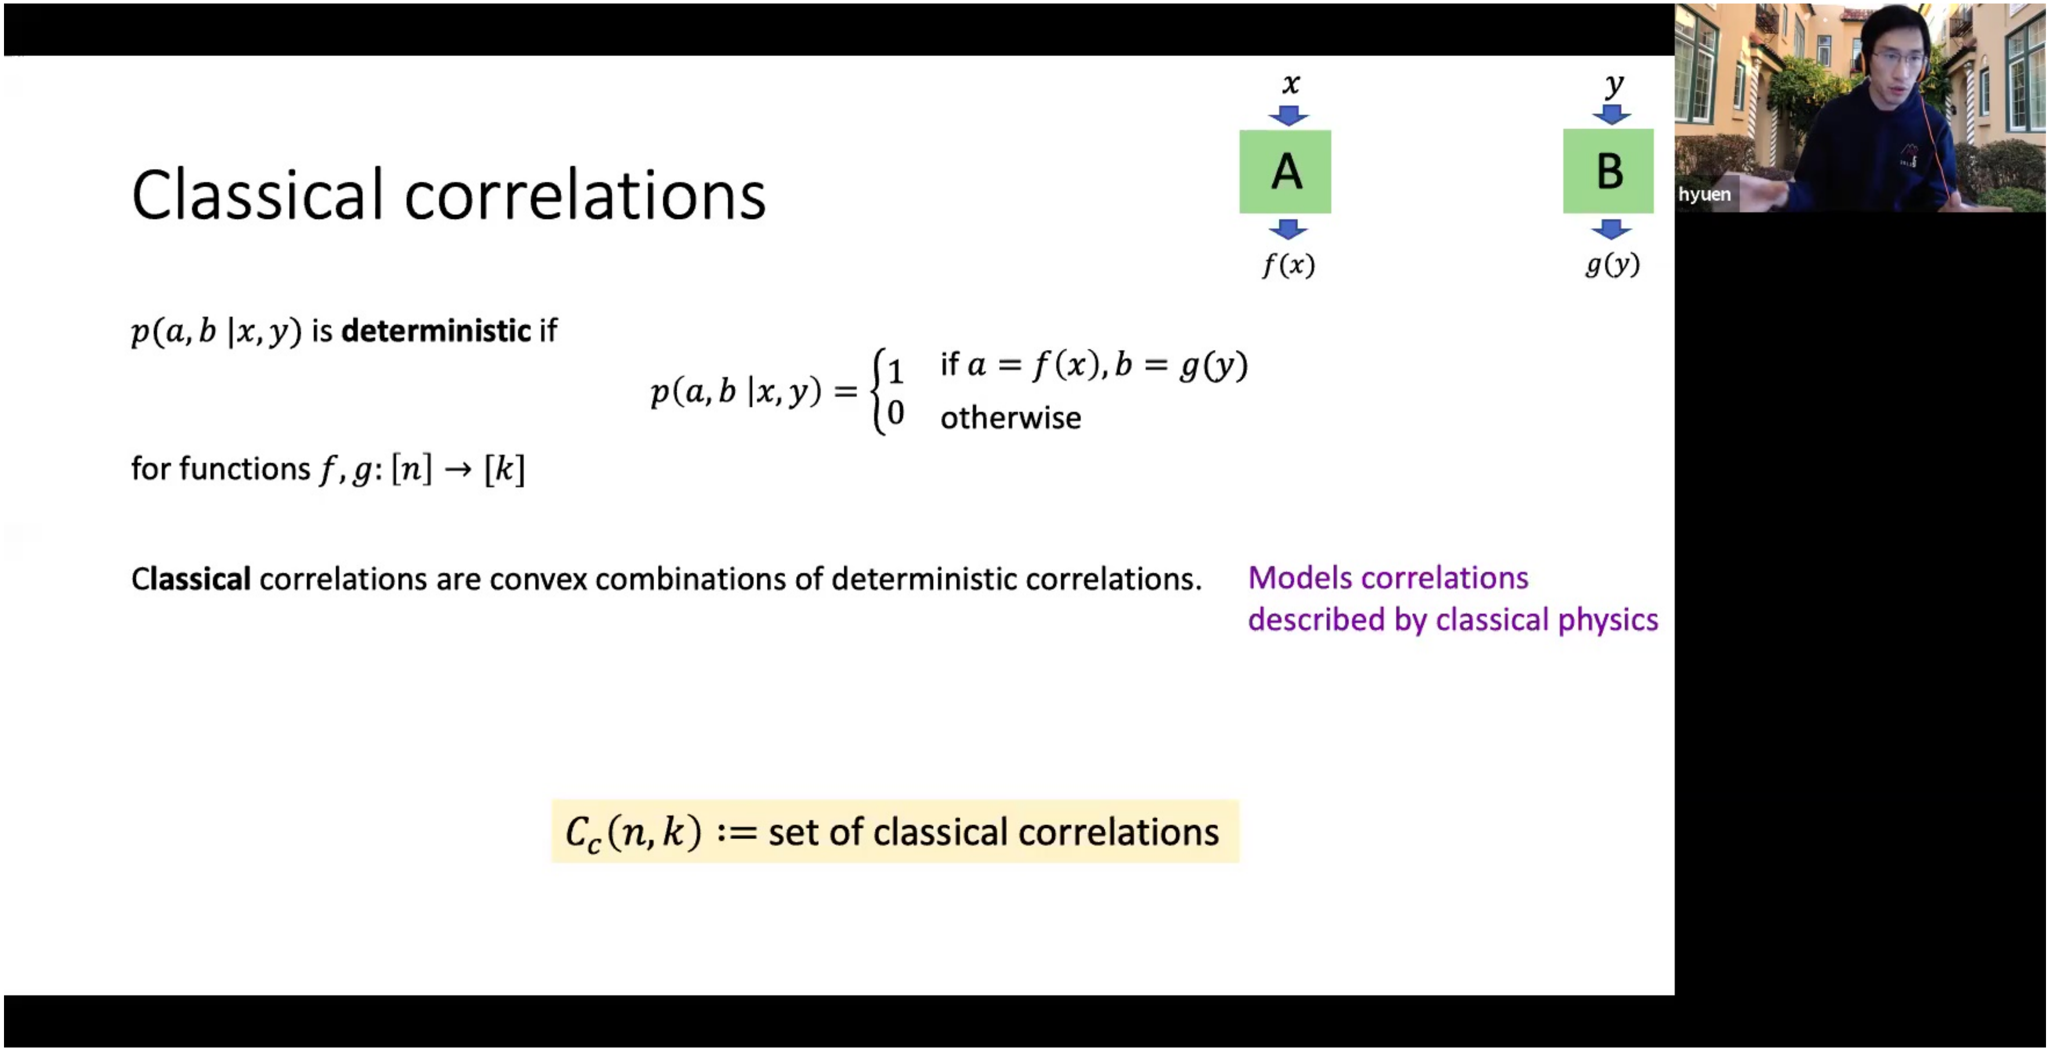

Classical correlations

$p(a, b \mid x, y)$ is **deterministic** if

$$p(a, b \mid x, y) = \begin{cases} 1 & \text{if } a = f(x), b = g(y) \\ 0 & \text{otherwise} \end{cases}$$

for functions $f, g: [n] \to [k]$

Classical correlations are convex combinations of deterministic correlations.

Models correlations described by classical physics

$C_c(n, k) :=$ set of classical correlations

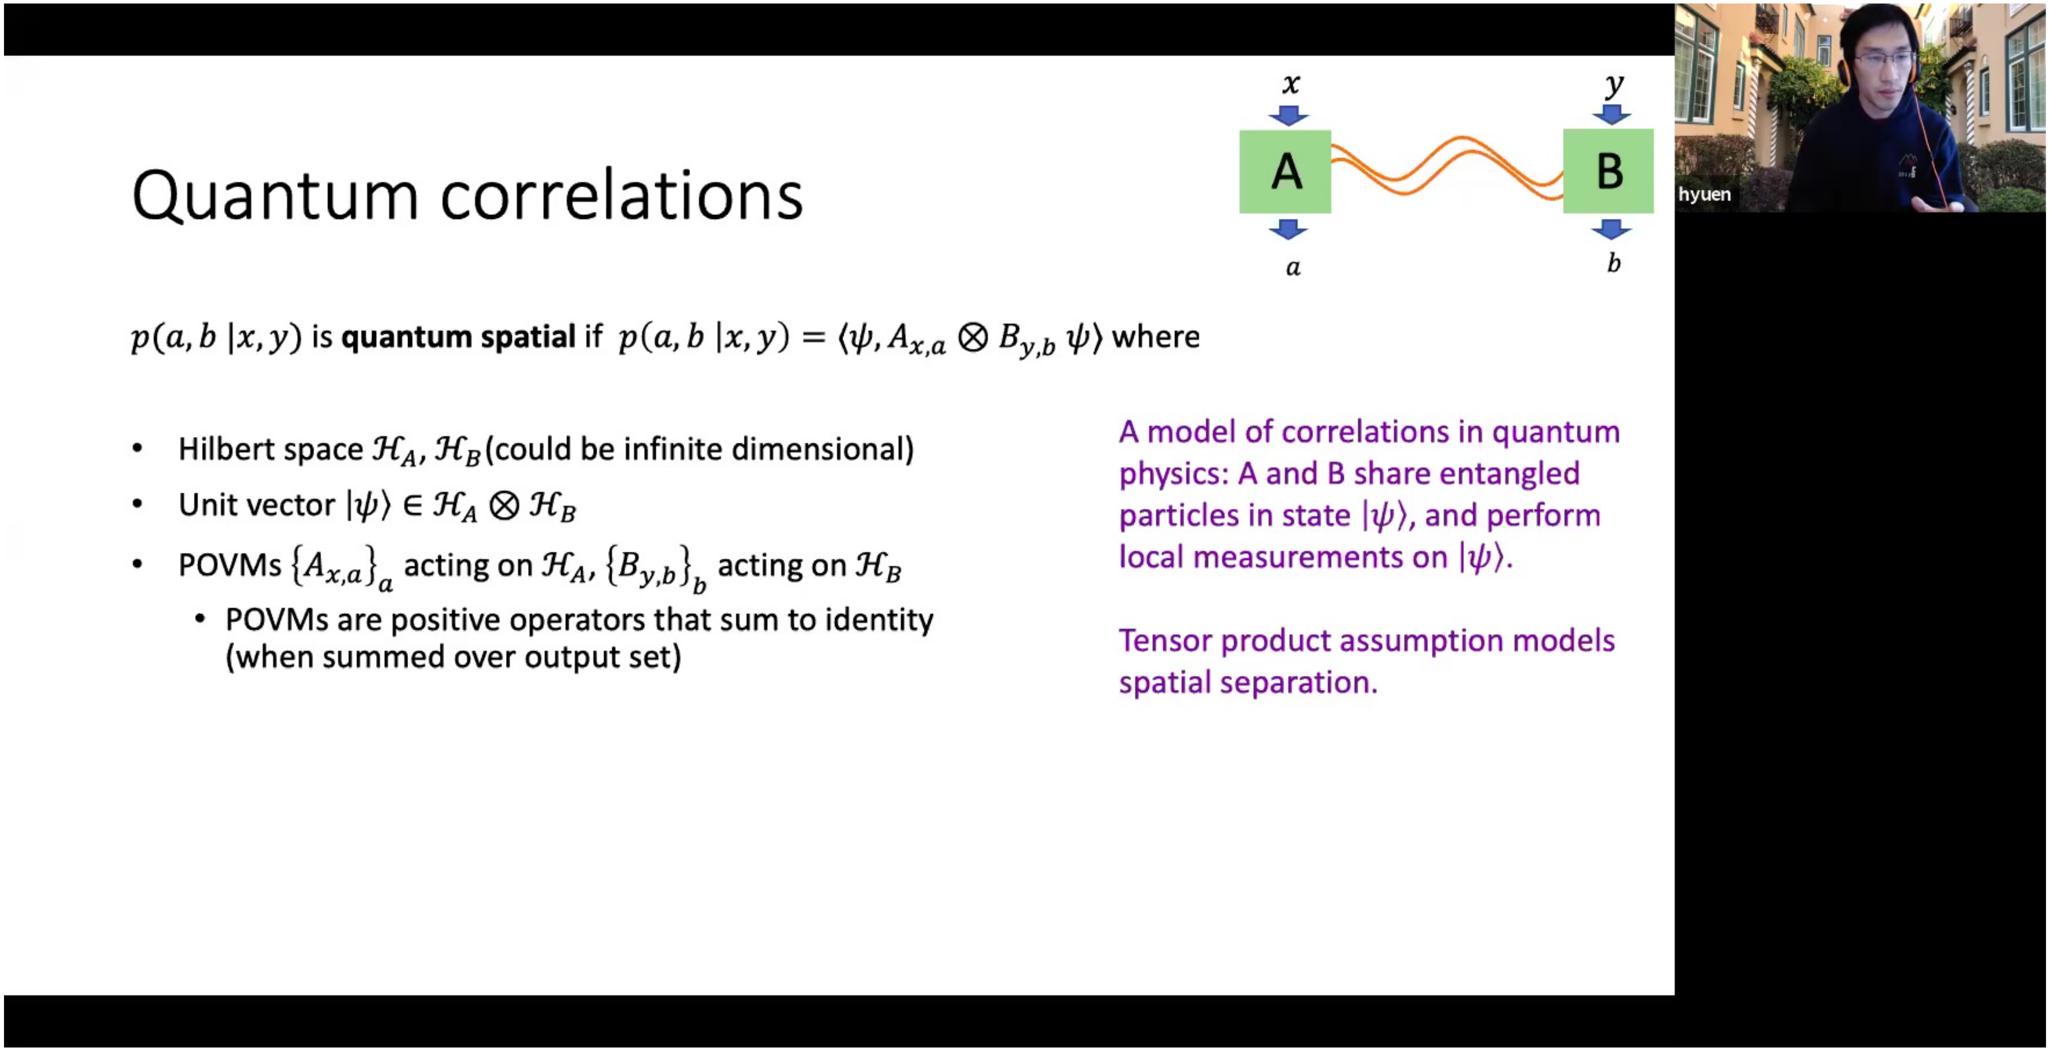

Quantum correlations

$p(a, b \mid x, y)$ is **quantum spatial** if $p(a, b \mid x, y) = \langle \psi, A_{x,a} \otimes B_{y,b}\, \psi \rangle$ where

- Hilbert space $\mathcal{H}_A, \mathcal{H}_B$(could be infinite dimensional)
- Unit vector $|\psi\rangle \in \mathcal{H}_A \otimes \mathcal{H}_B$
- POVMs $\{A_{x,a}\}_a$ acting on $\mathcal{H}_A$, $\{B_{y,b}\}_b$ acting on $\mathcal{H}_B$
  - POVMs are positive operators that sum to identity (when summed over output set)

A model of correlations in quantum physics: A and B share entangled particles in state $|\psi\rangle$, and perform local measurements on $|\psi\rangle$.

Tensor product assumption models spatial separation.

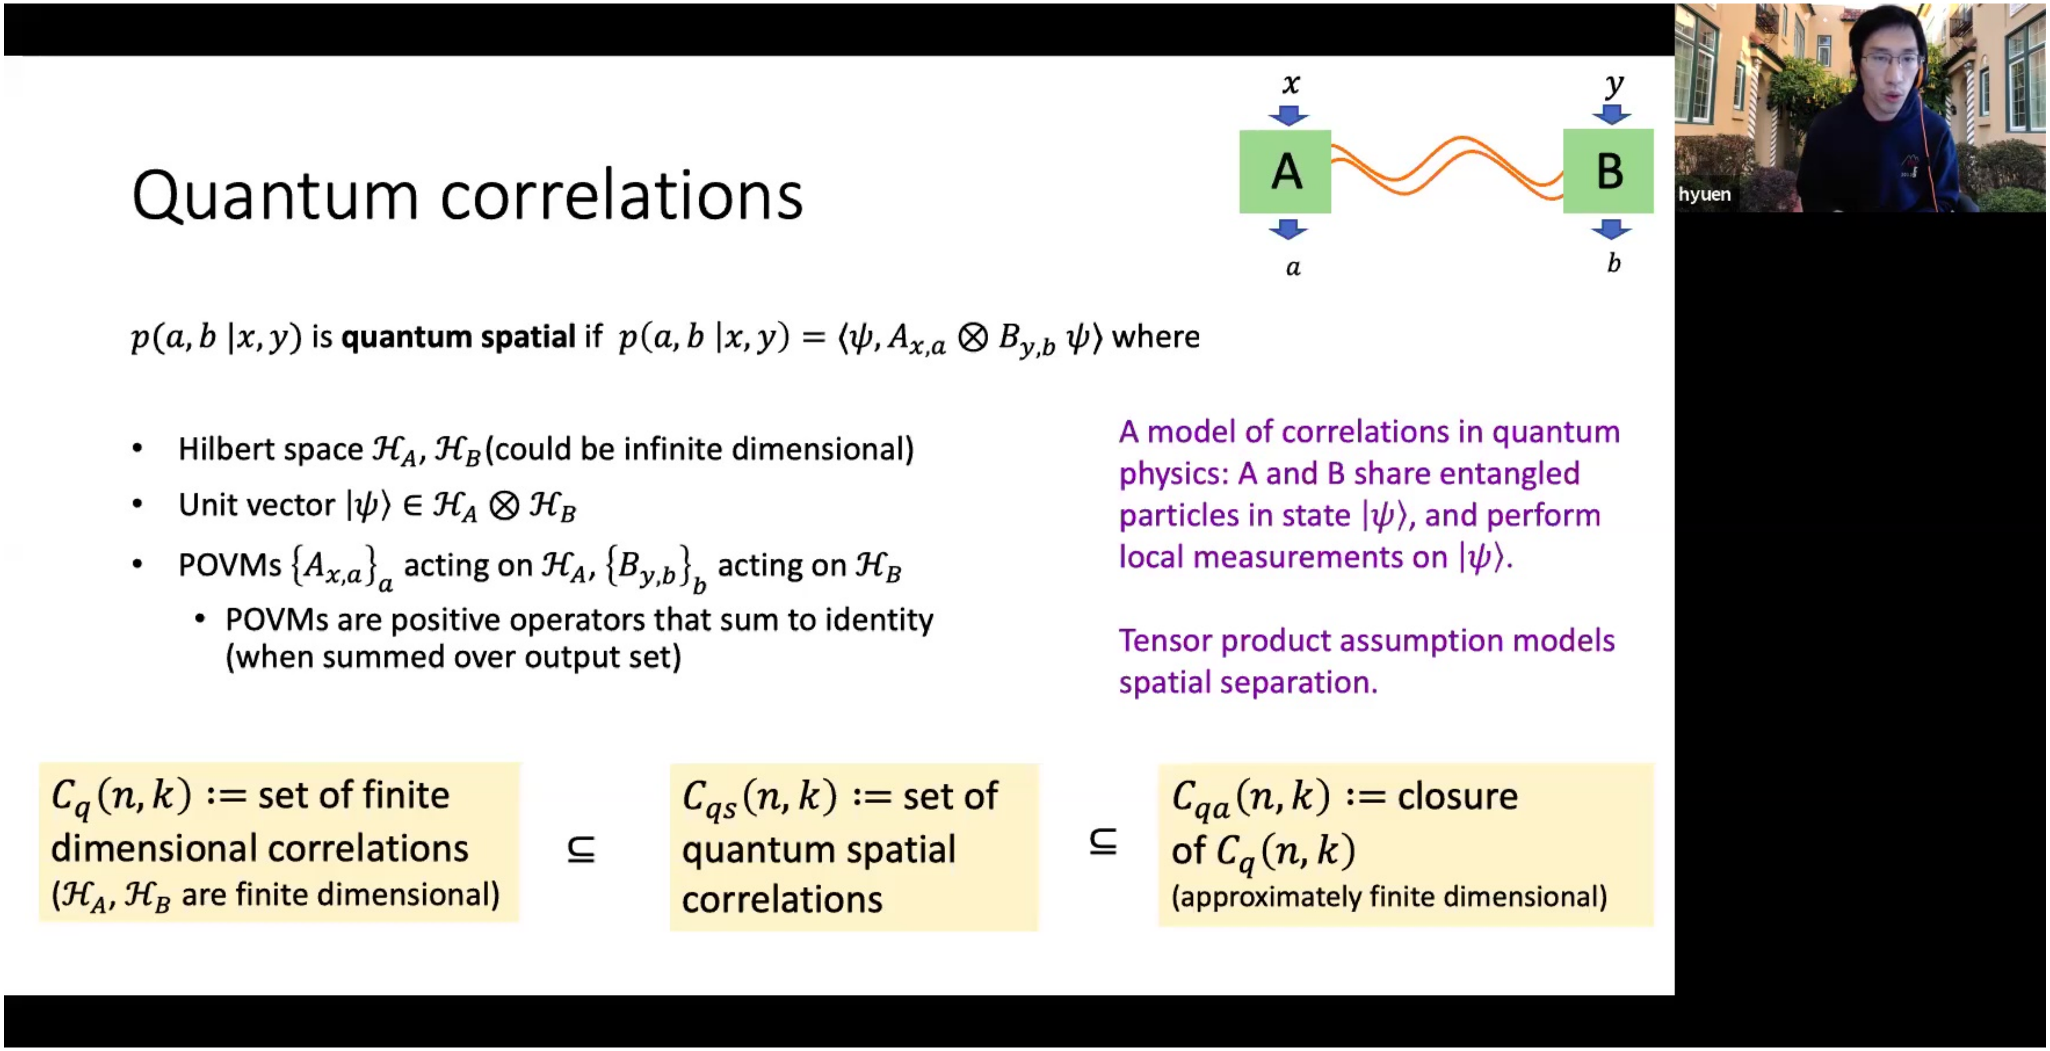

Quantum correlations

$p(a, b \,|x, y)$ is **quantum spatial** if $p(a, b\,|x, y) = \langle \psi, A_{x,a} \otimes B_{y,b}\, \psi \rangle$ where

- Hilbert space $\mathcal{H}_A, \mathcal{H}_B$(could be infinite dimensional)
- Unit vector $|\psi\rangle \in \mathcal{H}_A \otimes \mathcal{H}_B$
- POVMs $\{A_{x,a}\}_a$ acting on $\mathcal{H}_A$, $\{B_{y,b}\}_b$ acting on $\mathcal{H}_B$
  - POVMs are positive operators that sum to identity (when summed over output set)

A model of correlations in quantum physics: A and B share entangled particles in state $|\psi\rangle$, and perform local measurements on $|\psi\rangle$.

Tensor product assumption models spatial separation.

$C_q(n, k)$:= set of finite dimensional correlations ($\mathcal{H}_A, \mathcal{H}_B$ are finite dimensional) $\subseteq$ $C_{qs}(n, k)$:= set of quantum spatial correlations $\subseteq$ $C_{qa}(n, k)$:= closure of $C_q(n, k)$ (approximately finite dimensional)

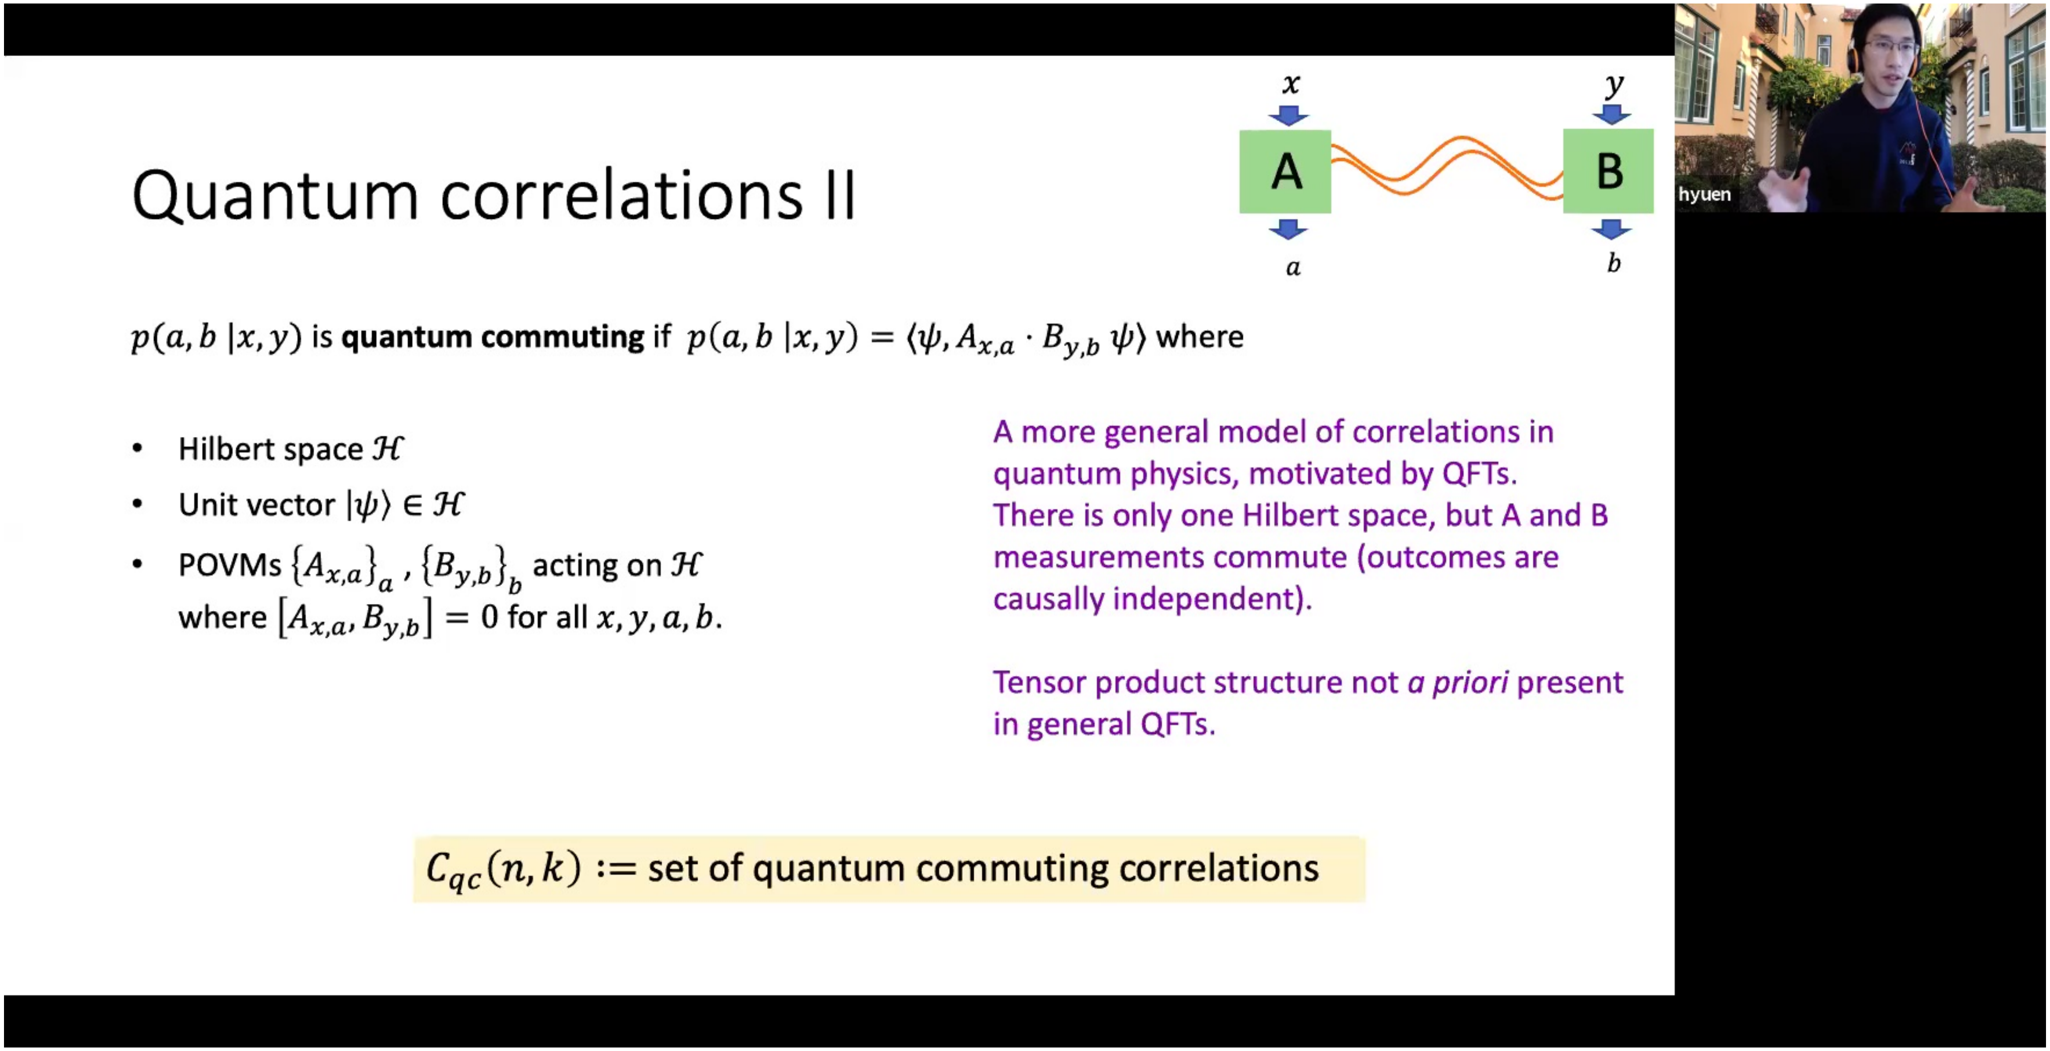

Quantum correlations II

$p(a, b\,|x, y)$ is **quantum commuting** if $p(a, b\,|x, y) = \langle \psi, A_{x,a} \cdot B_{y,b}\, \psi \rangle$ where

- Hilbert space $\mathcal{H}$
- Unit vector $|\psi\rangle \in \mathcal{H}$
- POVMs $\{A_{x,a}\}_a$, $\{B_{y,b}\}_b$ acting on $\mathcal{H}$ where $[A_{x,a}, B_{y,b}] = 0$ for all x, y, a, b.

A more general model of correlations in quantum physics, motivated by QFTs. There is only one Hilbert space, but A and B measurements commute (outcomes are causally independent).

Tensor product structure not *a priori* present in general QFTs.

$C_{qc}(n, k)$:= set of quantum commuting correlations

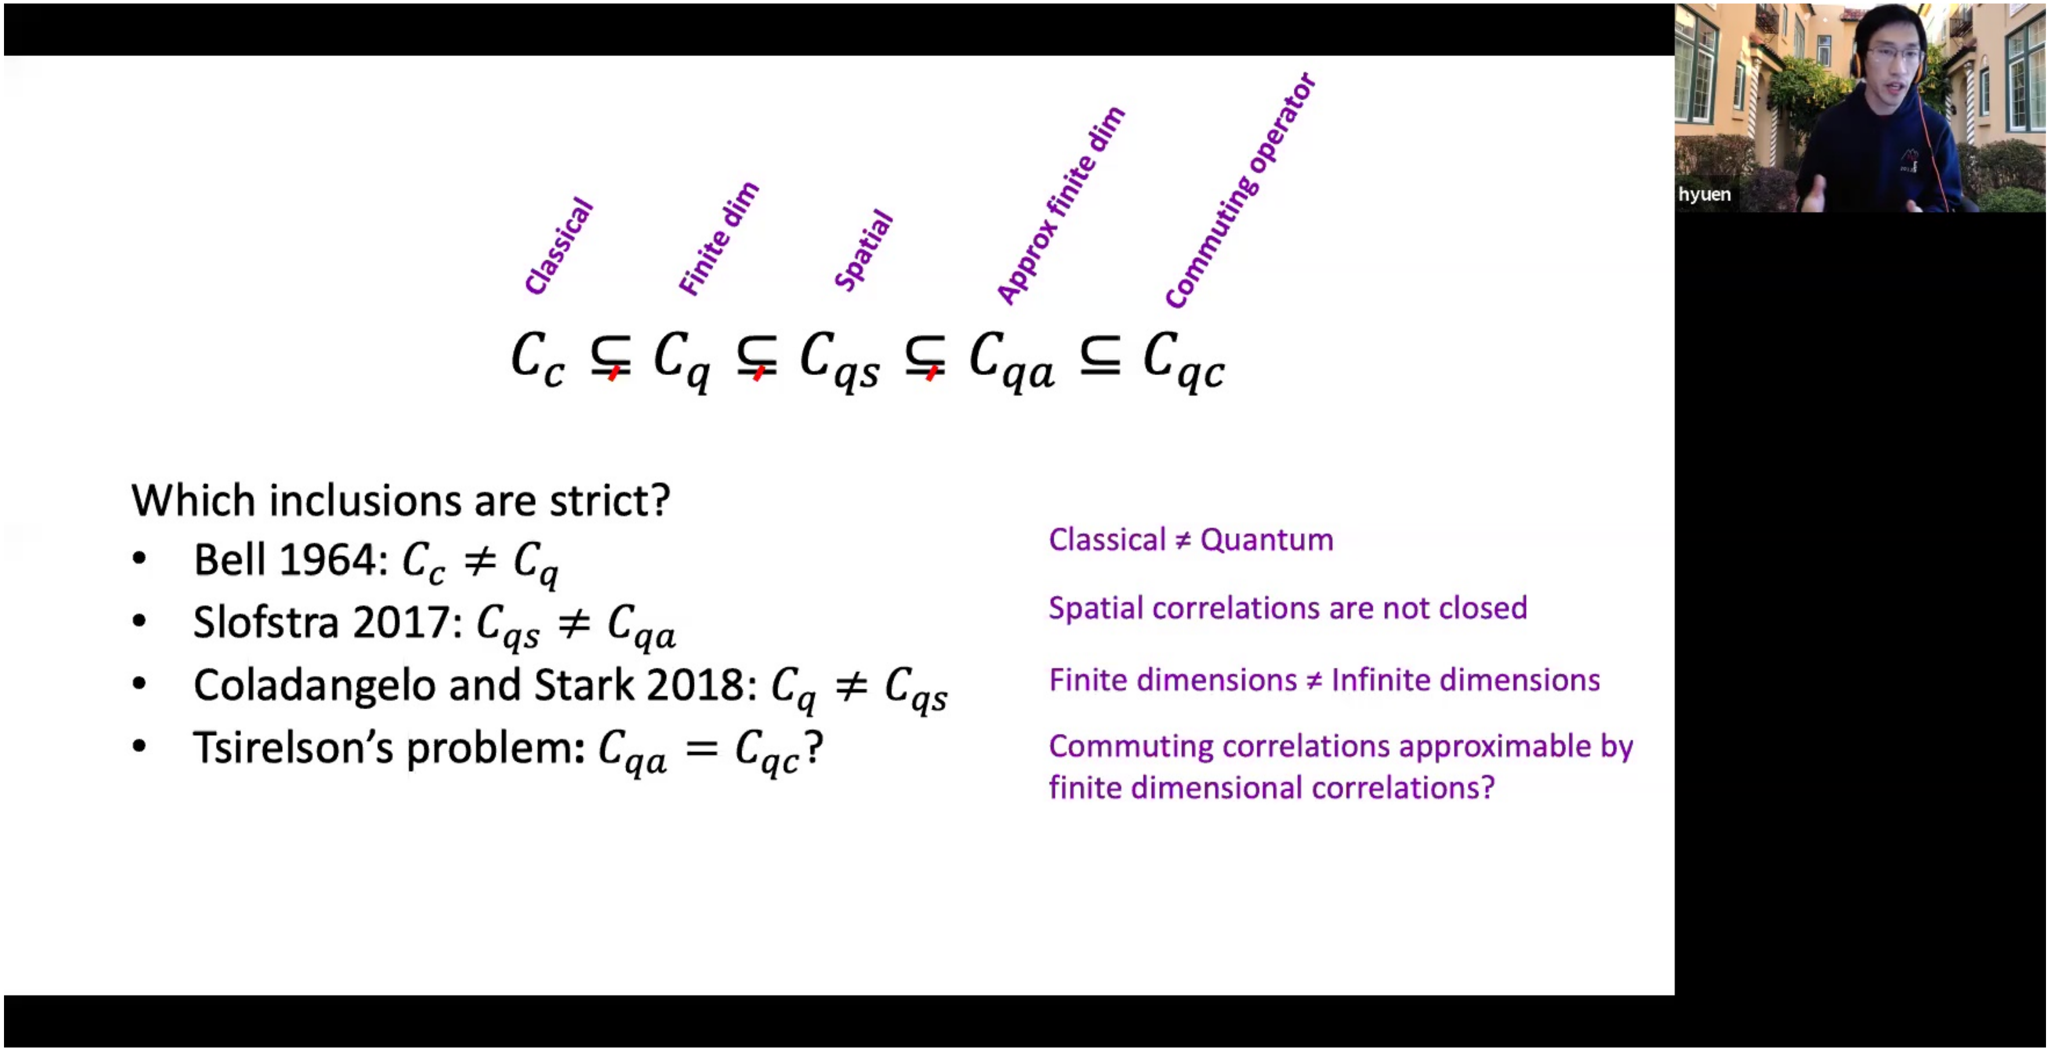

Classical Finite dim Spatial Approx finite dim Commuting operator

$$C_c \subsetneq C_q \subsetneq C_{qs} \subsetneq C_{qa} \subseteq C_{qc}$$

Which inclusions are strict?

- Bell 1964: $C_c \neq C_q$
- Slofstra 2017: $C_{qs} \neq C_{qa}$
- Coladangelo and Stark 2018: $C_q \neq C_{qs}$
- Tsirelson's problem: $C_{qa} = C_{qc}$?

Classical ≠ Quantum

Spatial correlations are not closed

Finite dimensions ≠ Infinite dimensions

Commuting correlations approximable by finite dimensional correlations?

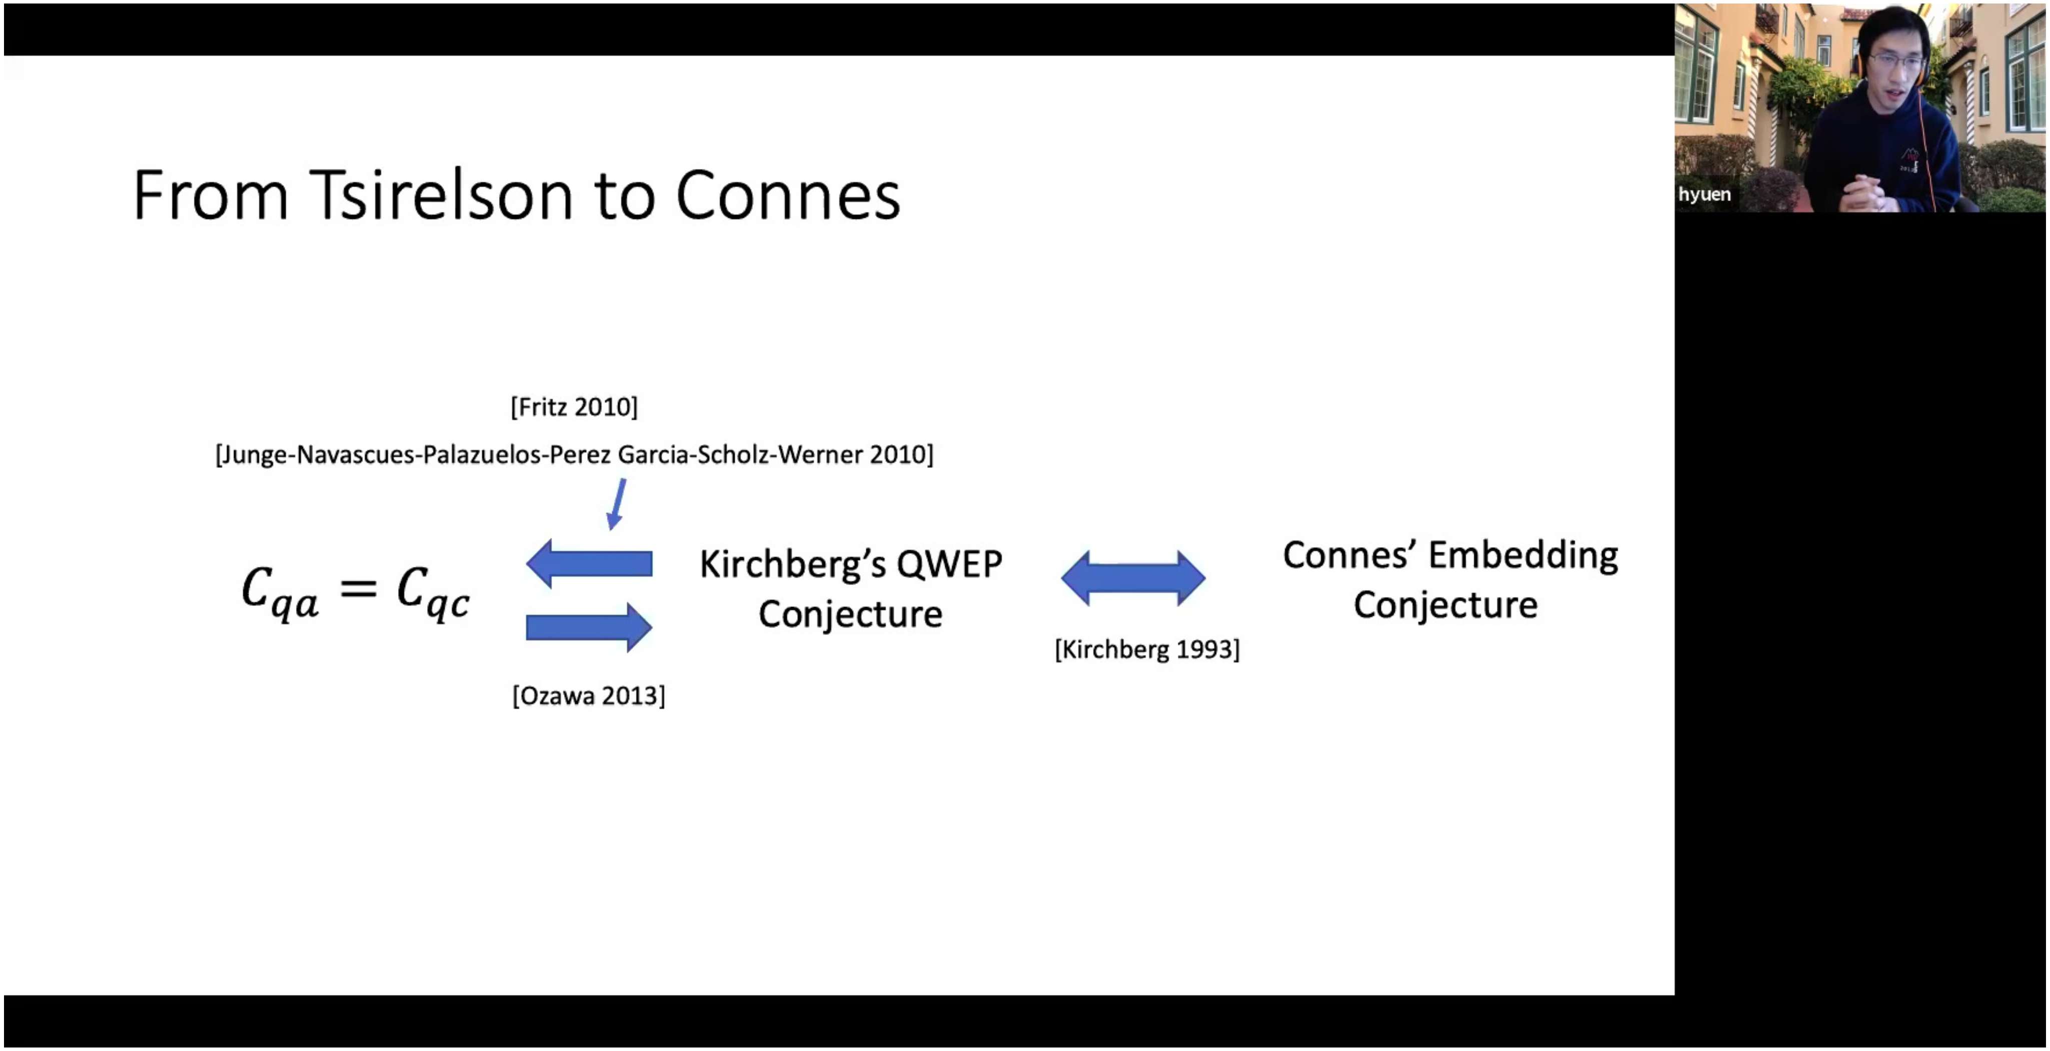

From Tsirelson to Connes

[Fritz 2010]

[Junge-Navascues-Palazuelos-Perez Garcia-Scholz-Werner 2010]

$C_{qa} = C_{qc}$

Kirchberg's QWEP Conjecture

Connes' Embedding Conjecture

[Ozawa 2013]

[Kirchberg 1993]

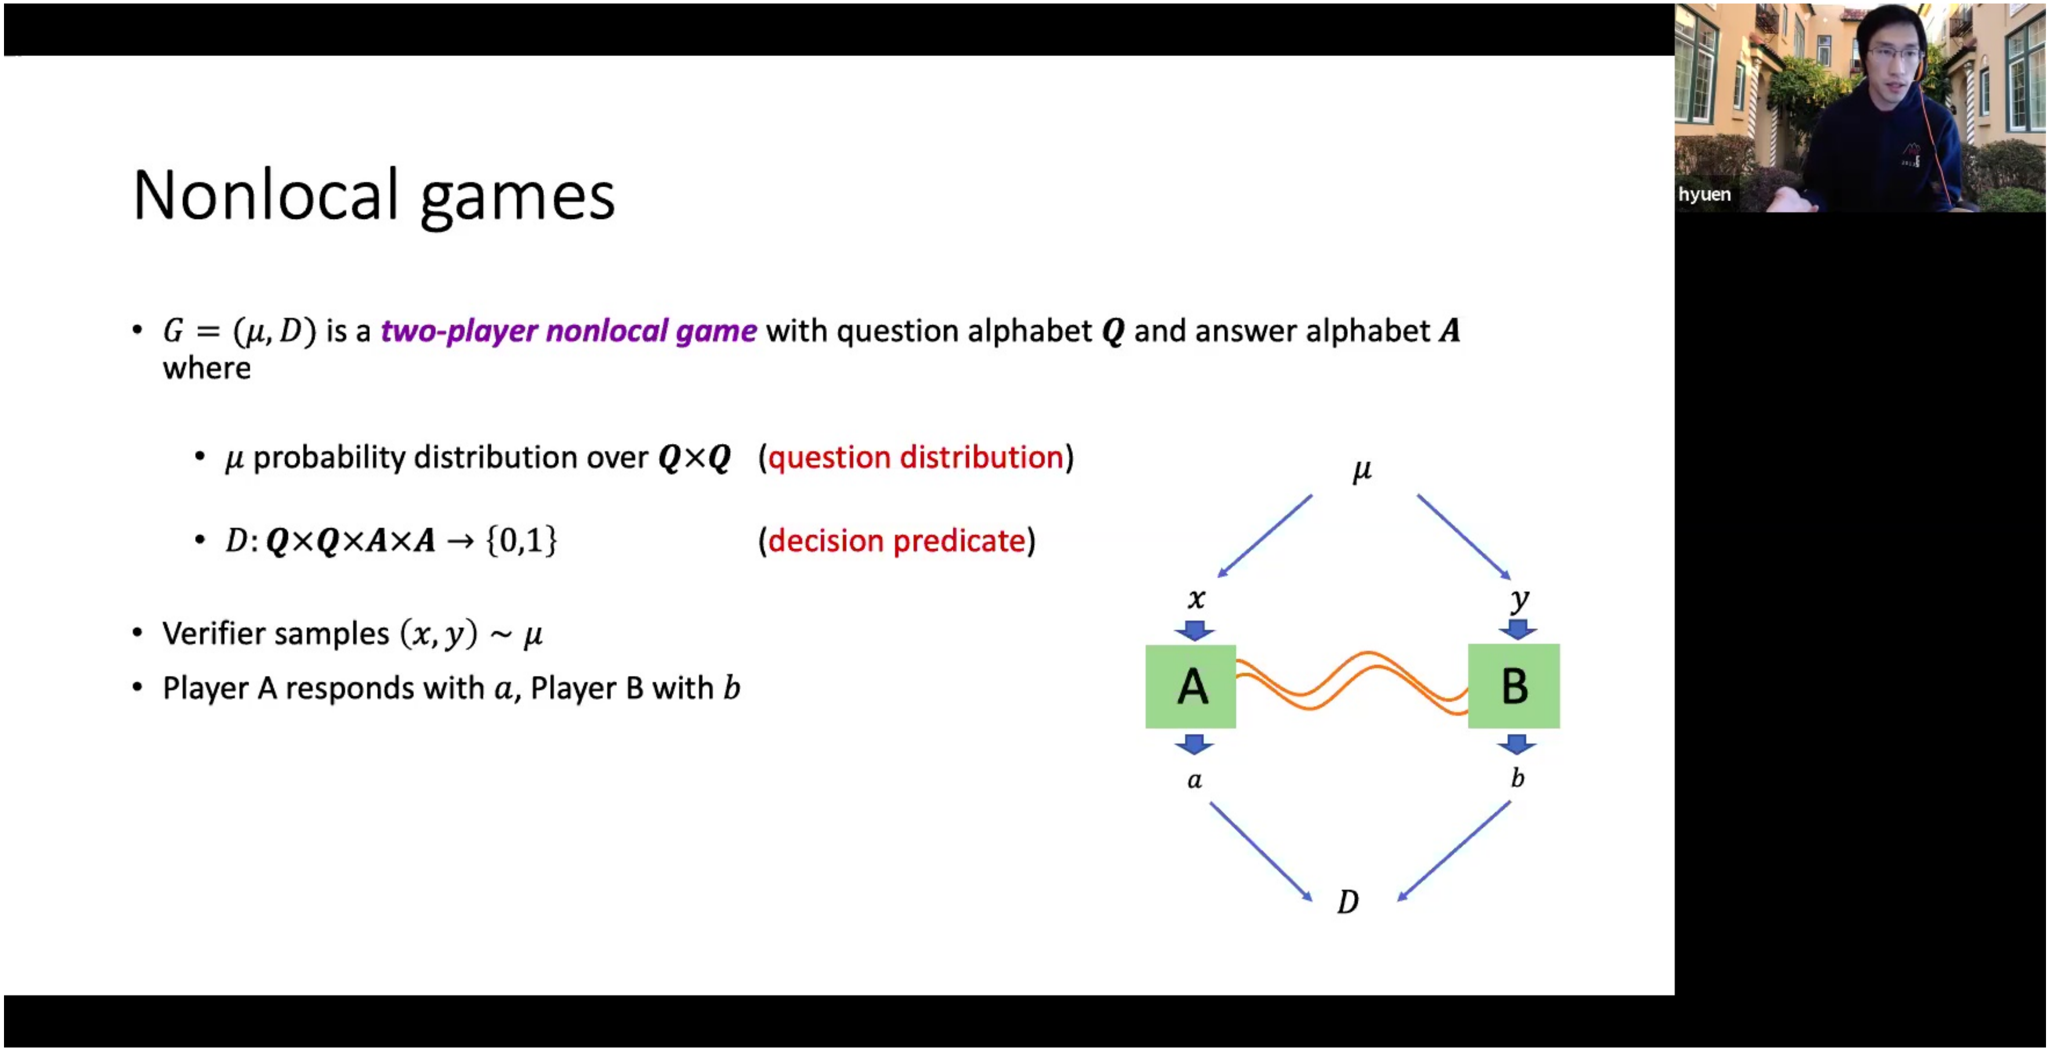

Nonlocal games

- $G = (\mu, D)$ is a ***two-player nonlocal game*** with question alphabet $\boldsymbol{Q}$ and answer alphabet $\boldsymbol{A}$ where
  - μ probability distribution over $\boldsymbol{Q} \times \boldsymbol{Q}$ (question distribution)
  - $D: \boldsymbol{Q} \times \boldsymbol{Q} \times \boldsymbol{A} \times \boldsymbol{A} \rightarrow \{0,1\}$ (decision predicate)
- Verifier samples $(x, y) \sim \mu$
- Player A responds with a, Player B with b

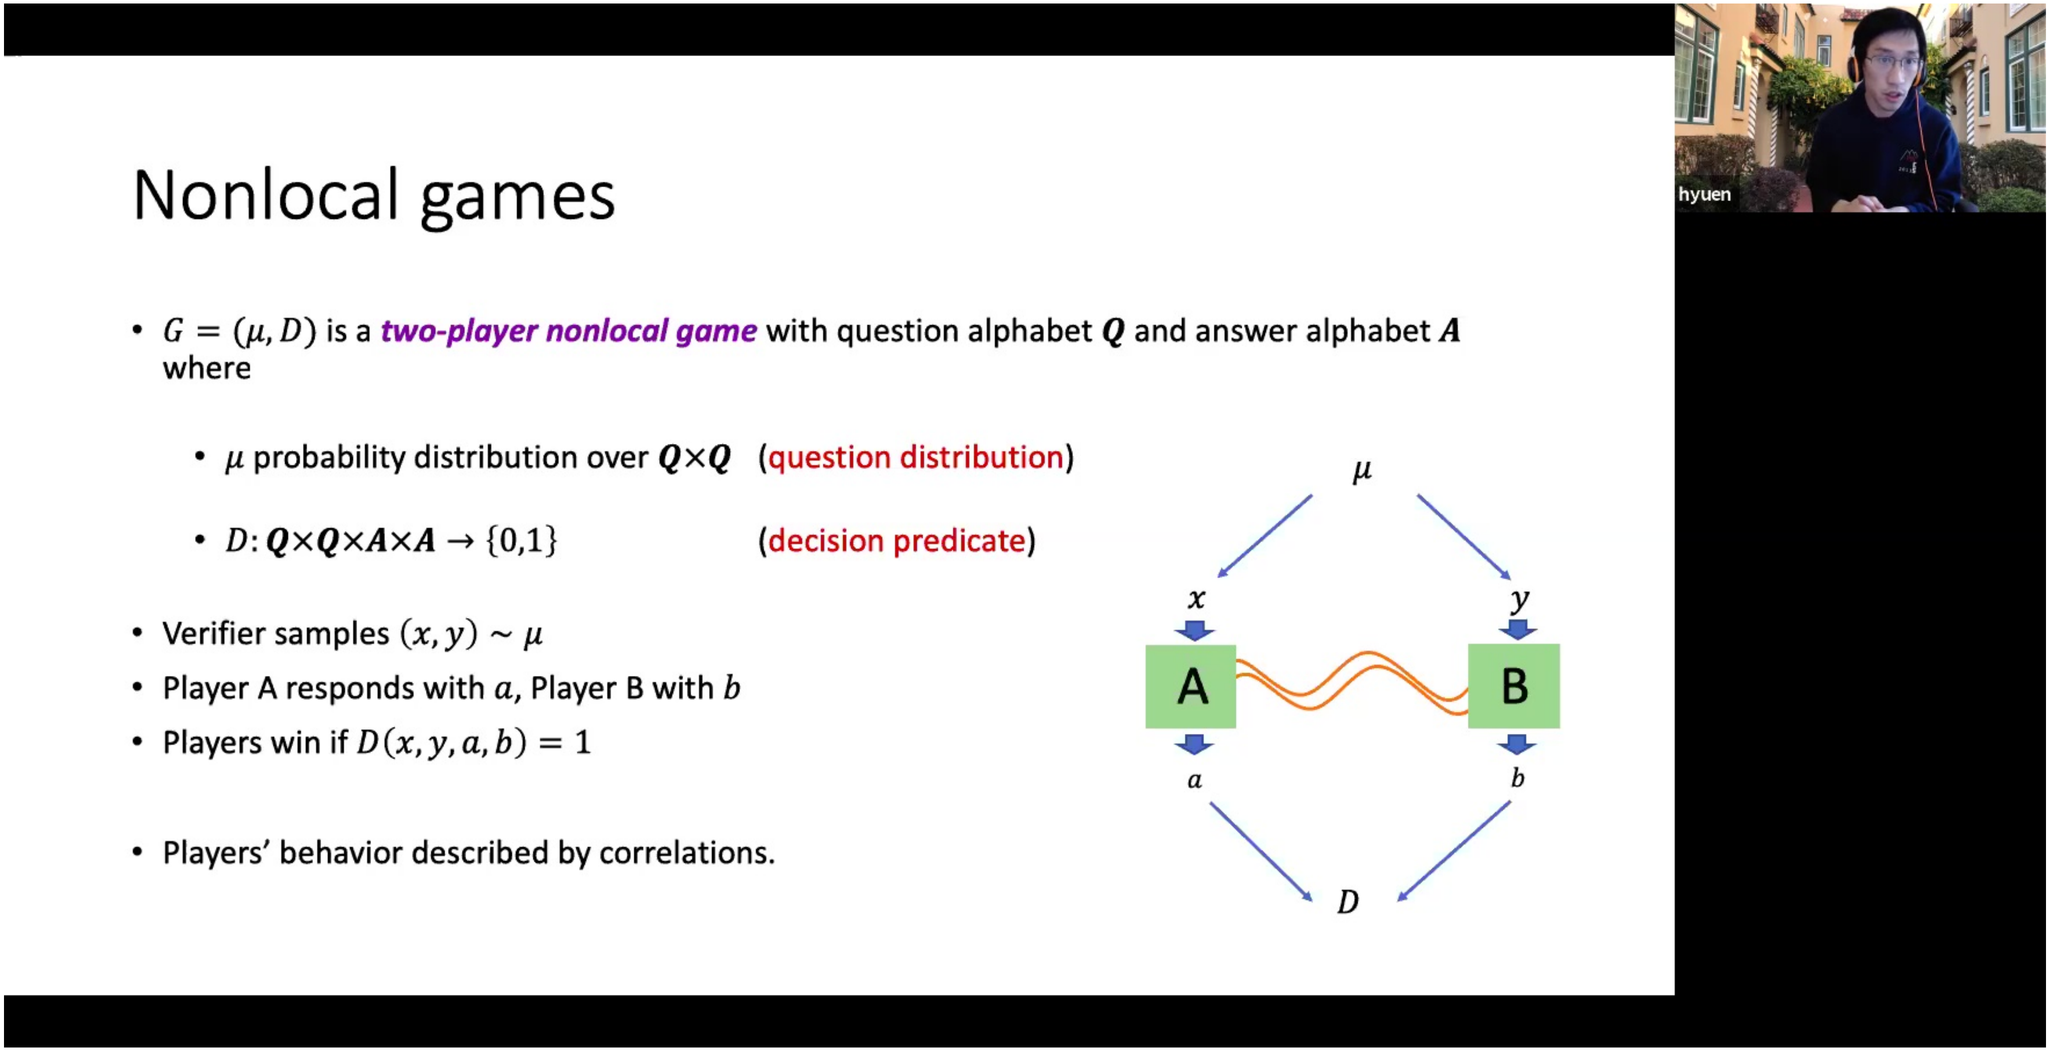

Nonlocal games

- $G = (\mu, D)$ is a ***two-player nonlocal game*** with question alphabet $\boldsymbol{Q}$ and answer alphabet $\boldsymbol{A}$ where
  - μ probability distribution over $\boldsymbol{Q}\times\boldsymbol{Q}$ (question distribution)
  - $D: \boldsymbol{Q}\times\boldsymbol{Q}\times\boldsymbol{A}\times\boldsymbol{A} \to \{0,1\}$ (decision predicate)
- Verifier samples $(x, y) \sim \mu$
- Player A responds with a, Player B with b
- Players win if $D(x, y, a, b) = 1$
- Players' behavior described by correlations.

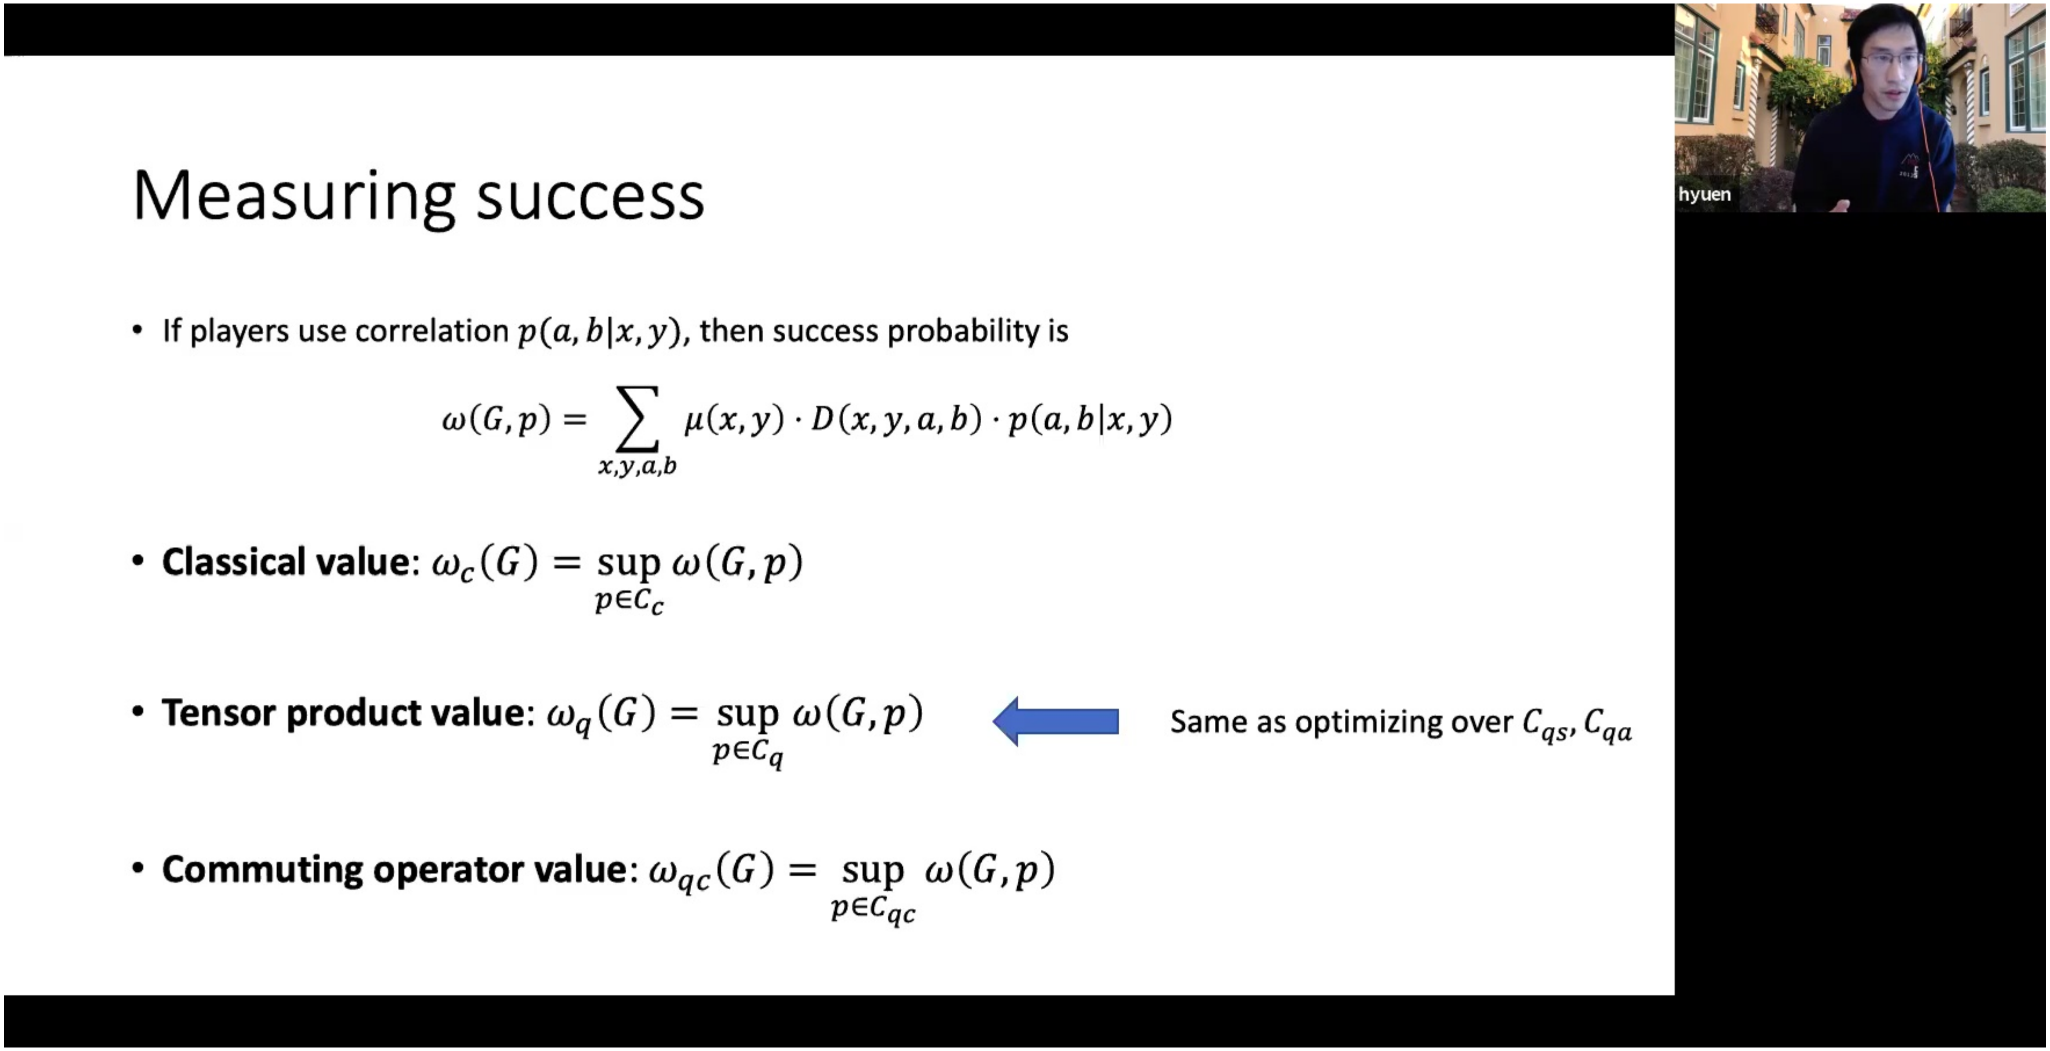

Measuring success

- If players use correlation $p(a,b|x,y)$, then success probability is

$$\omega(G,p) = \sum_{x,y,a,b} \mu(x,y) \cdot D(x,y,a,b) \cdot p(a,b|x,y)$$

- **Classical value**: $\omega_c(G) = \sup_{p \in C_c} \omega(G,p)$

- **Tensor product value**: $\omega_q(G) = \sup_{p \in C_q} \omega(G,p)$

Same as optimizing over C_{qs}, C_{qa}

- **Commuting operator value**: $\omega_{qc}(G) = \sup_{p \in C_{qc}} \omega(G,p)$

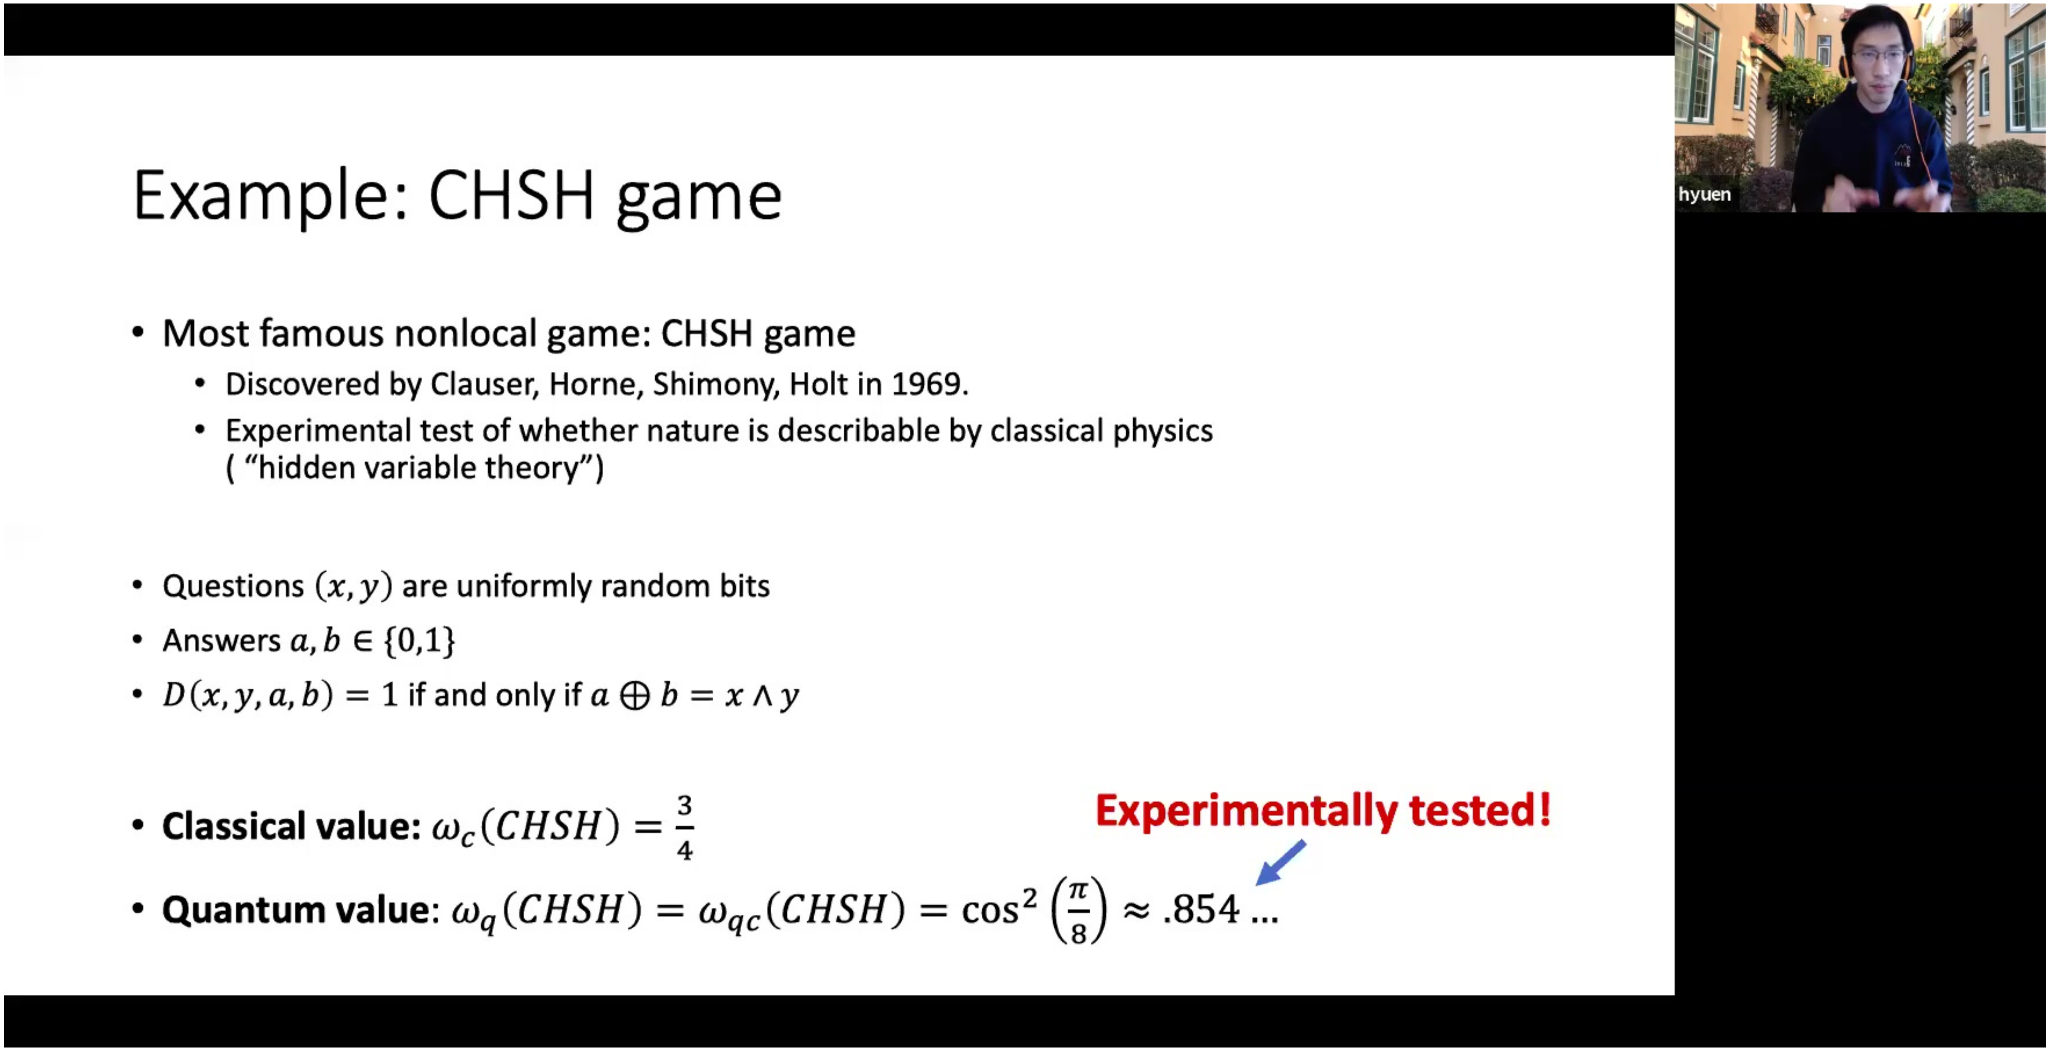

Example: CHSH game

- Most famous nonlocal game: CHSH game
  - Discovered by Clauser, Horne, Shimony, Holt in 1969.
  - Experimental test of whether nature is describable by classical physics (“hidden variable theory”)

- Questions (x, y) are uniformly random bits
- Answers $a, b \in \{0,1\}$
- $D(x, y, a, b) = 1$ if and only if $a \oplus b = x \wedge y$

- **Classical value:** $\omega_c(CHSH) = \frac{3}{4}$
- **Quantum value**: $\omega_q(CHSH) = \omega_{qc}(CHSH) = \cos^2\left(\frac{\pi}{8}\right) \approx .854 \ldots$

Experimentally tested!

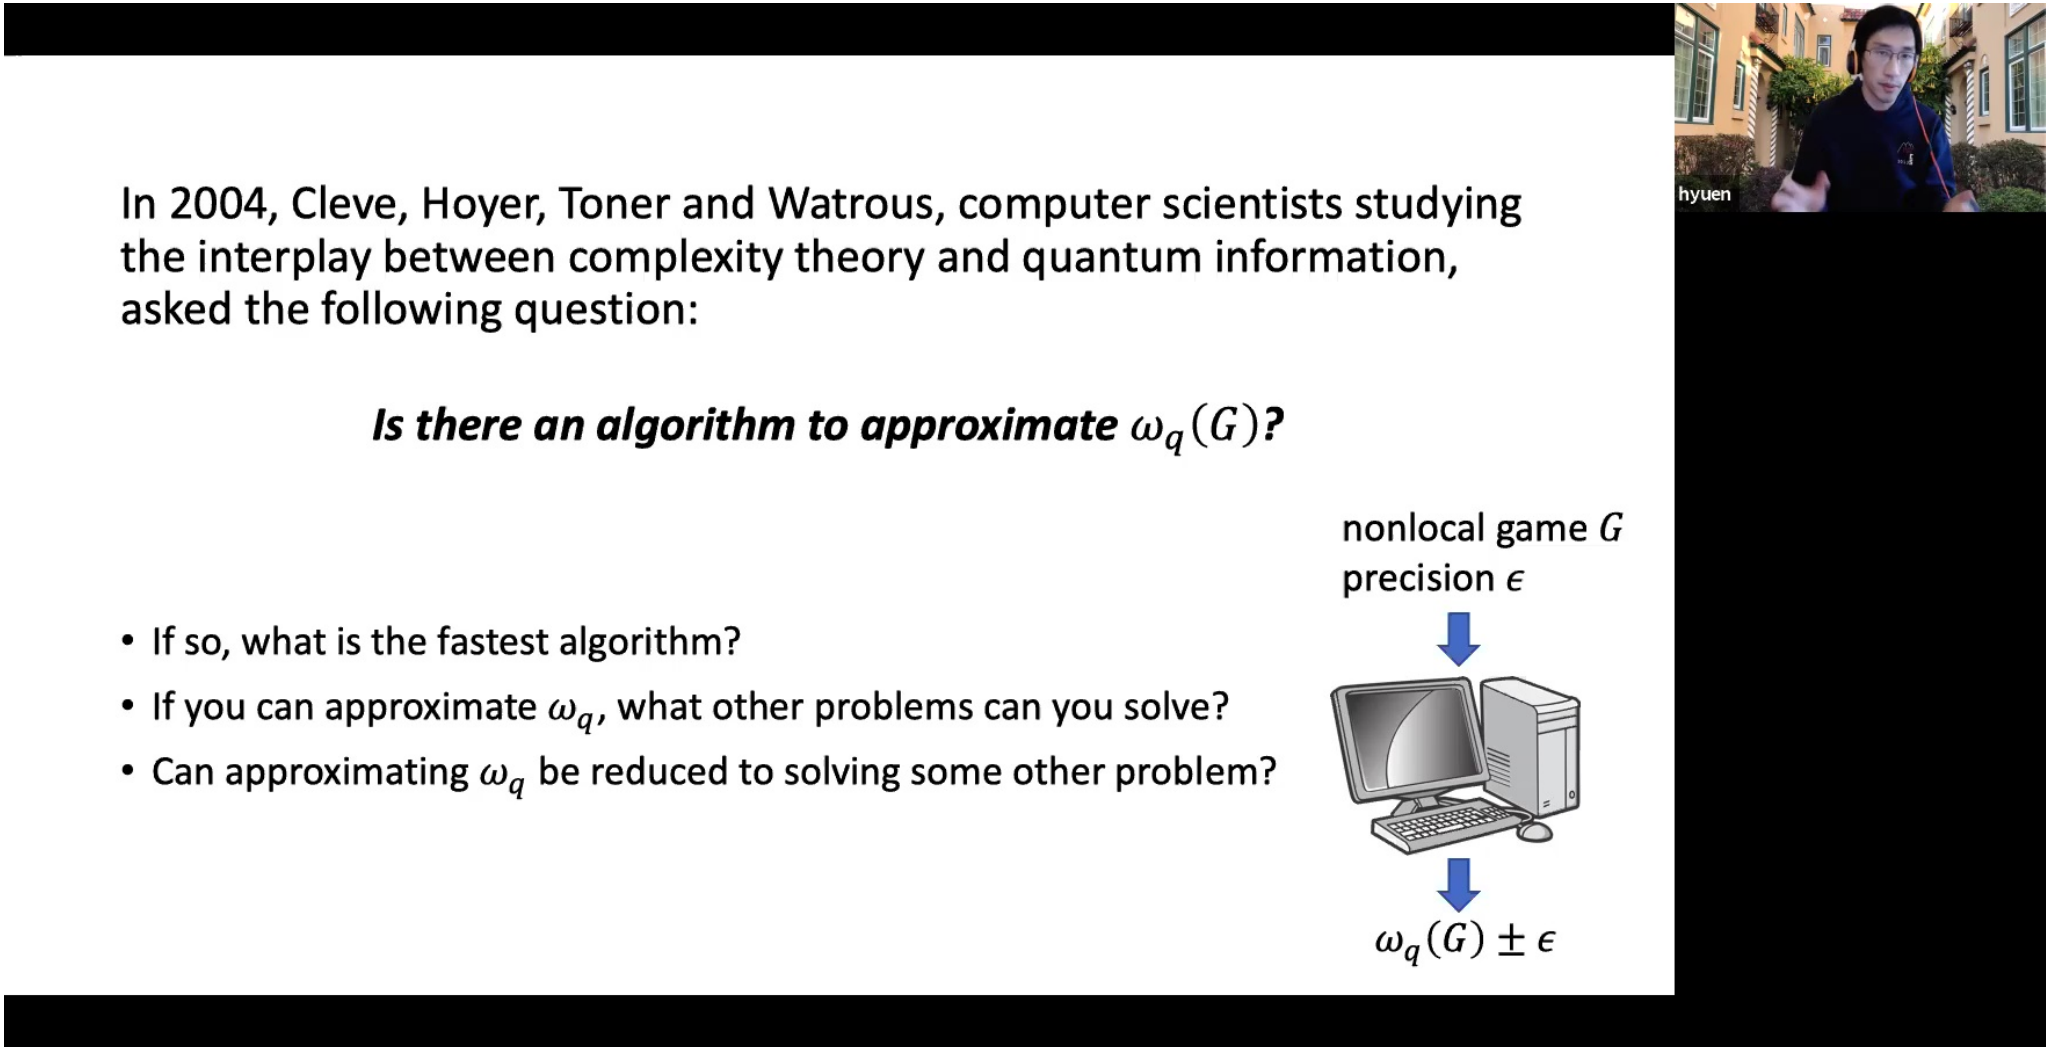

In 2004, Cleve, Hoyer, Toner and Watrous, computer scientists studying the interplay between complexity theory and quantum information, asked the following question:

Is there an algorithm to approximate $\omega_q(G)$?

nonlocal game G
precision ϵ

$\omega_q(G) \pm \epsilon$

- If so, what is the fastest algorithm?
- If you can approximate ω_q, what other problems can you solve?
- Can approximating ω_q be reduced to solving some other problem?

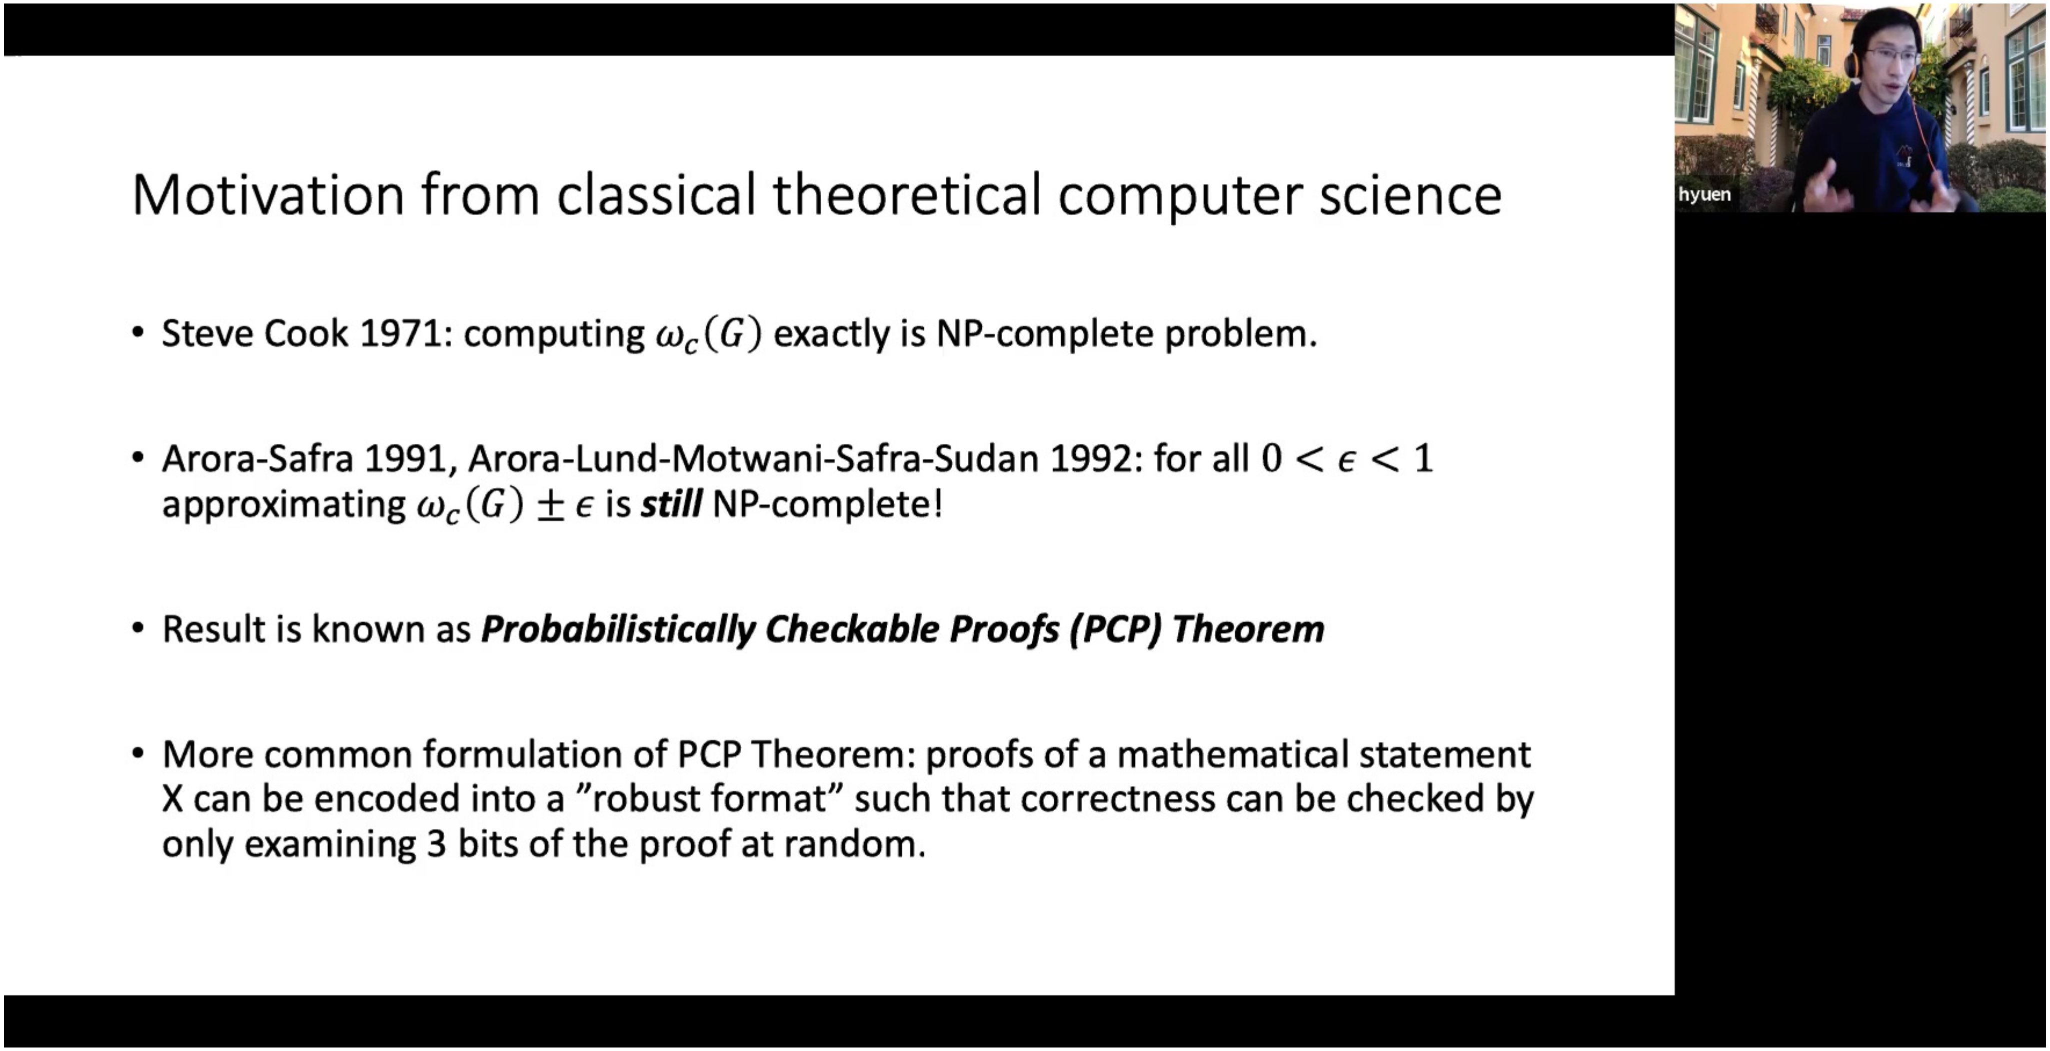

Motivation from classical theoretical computer science

- Steve Cook 1971: computing $\omega_c(G)$ exactly is NP-complete problem.

- Arora-Safra 1991, Arora-Lund-Motwani-Safra-Sudan 1992: for all $0 < \epsilon < 1$ approximating $\omega_c(G) \pm \epsilon$ is ***still*** NP-complete!

- Result is known as ***Probabilistically Checkable Proofs (PCP) Theorem***

- More common formulation of PCP Theorem: proofs of a mathematical statement X can be encoded into a ”robust format” such that correctness can be checked by only examining 3 bits of the proof at random.

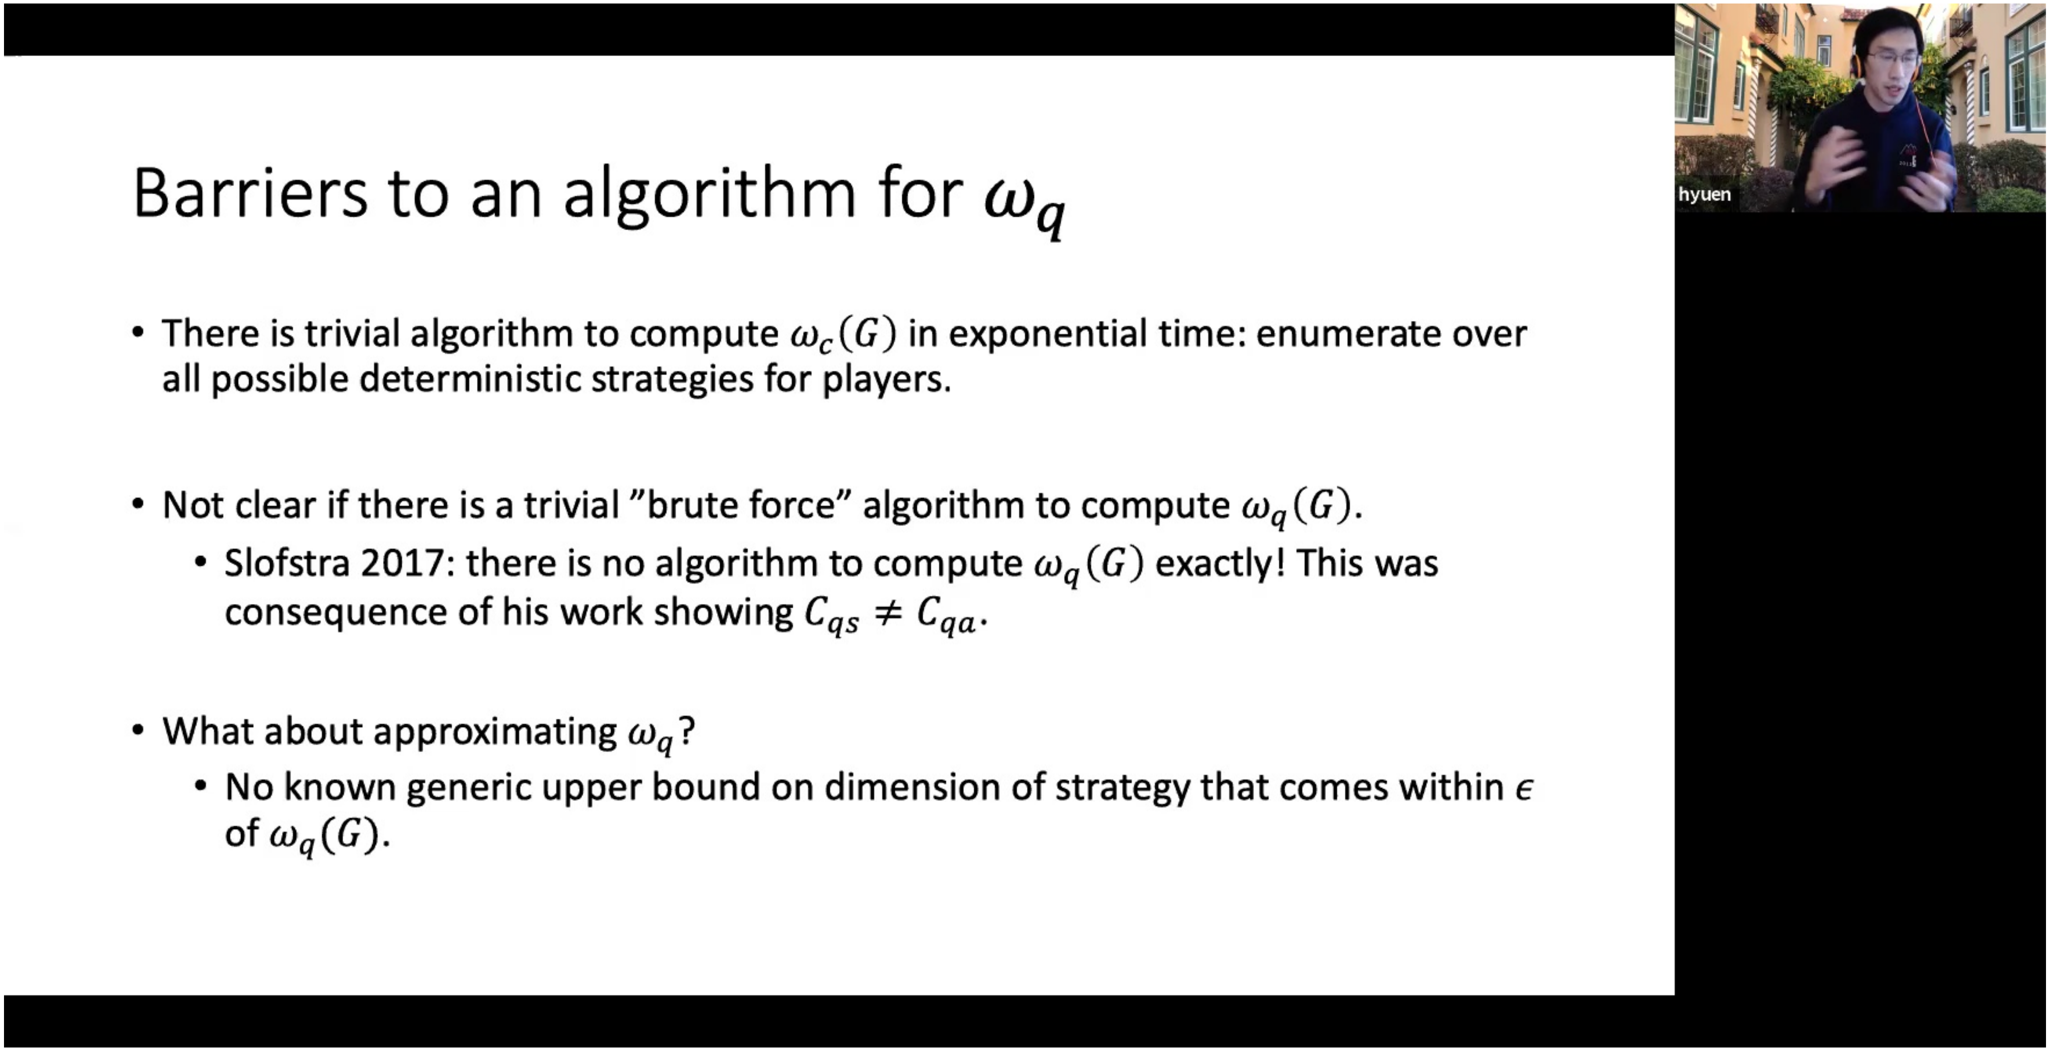

Barriers to an algorithm for ω_q

- There is trivial algorithm to compute $\omega_c(G)$ in exponential time: enumerate over all possible deterministic strategies for players.

- Not clear if there is a trivial ”brute force” algorithm to compute $\omega_q(G)$.
  - Slofstra 2017: there is no algorithm to compute $\omega_q(G)$ exactly! This was consequence of his work showing $C_{qs} \neq C_{qa}$.

- What about approximating ω_q?
  - No known generic upper bound on dimension of strategy that comes within ϵ of $\omega_q(G)$.

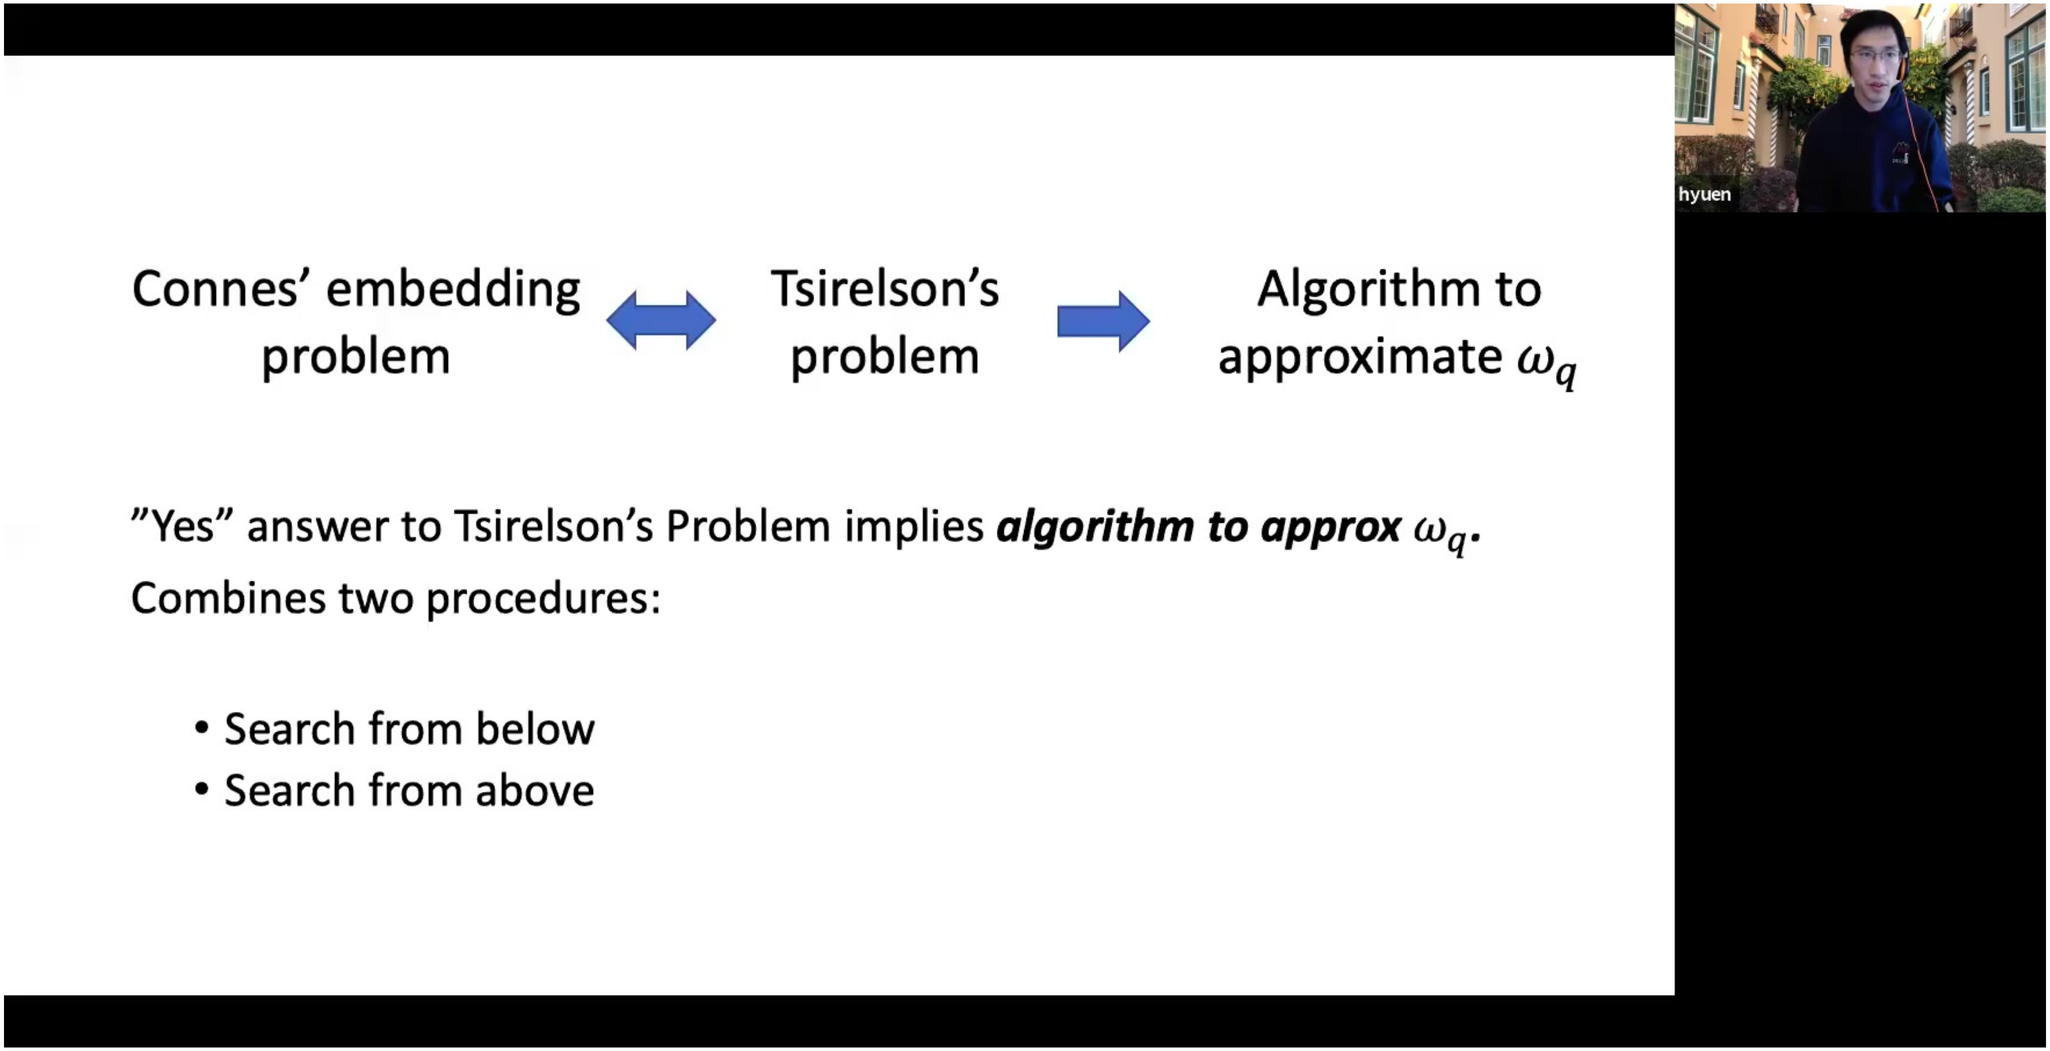

Connes' embedding problem ↔ Tsirelson's problem → Algorithm to approximate ω_q

”Yes” answer to Tsirelson's Problem implies ***algorithm to approx*** ω_q.

Combines two procedures:

- Search from below
- Search from above

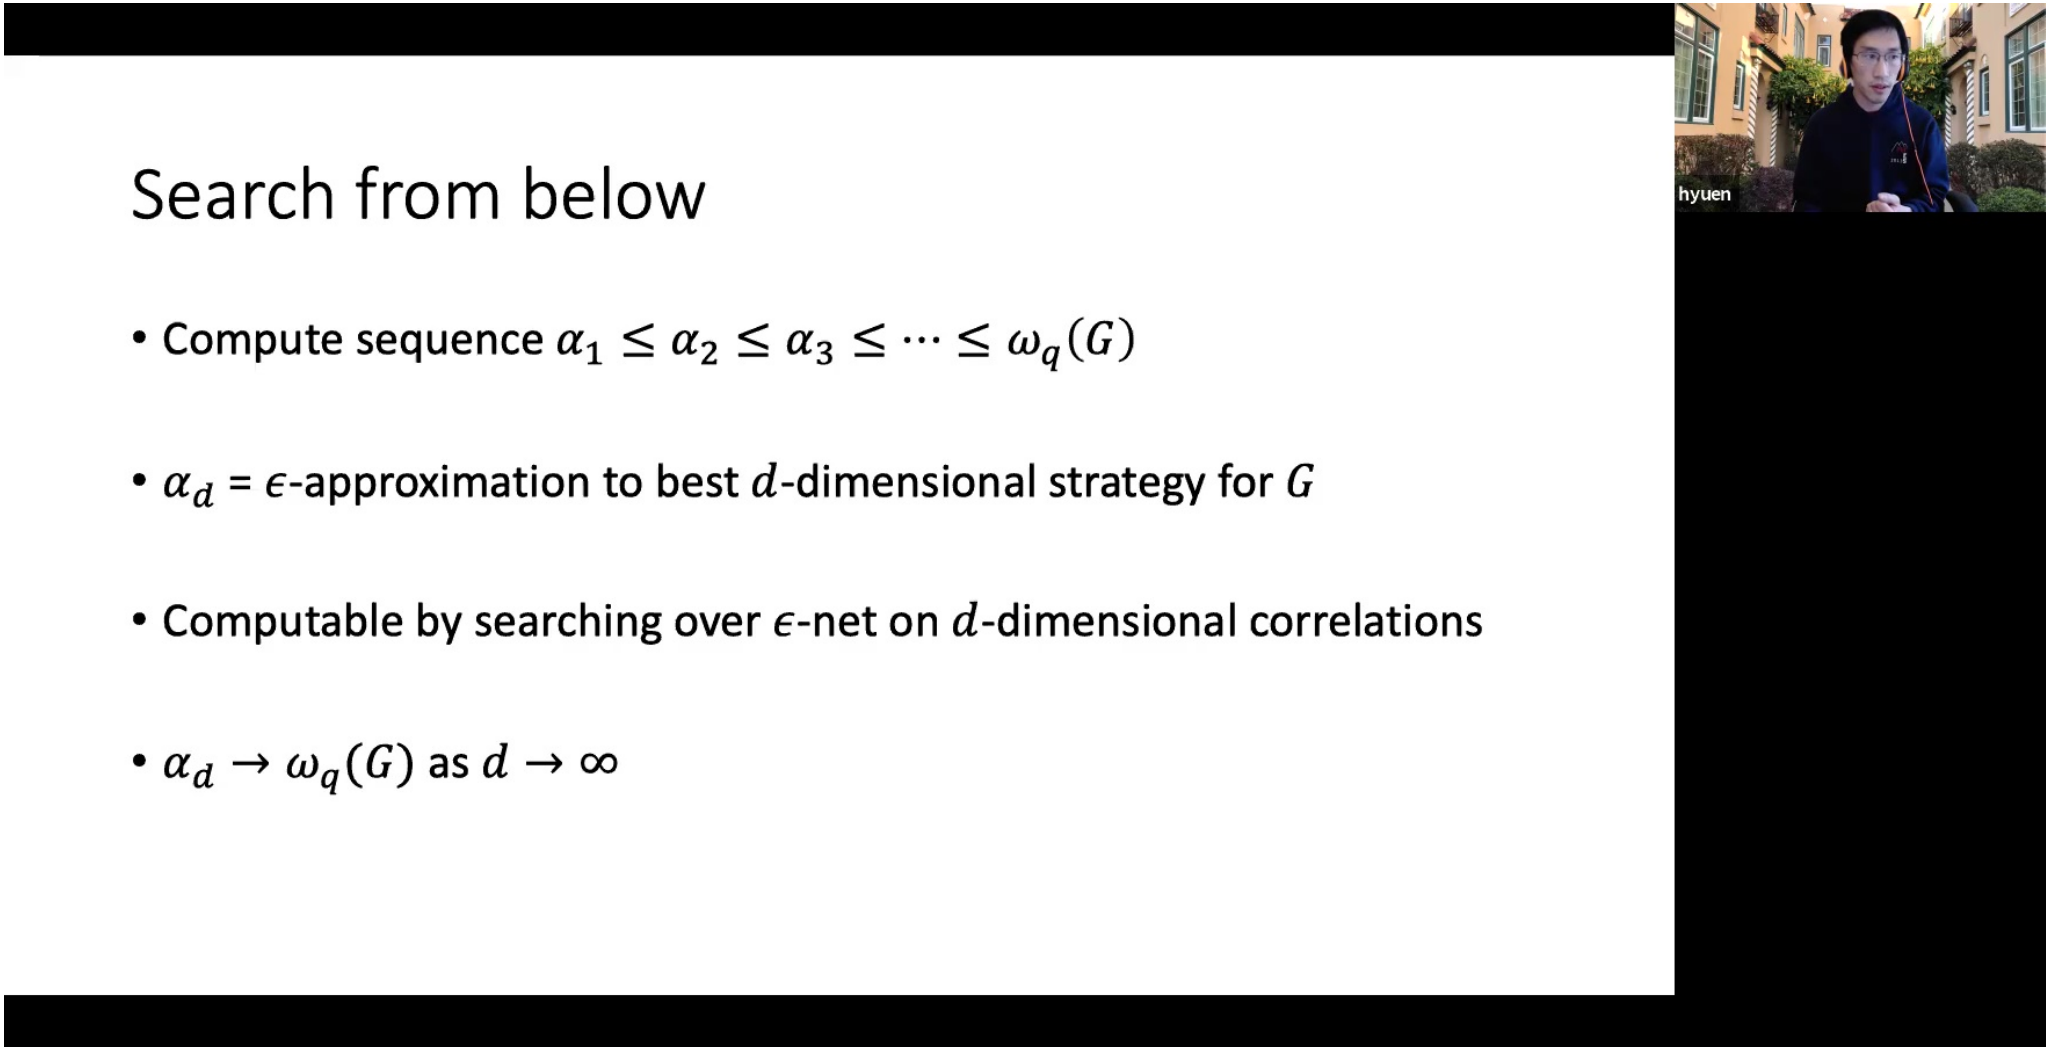

Search from below

- Compute sequence $\alpha_1 \leq \alpha_2 \leq \alpha_3 \leq \cdots \leq \omega_q(G)$
- α_d = ϵ-approximation to best d-dimensional strategy for G
- Computable by searching over ϵ-net on d-dimensional correlations
- $\alpha_d \to \omega_q(G)$ as $d \to \infty$

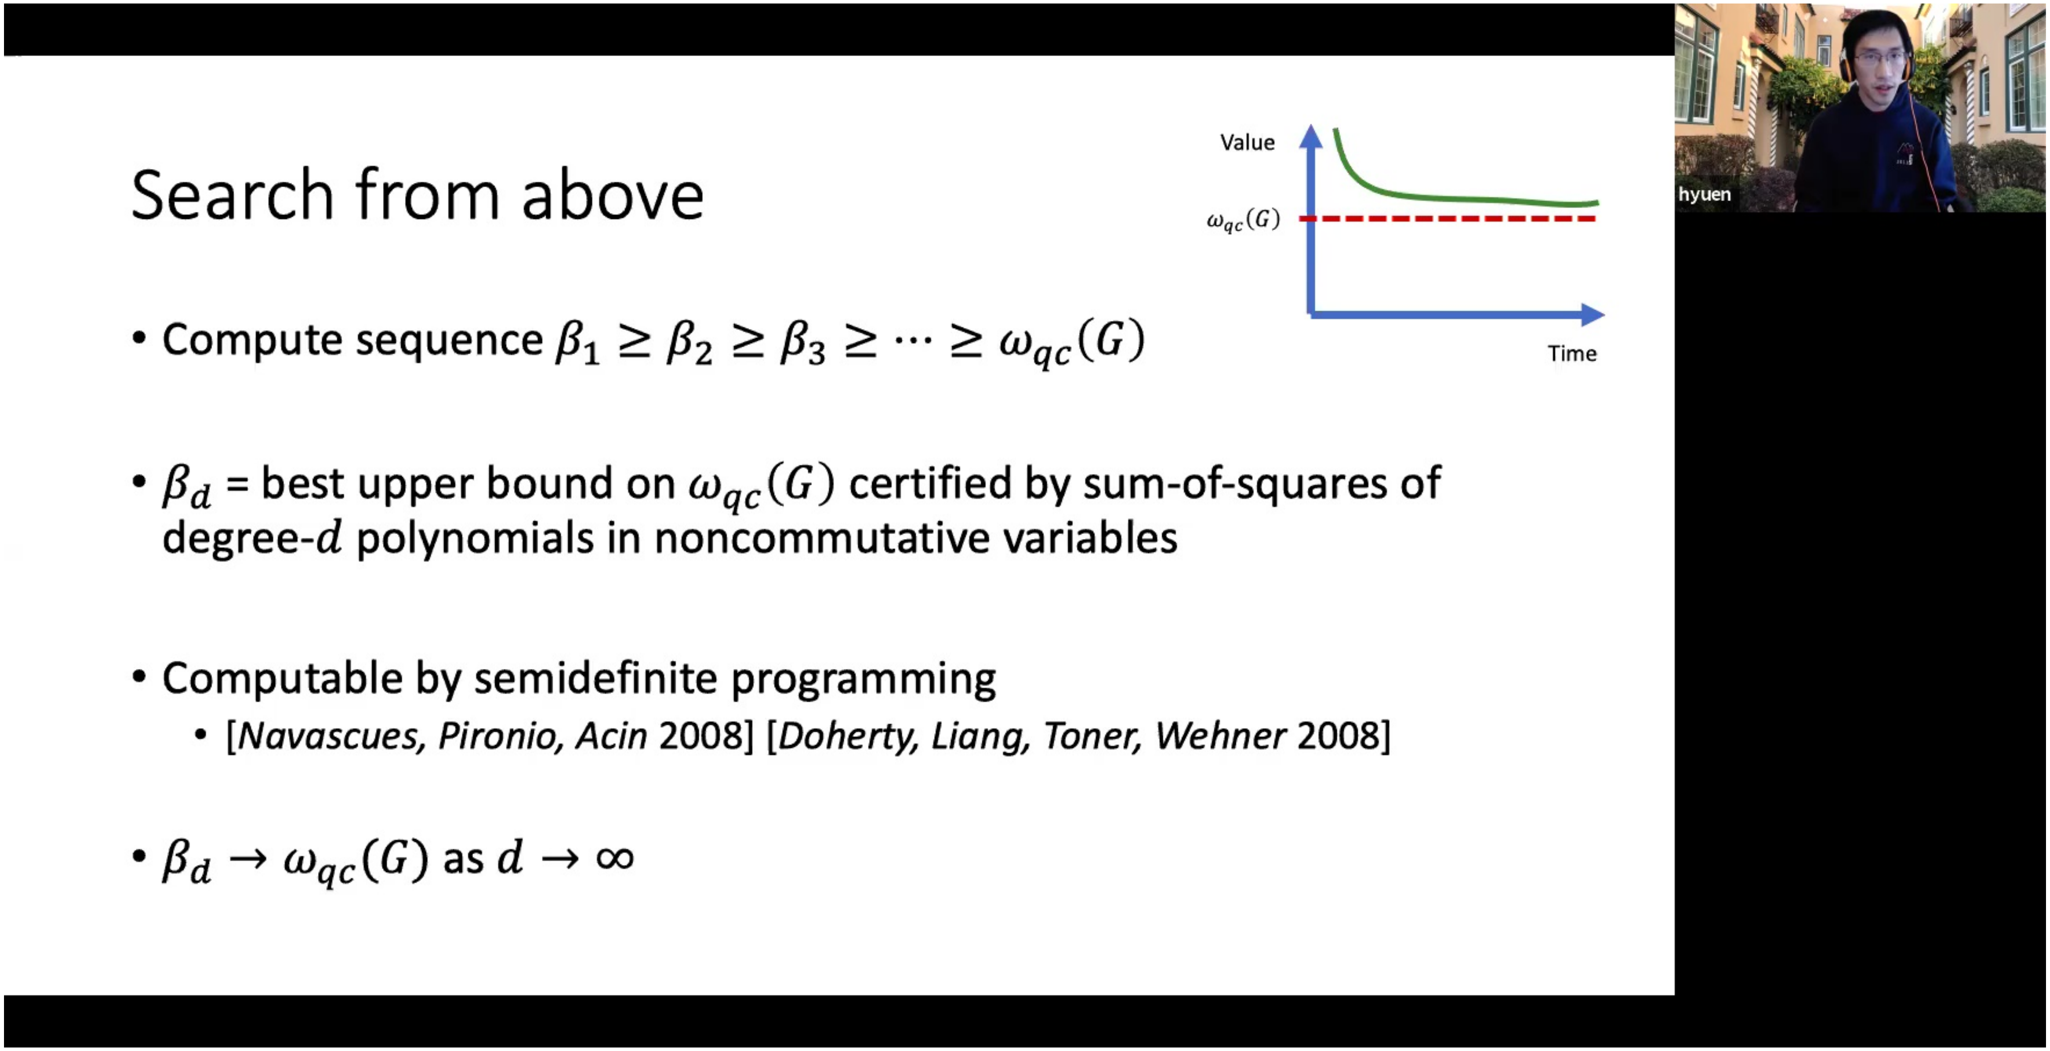

Search from above

- Compute sequence $\beta_1 \geq \beta_2 \geq \beta_3 \geq \cdots \geq \omega_{qc}(G)$

- β_d = best upper bound on $\omega_{qc}(G)$ certified by sum-of-squares of degree-d polynomials in noncommutative variables

- Computable by semidefinite programming
  - [*Navascues, Pironio, Acin* 2008] [*Doherty, Liang, Toner, Wehner* 2008]

- $\beta_d \to \omega_{qc}(G)$ as $d \to \infty$

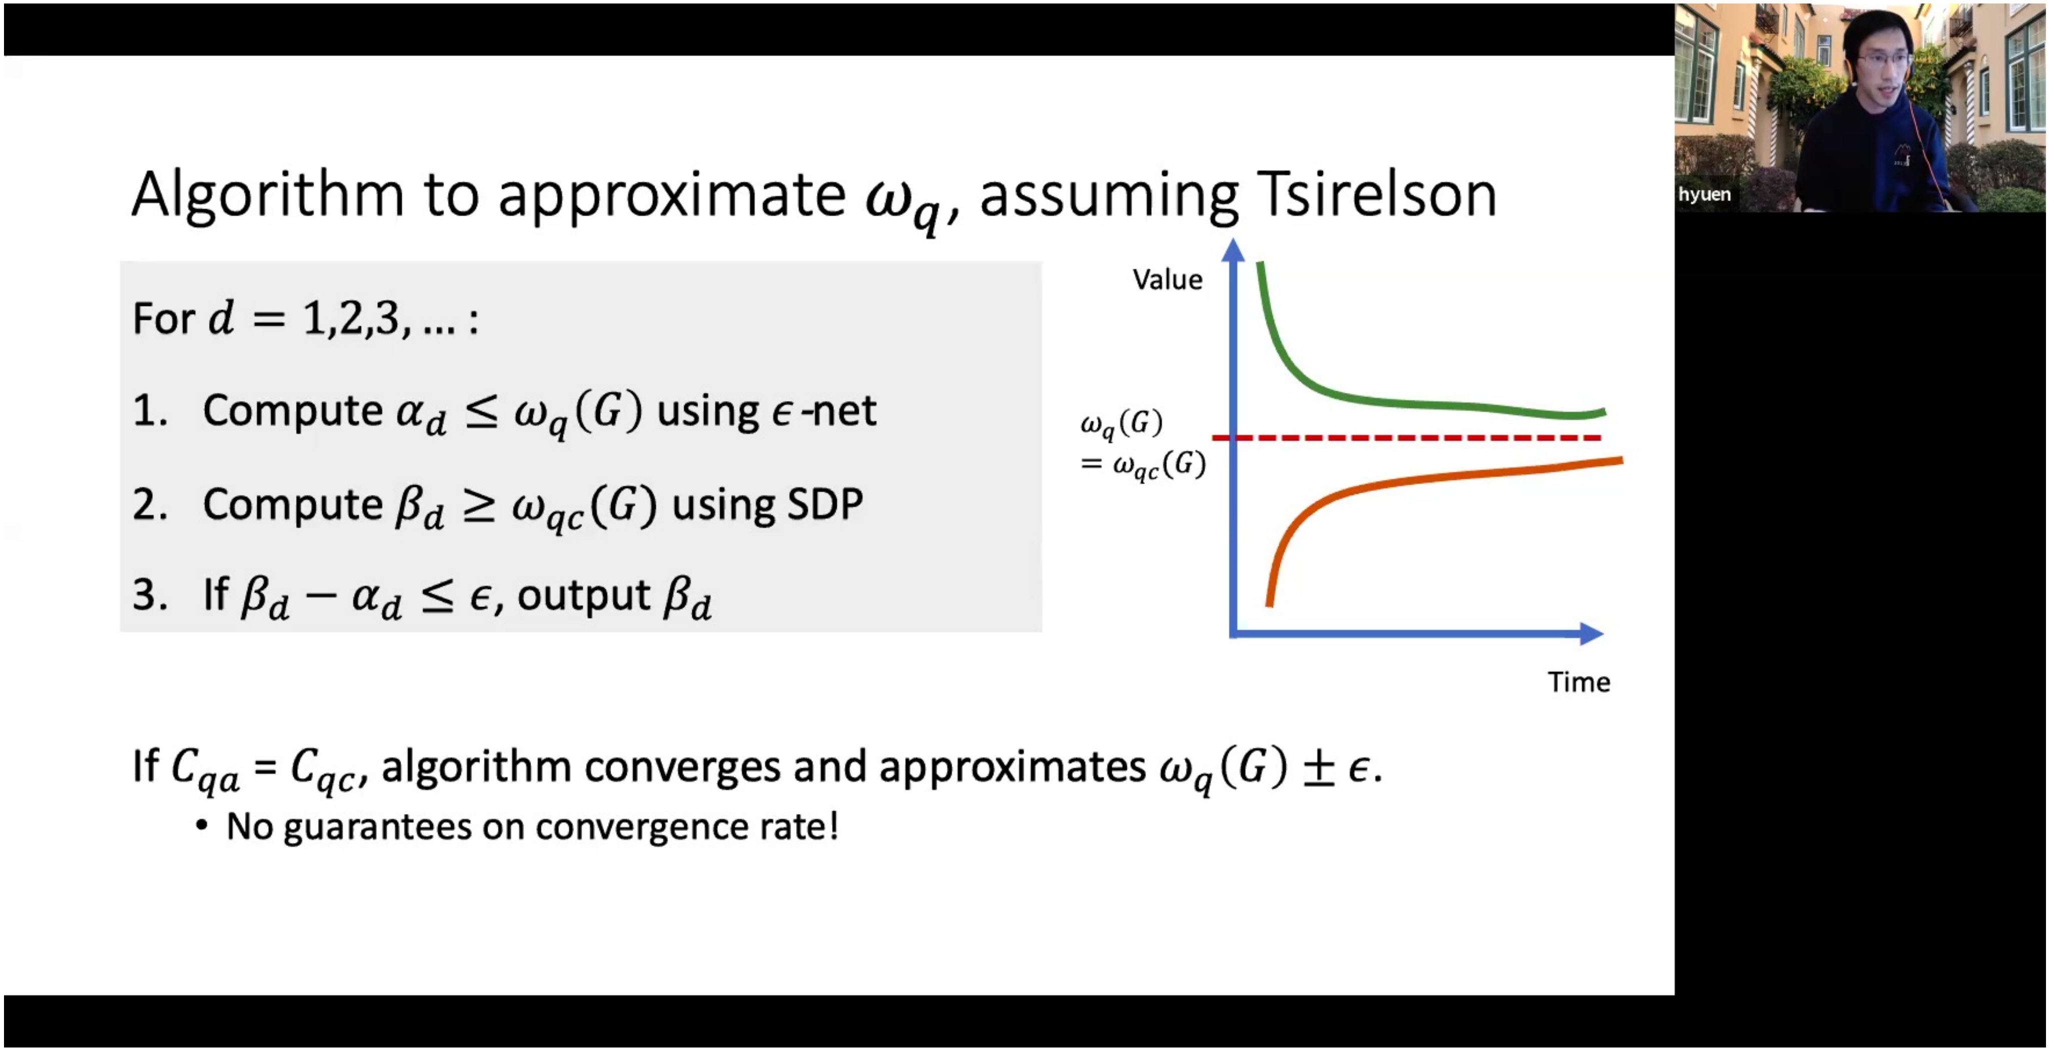

Algorithm to approximate ω_q, assuming Tsirelson

```
For $d = 1,2,3,\ldots$ :
1. Compute $\alpha_d \leq \omega_q(G)$ using $\epsilon$-net
2. Compute $\beta_d \geq \omega_{qc}(G)$ using SDP
3. If $\beta_d - \alpha_d \leq \epsilon$, output $\beta_d$
```

If C_{qa} = C_{qc}, algorithm converges and approximates $\omega_q(G) \pm \epsilon$.

- No guarantees on convergence rate!

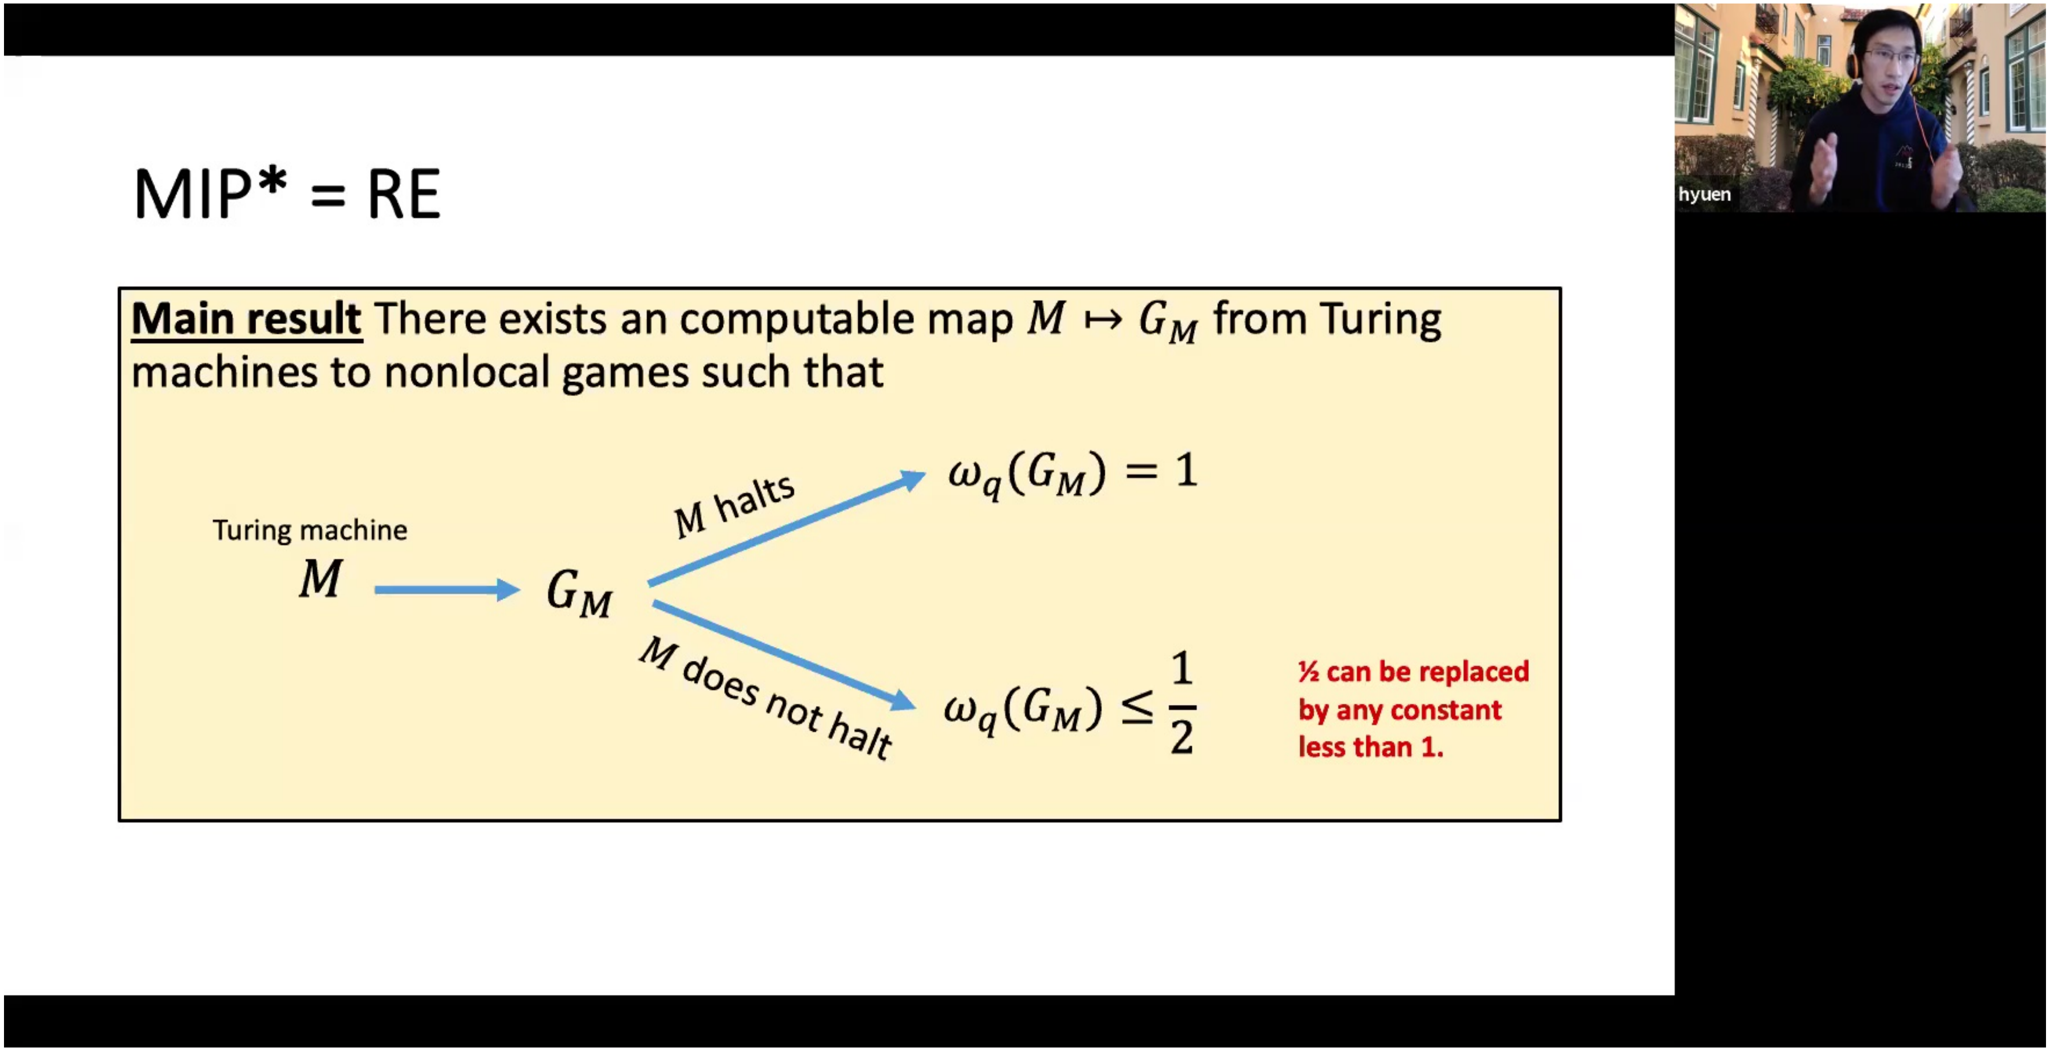

MIP* = RE

<u>Main result</u> There exists an computable map $M \mapsto G_M$ from Turing machines to nonlocal games such that

Turing machine

M

G_M

M halts

$\omega_q(G_M) = 1$

M does not halt

$\omega_q(G_M) \leq \frac{1}{2}$

½ can be replaced by any constant less than 1.

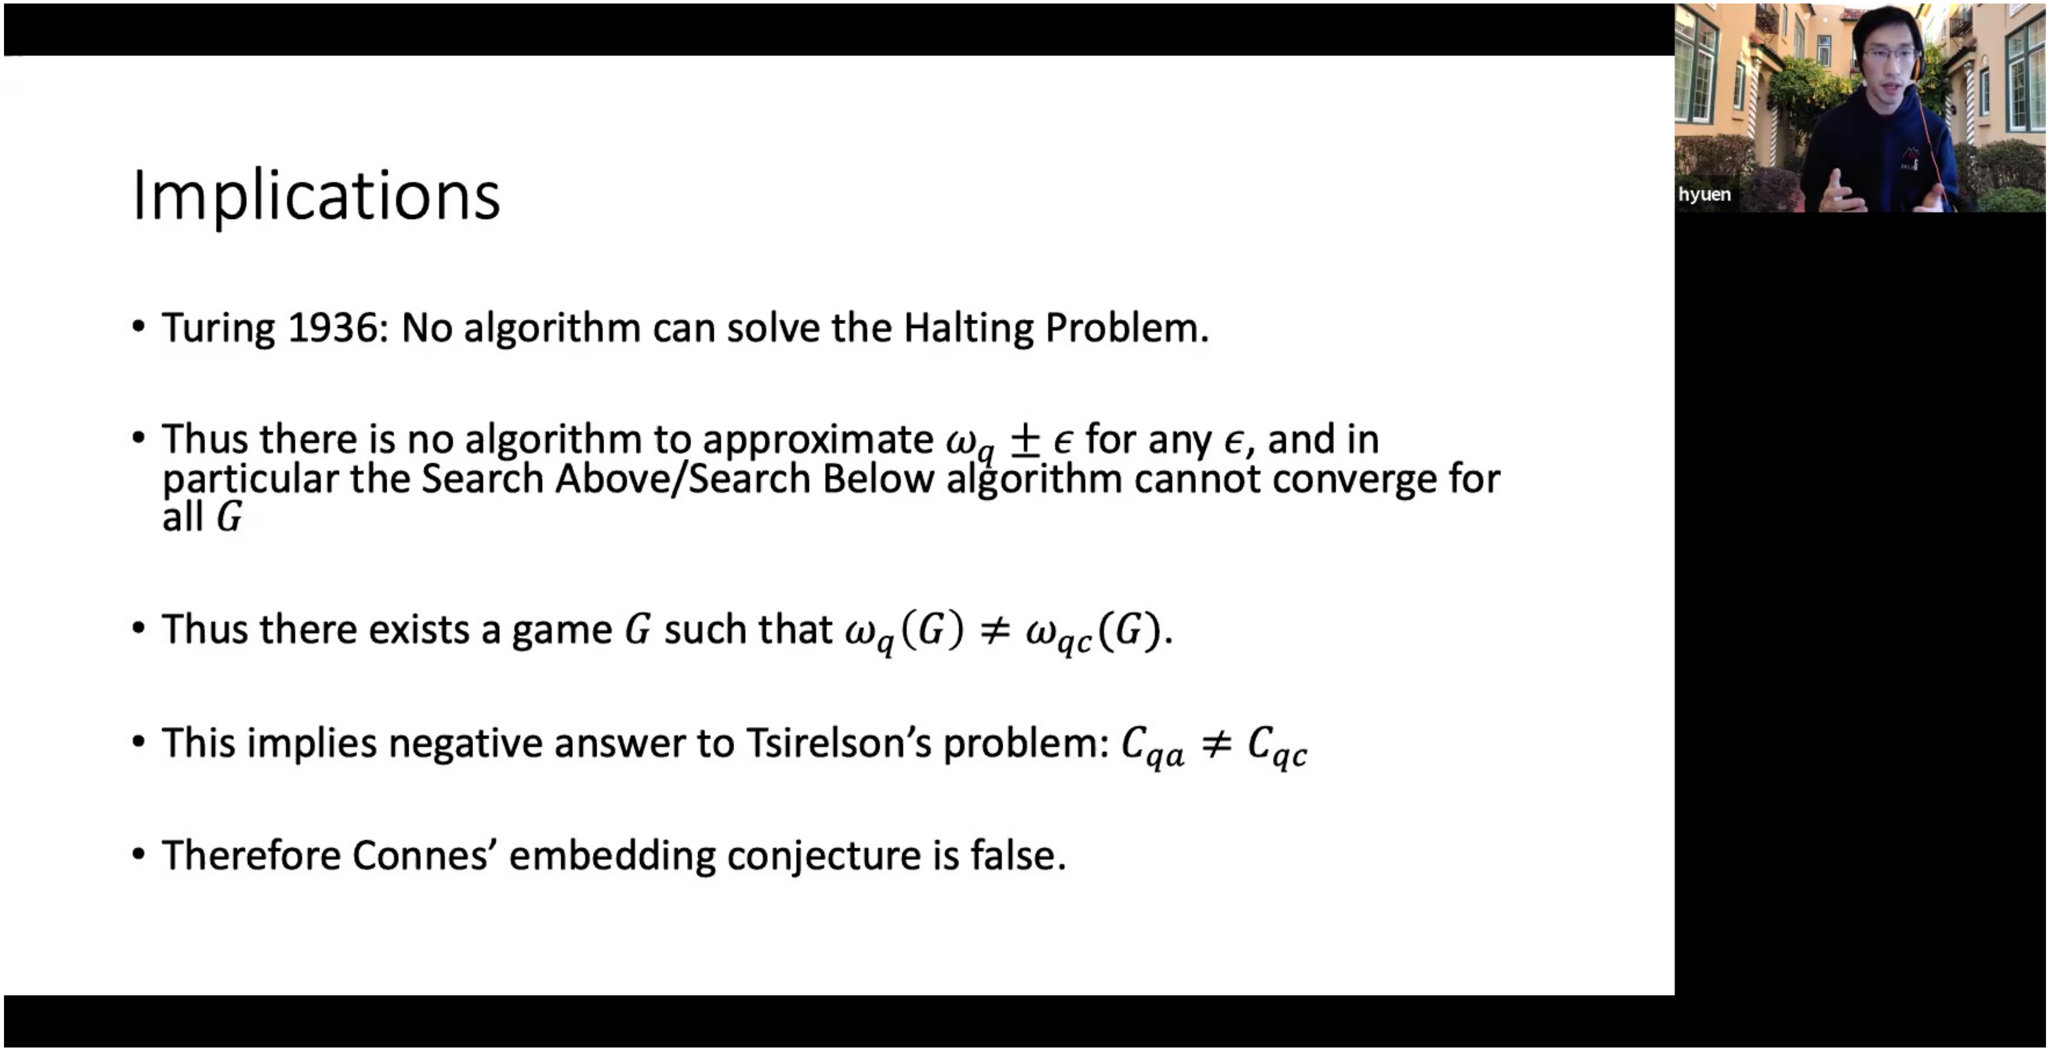

Implications

- Turing 1936: No algorithm can solve the Halting Problem.
- Thus there is no algorithm to approximate $\omega_q \pm \epsilon$ for any ϵ, and in particular the Search Above/Search Below algorithm cannot converge for all G
- Thus there exists a game G such that $\omega_q(G) \neq \omega_{qc}(G)$.
- This implies negative answer to Tsirelson's problem: $C_{qa} \neq C_{qc}$
- Therefore Connes' embedding conjecture is false.

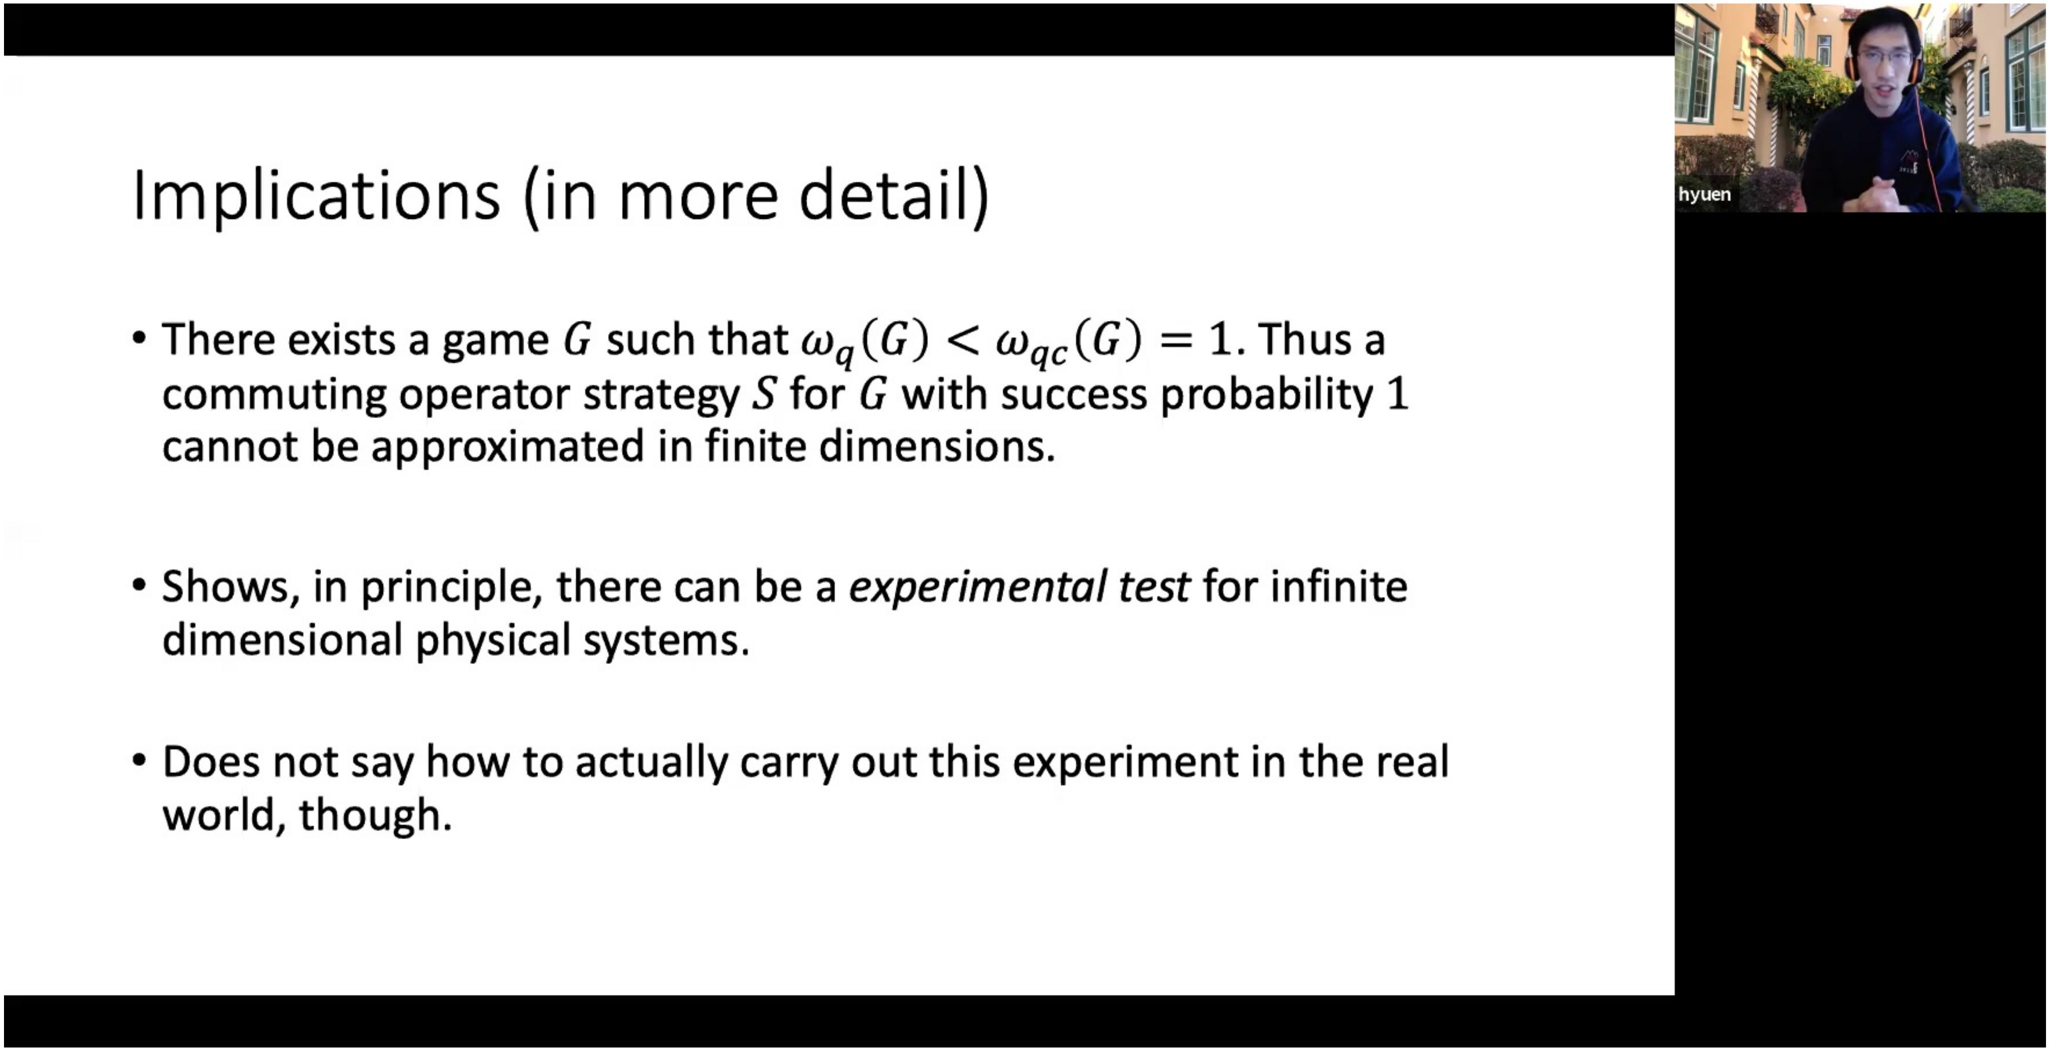

Implications (in more detail)

- There exists a game G such that $\omega_q(G) < \omega_{qc}(G) = 1$. Thus a commuting operator strategy S for G with success probability 1 cannot be approximated in finite dimensions.
- Shows, in principle, there can be a *experimental test* for infinite dimensional physical systems.
- Does not say how to actually carry out this experiment in the real world, though.

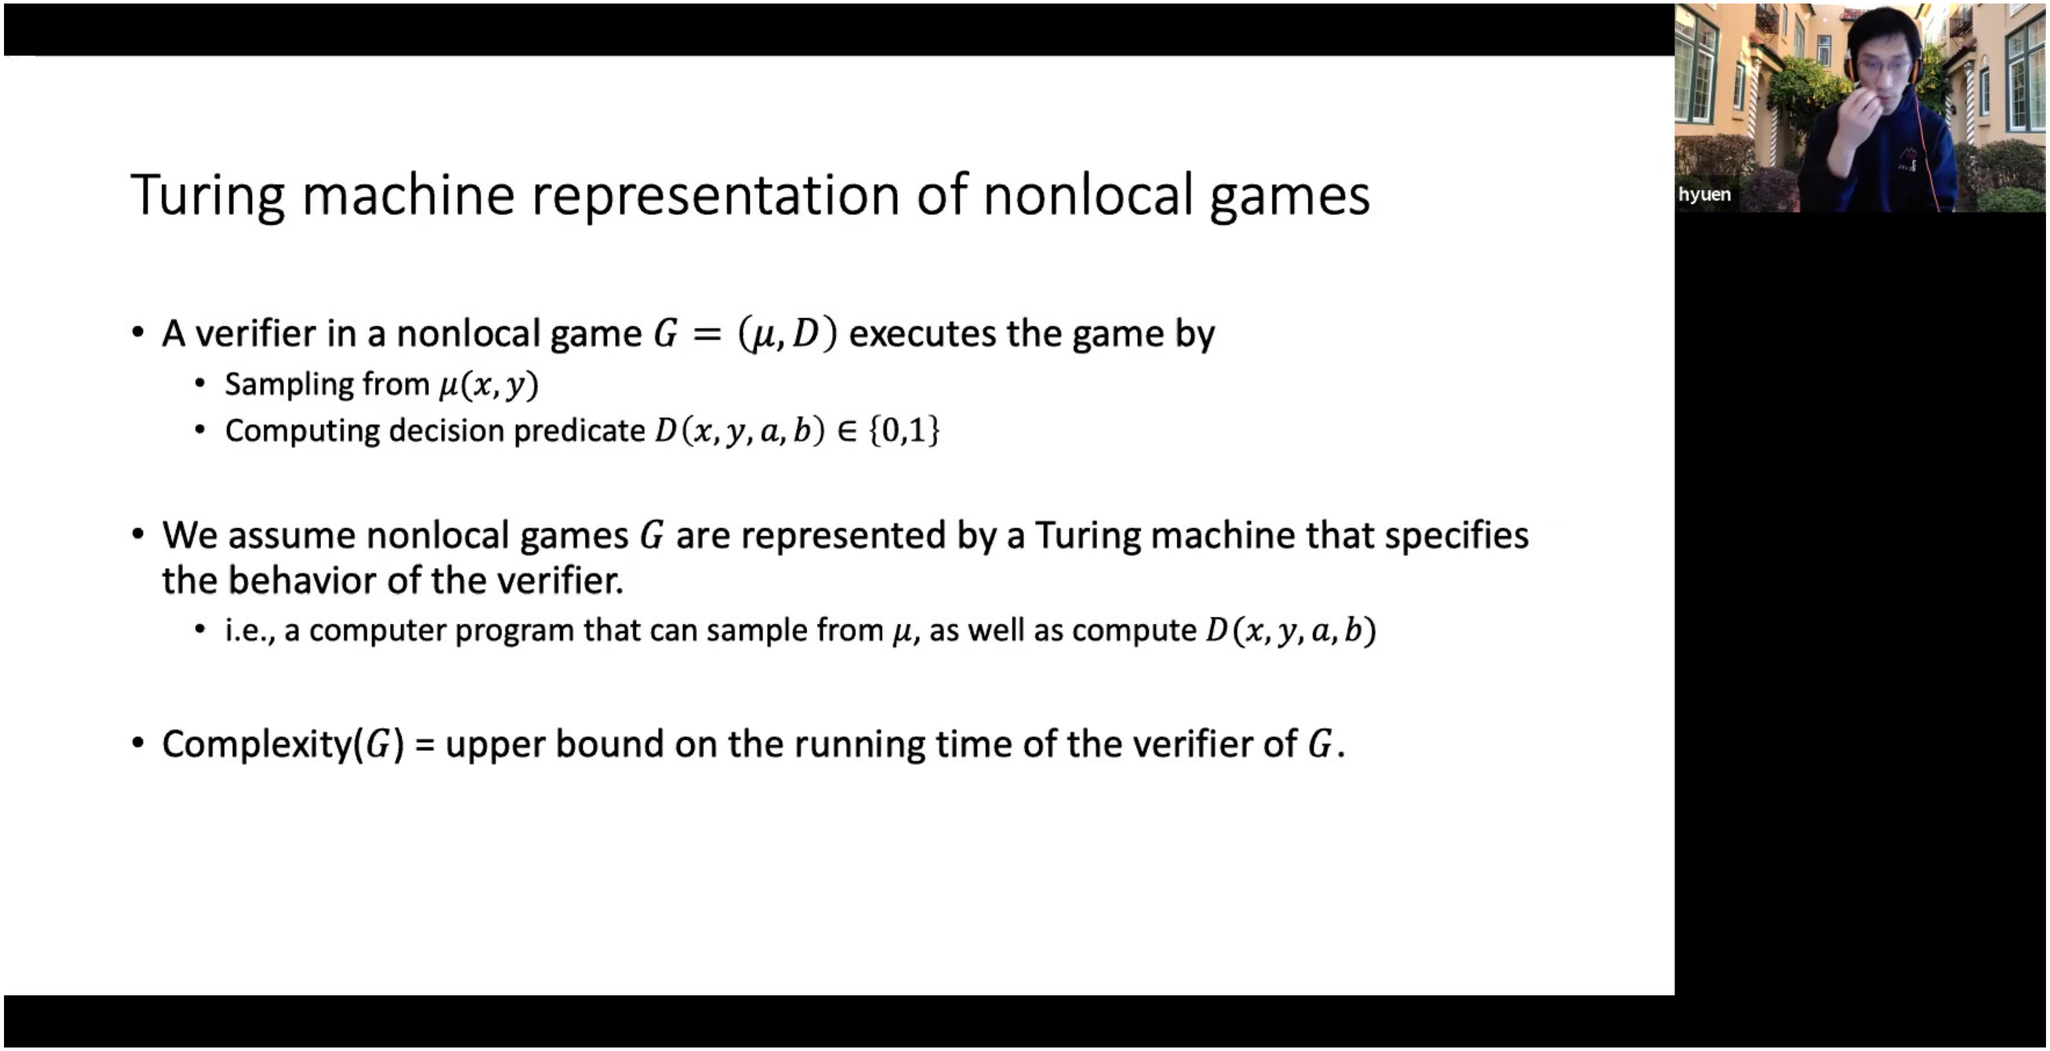

Turing machine representation of nonlocal games

- A verifier in a nonlocal game $G = (\mu, D)$ executes the game by
  - Sampling from $\mu(x, y)$
  - Computing decision predicate $D(x, y, a, b) \in \{0,1\}$

- We assume nonlocal games G are represented by a Turing machine that specifies the behavior of the verifier.
  - i.e., a computer program that can sample from μ, as well as compute $D(x, y, a, b)$

- Complexity(G) = upper bound on the running time of the verifier of G.

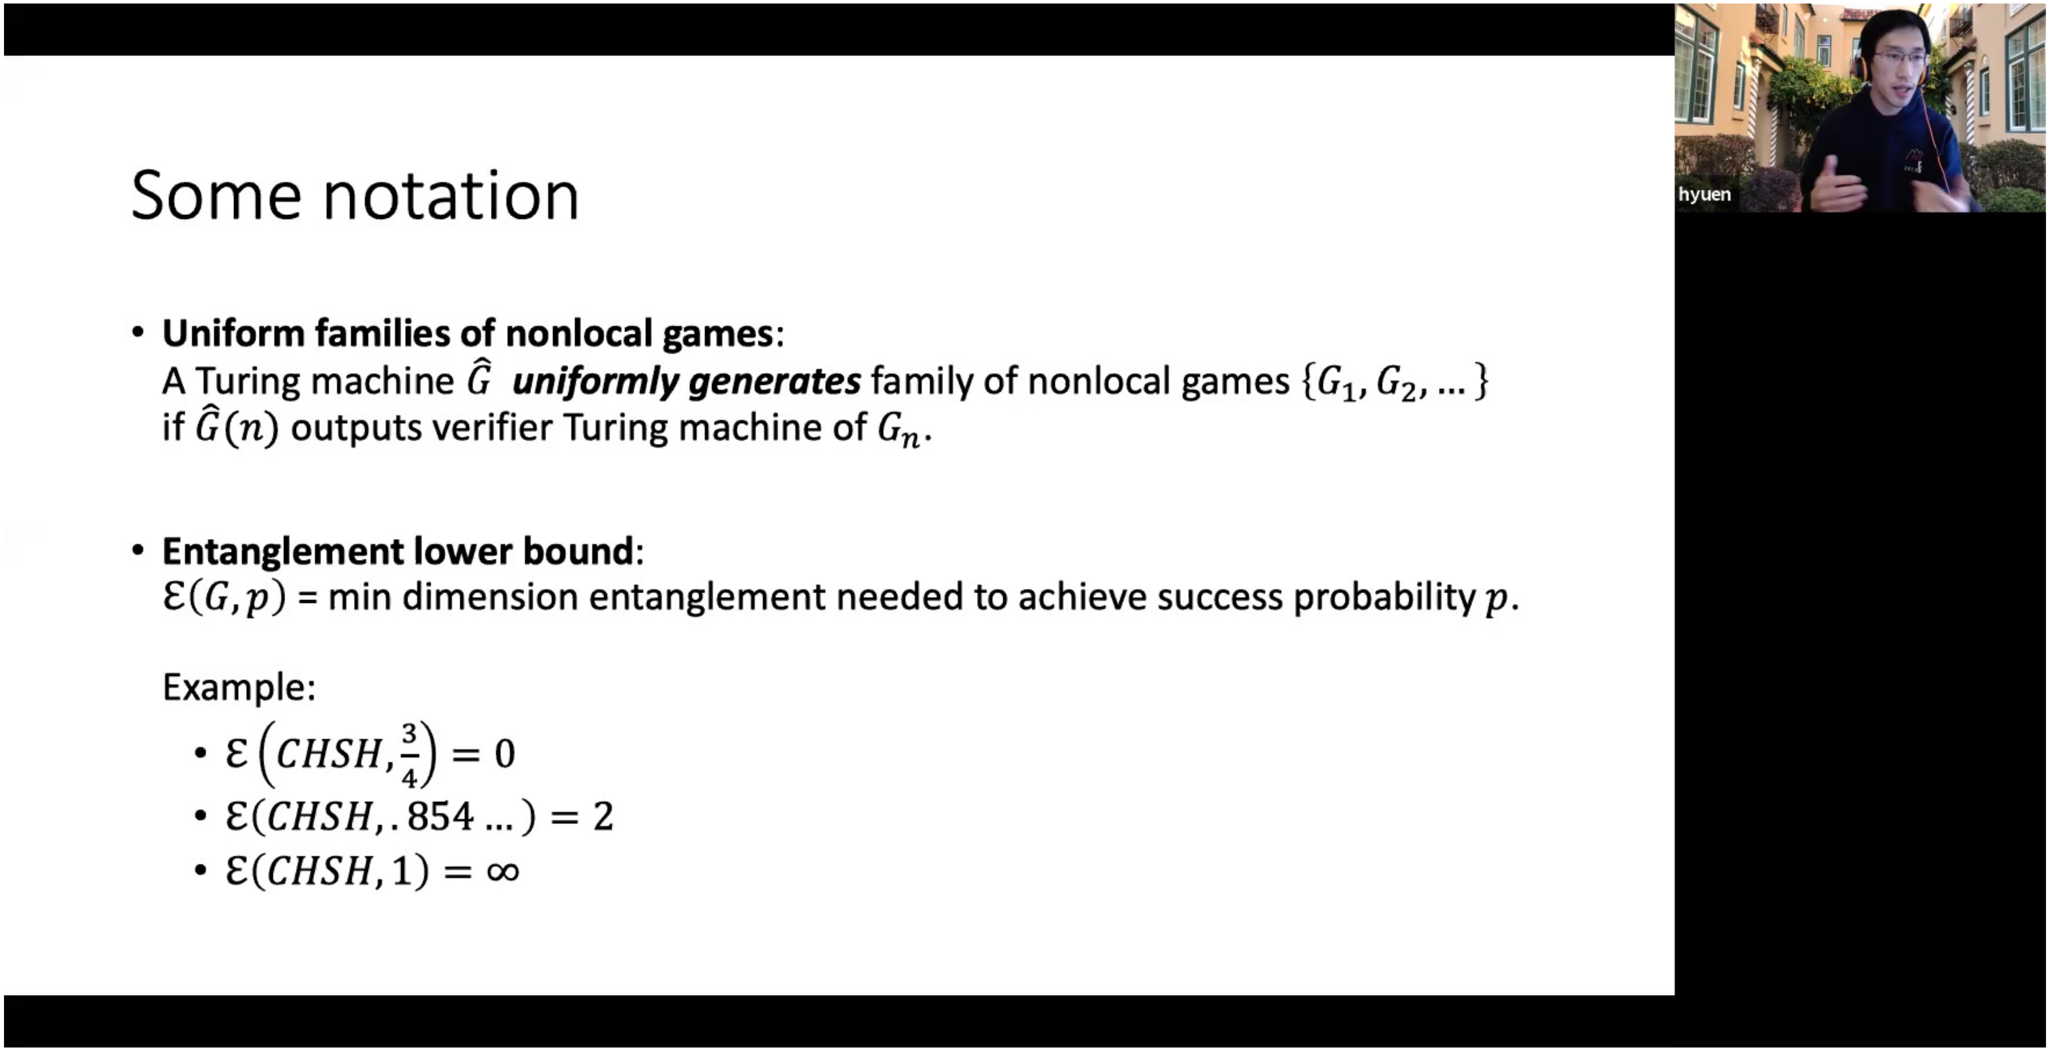

Some notation

- **Uniform families of nonlocal games**:
  A Turing machine $\widehat{G}$ ***uniformly generates*** family of nonlocal games $\{G_1, G_2, \ldots\}$ if $\widehat{G}(n)$ outputs verifier Turing machine of G_n.

- **Entanglement lower bound**:
  $\mathcal{E}(G, p)$ = min dimension entanglement needed to achieve success probability p.

  Example:
  - $\mathcal{E}\left(CHSH, \frac{3}{4}\right) = 0$
  - $\mathcal{E}(CHSH, .854\ldots) = 2$
  - $\mathcal{E}(CHSH, 1) = \infty$

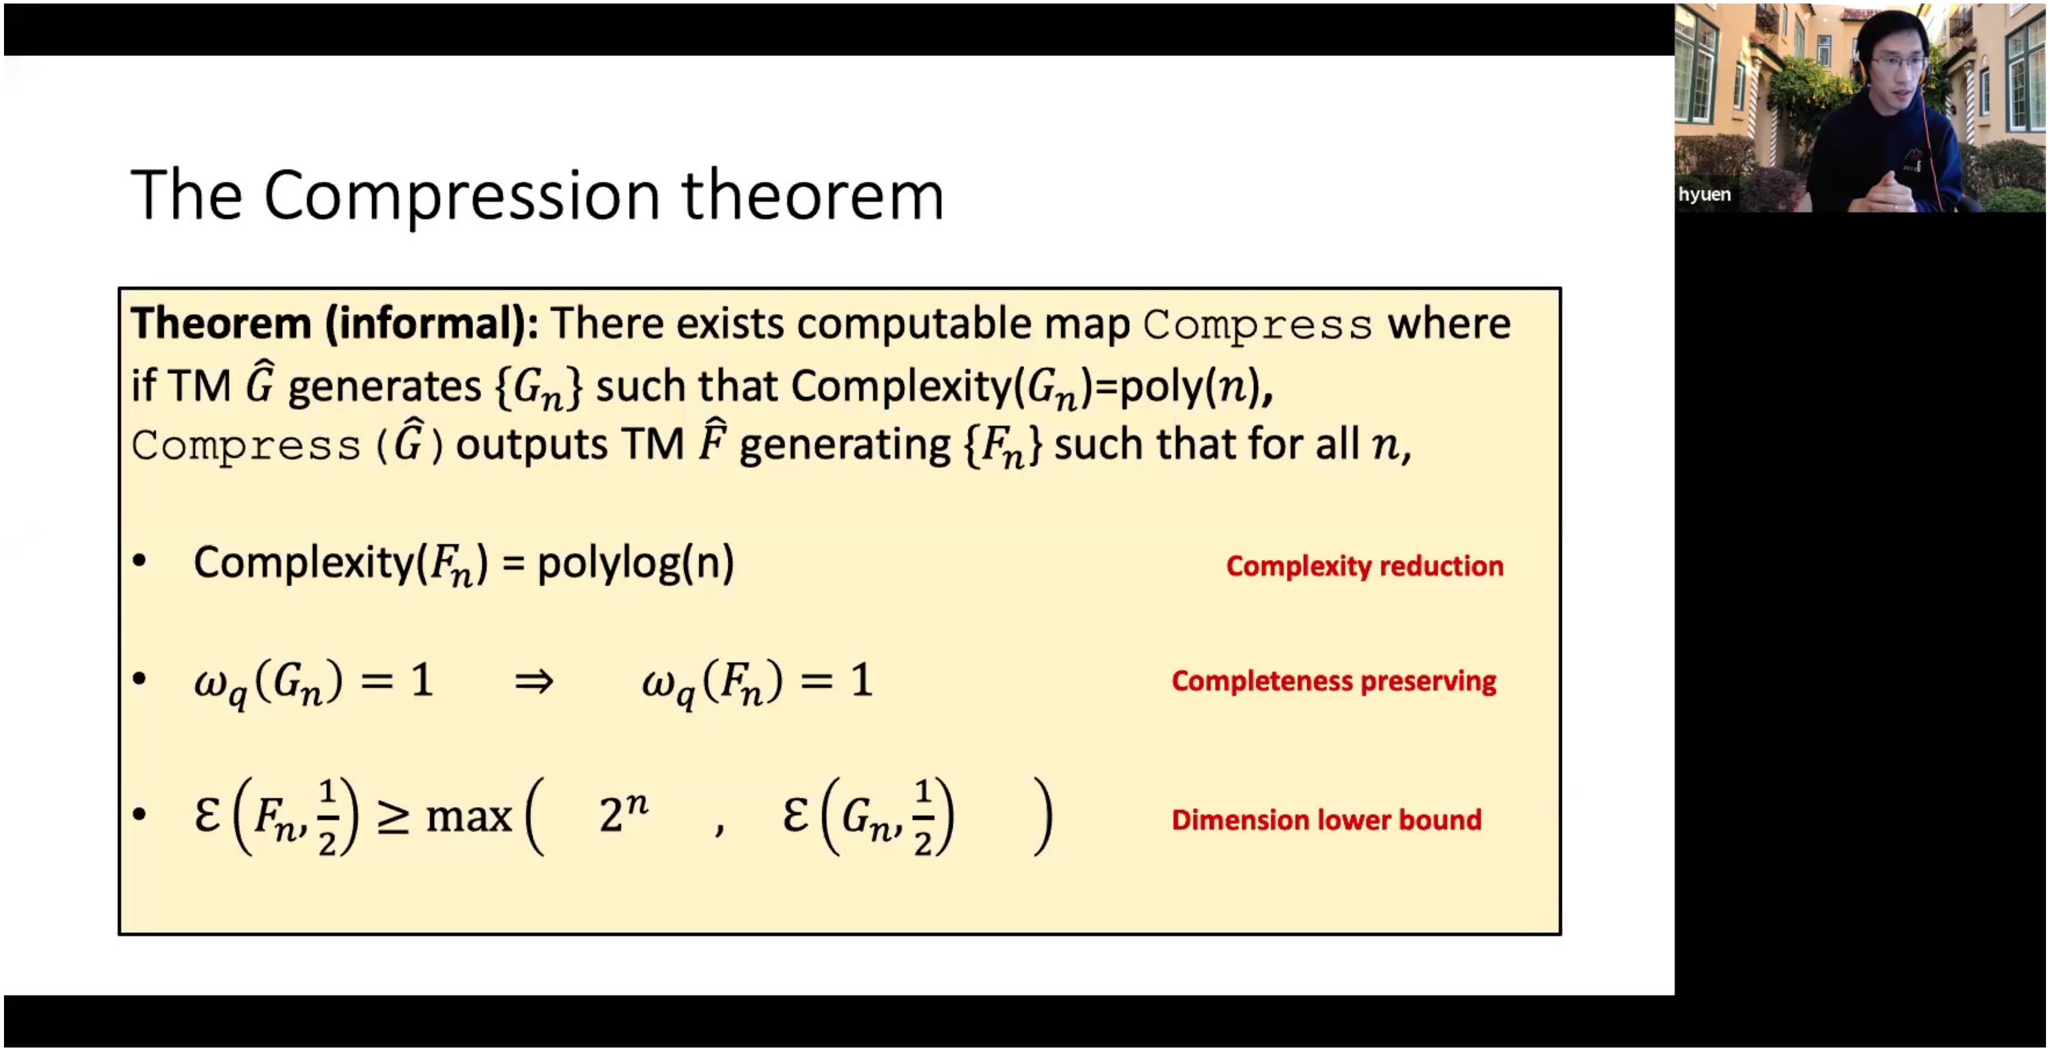

The Compression theorem

Theorem (informal): There exists computable map `Compress` where if TM $\hat{G}$ generates $\{G_n\}$ such that Complexity(G_n)=poly(n), `Compress` ($\hat{G}$) outputs TM $\hat{F}$ generating $\{F_n\}$ such that for all n,

- Complexity(F_n) = polylog(n) **Complexity reduction**
- $\omega_q(G_n) = 1 \quad \Rightarrow \quad \omega_q(F_n) = 1$ **Completeness preserving**
- $\mathcal{E}\left(F_n, \frac{1}{2}\right) \geq \max\left(\quad 2^n \quad , \quad \mathcal{E}\left(G_n, \frac{1}{2}\right) \quad \right)$ **Dimension lower bound**

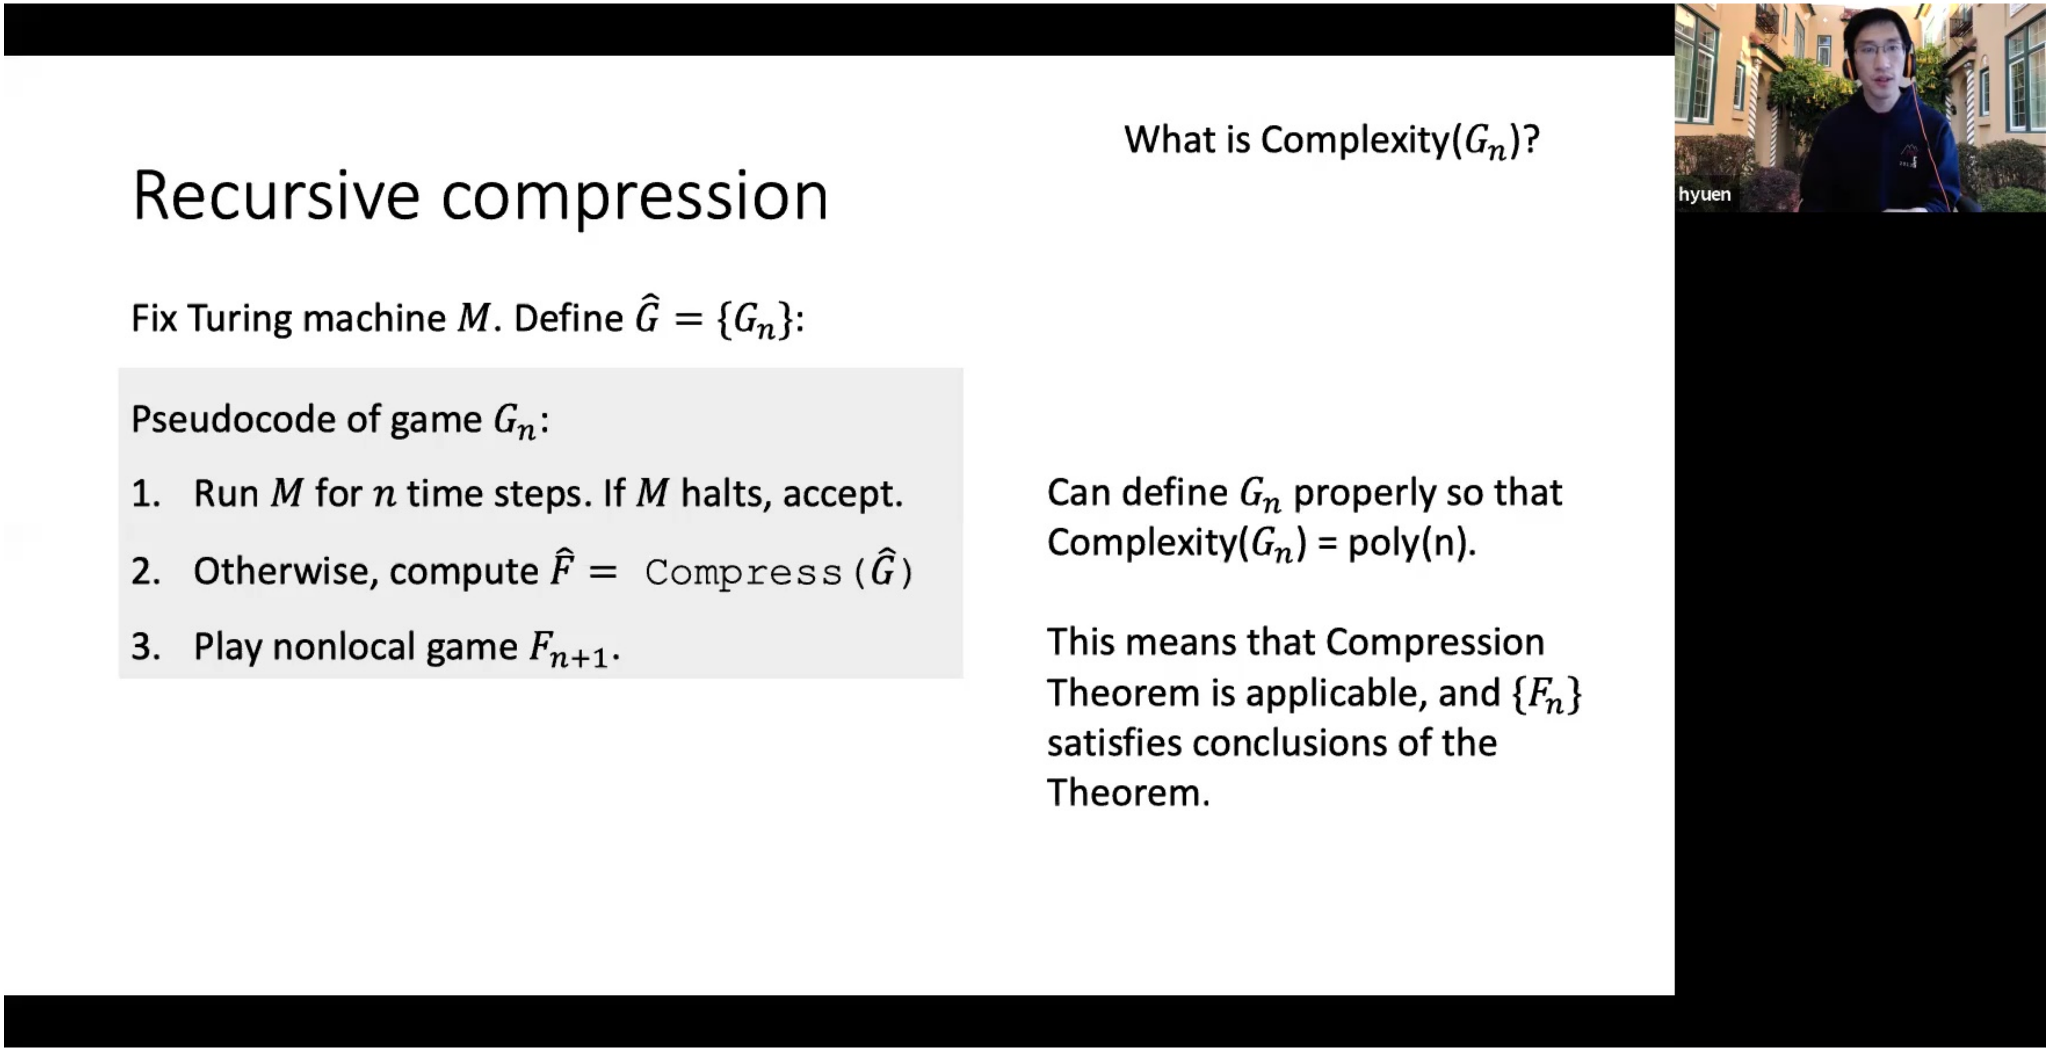

Recursive compression

What is Complexity(G_n)?

Fix Turing machine M. Define $\hat{G} = \{G_n\}$:

```
Pseudocode of game $G_n$:
1. Run $M$ for $n$ time steps. If $M$ halts, accept.
2. Otherwise, compute $\hat{F} =$ Compress($\hat{G}$)
3. Play nonlocal game $F_{n+1}$.
```

Can define G_n properly so that Complexity(G_n) = poly(n).

This means that Compression Theorem is applicable, and $\{F_n\}$ satisfies conclusions of the Theorem.

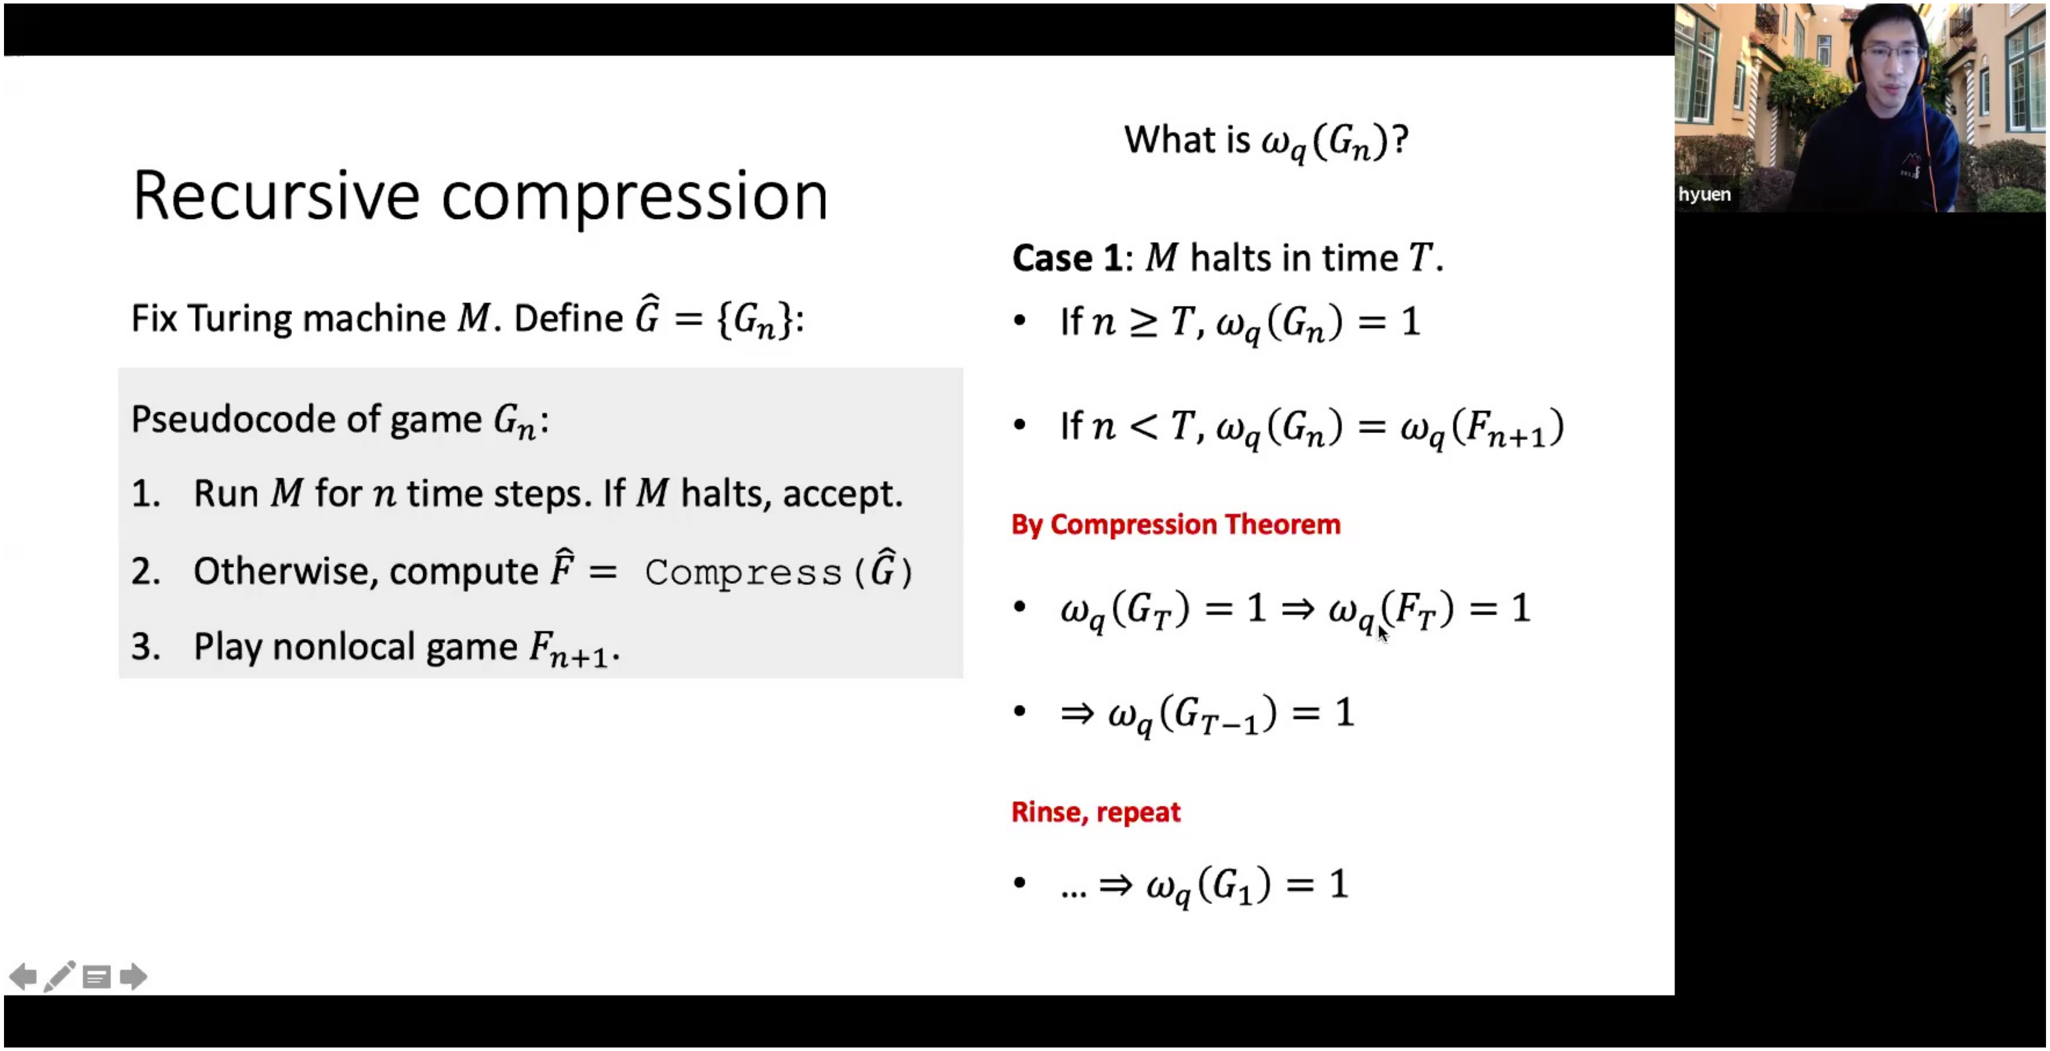

Recursive compression

Fix Turing machine M. Define $\hat{G} = \{G_n\}$:

Pseudocode of game G_n:

```
1. Run $M$ for $n$ time steps. If $M$ halts, accept.
2. Otherwise, compute $\hat{F} = \texttt{Compress}(\hat{G})$
3. Play nonlocal game $F_{n+1}$.
```

What is $\omega_q(G_n)$?

Case 1: M halts in time T.

- If $n \geq T$, $\omega_q(G_n) = 1$
- If $n < T$, $\omega_q(G_n) = \omega_q(F_{n+1})$

By Compression Theorem

- $\omega_q(G_T) = 1 \Rightarrow \omega_q(F_T) = 1$
- $\Rightarrow \omega_q(G_{T-1}) = 1$

Rinse, repeat

- ... $\Rightarrow \omega_q(G_1) = 1$

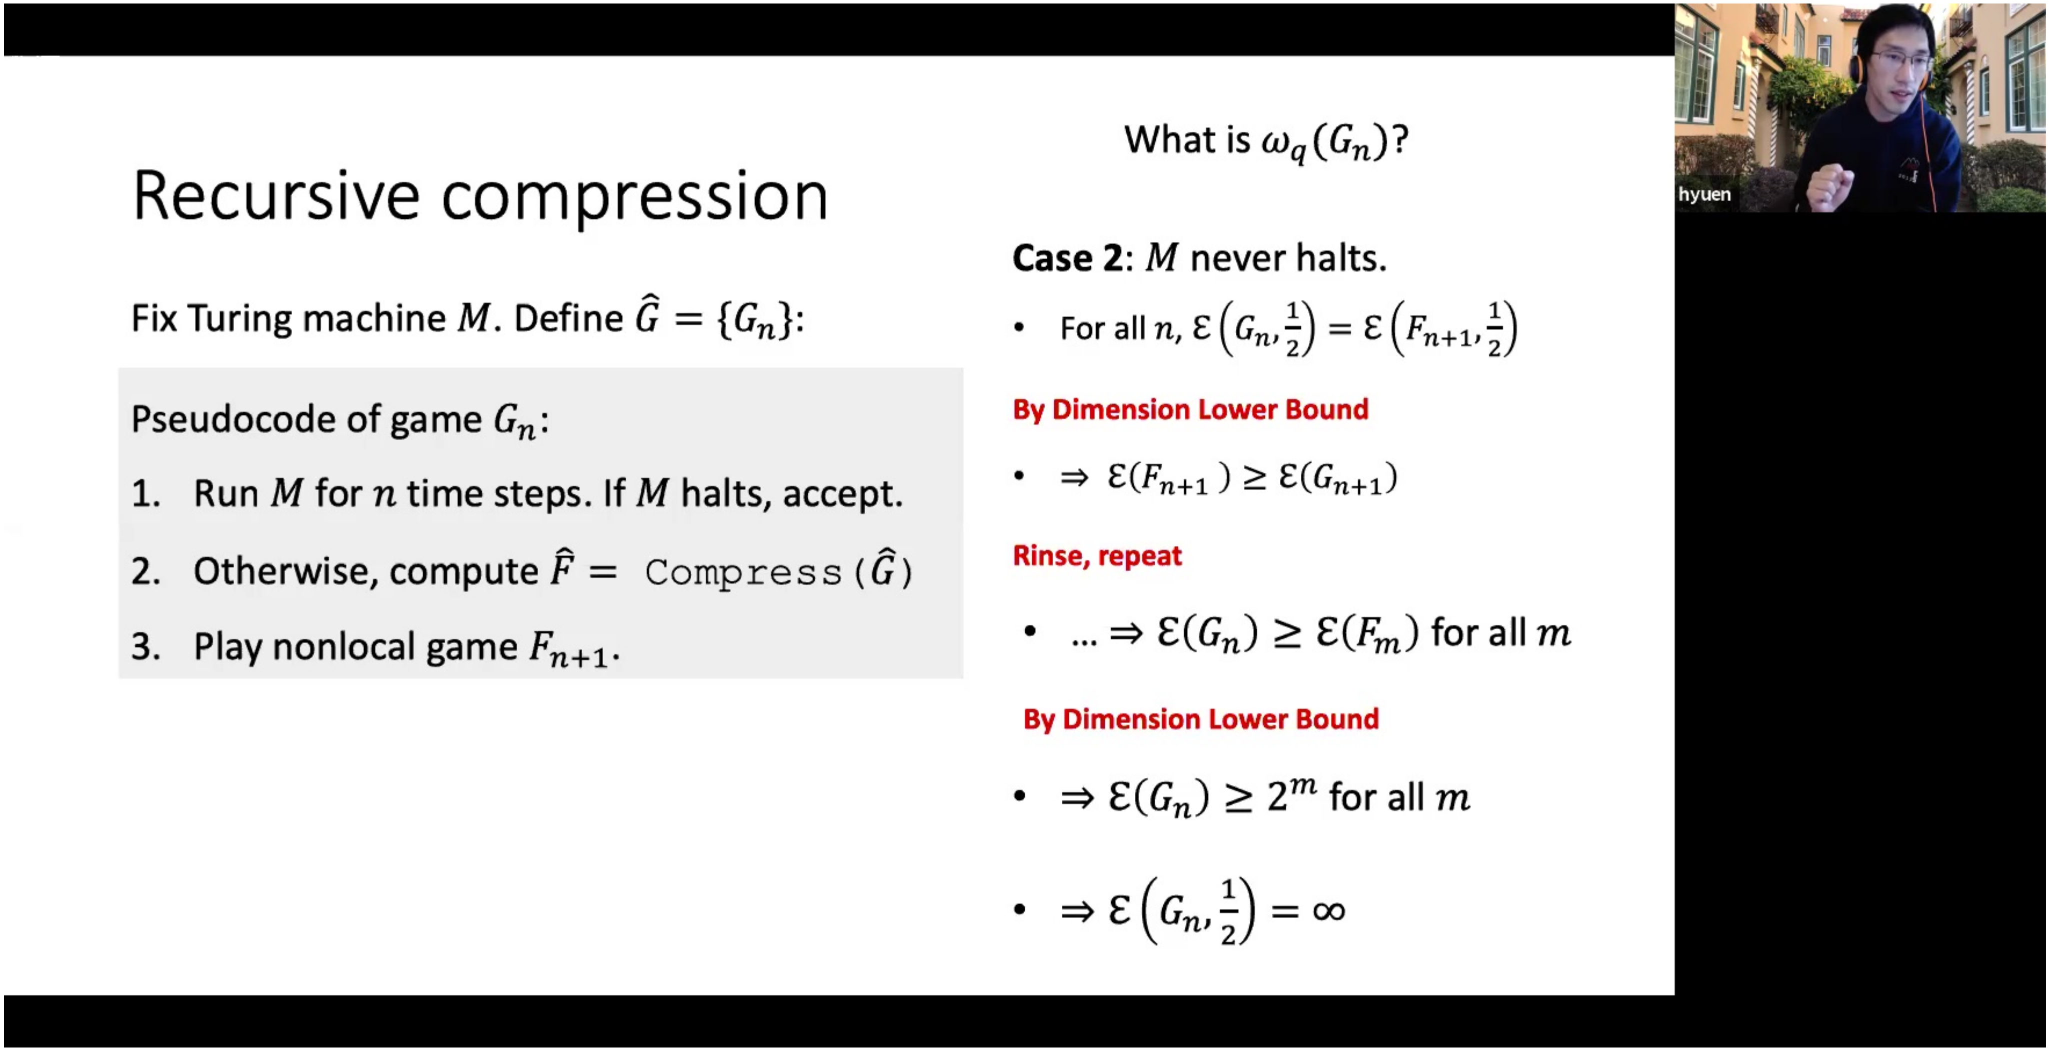

Recursive compression

Fix Turing machine M. Define $\hat{G} = \{G_n\}$:

Pseudocode of game G_n:

```
1. Run $M$ for $n$ time steps. If $M$ halts, accept.
2. Otherwise, compute $\hat{F} = \texttt{Compress}(\hat{G})$
3. Play nonlocal game $F_{n+1}$.
```

What is $\omega_q(G_n)$?

Case 2: M never halts.

- For all n, $\mathcal{E}\left(G_n, \frac{1}{2}\right) = \mathcal{E}\left(F_{n+1}, \frac{1}{2}\right)$

By Dimension Lower Bound

- $\Rightarrow \mathcal{E}(F_{n+1}) \geq \mathcal{E}(G_{n+1})$

Rinse, repeat

- $\ldots \Rightarrow \mathcal{E}(G_n) \geq \mathcal{E}(F_m)$ for all m

By Dimension Lower Bound

- $\Rightarrow \mathcal{E}(G_n) \geq 2^m$ for all m
- $\Rightarrow \mathcal{E}\left(G_n, \frac{1}{2}\right) = \infty$

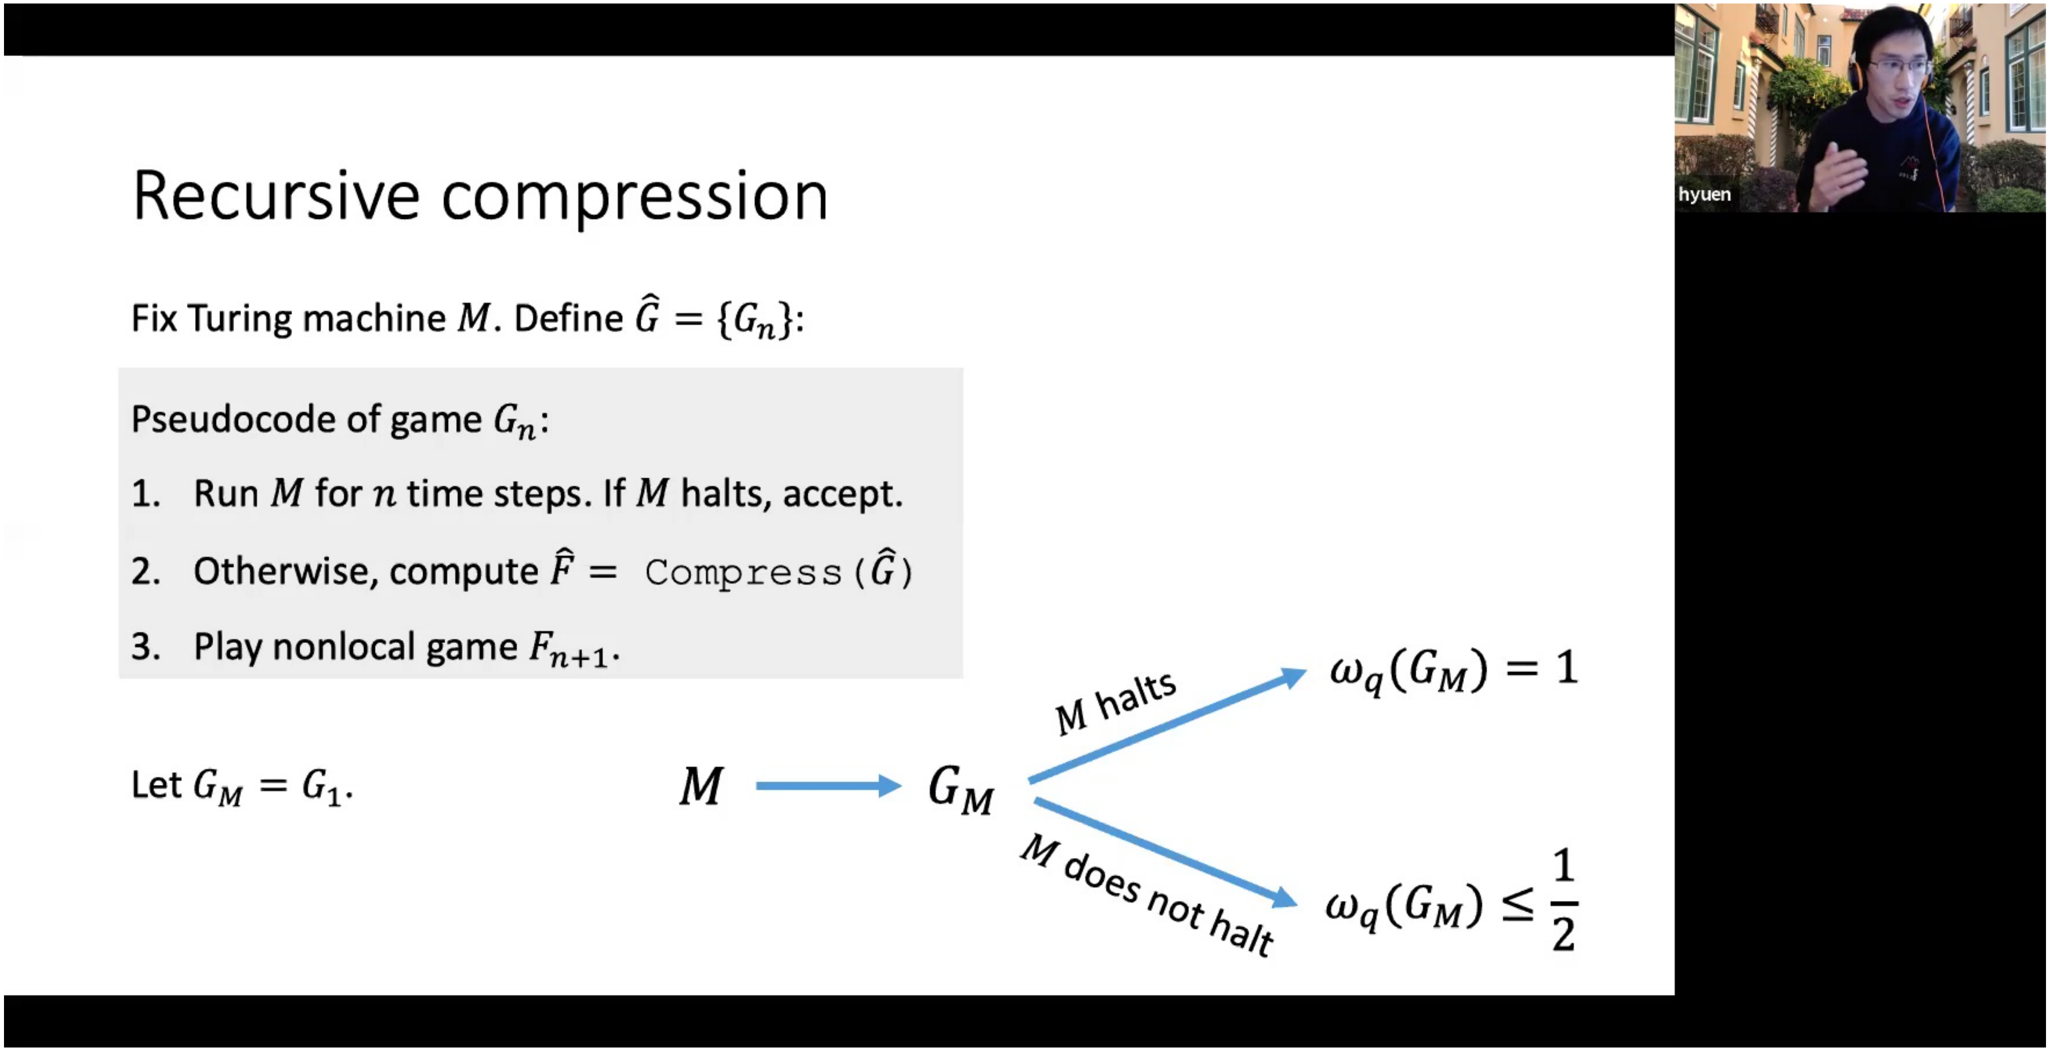

Recursive compression

Fix Turing machine M. Define $\hat{G} = \{G_n\}$:

Pseudocode of game G_n:

```
1. Run $M$ for $n$ time steps. If $M$ halts, accept.
2. Otherwise, compute $\hat{F} = $ Compress($\hat{G}$)
3. Play nonlocal game $F_{n+1}$.
```

Let $G_M = G_1$.

$$\begin{tikzcd} & & \omega_q(G_M) = 1 \\ M \arrow[r] & G_M \arrow[ur, "M \text{ halts}"] \arrow[dr, "M \text{ does not halt}"'] & \\ & & \omega_q(G_M) \leq \frac{1}{2} \end{tikzcd}$$

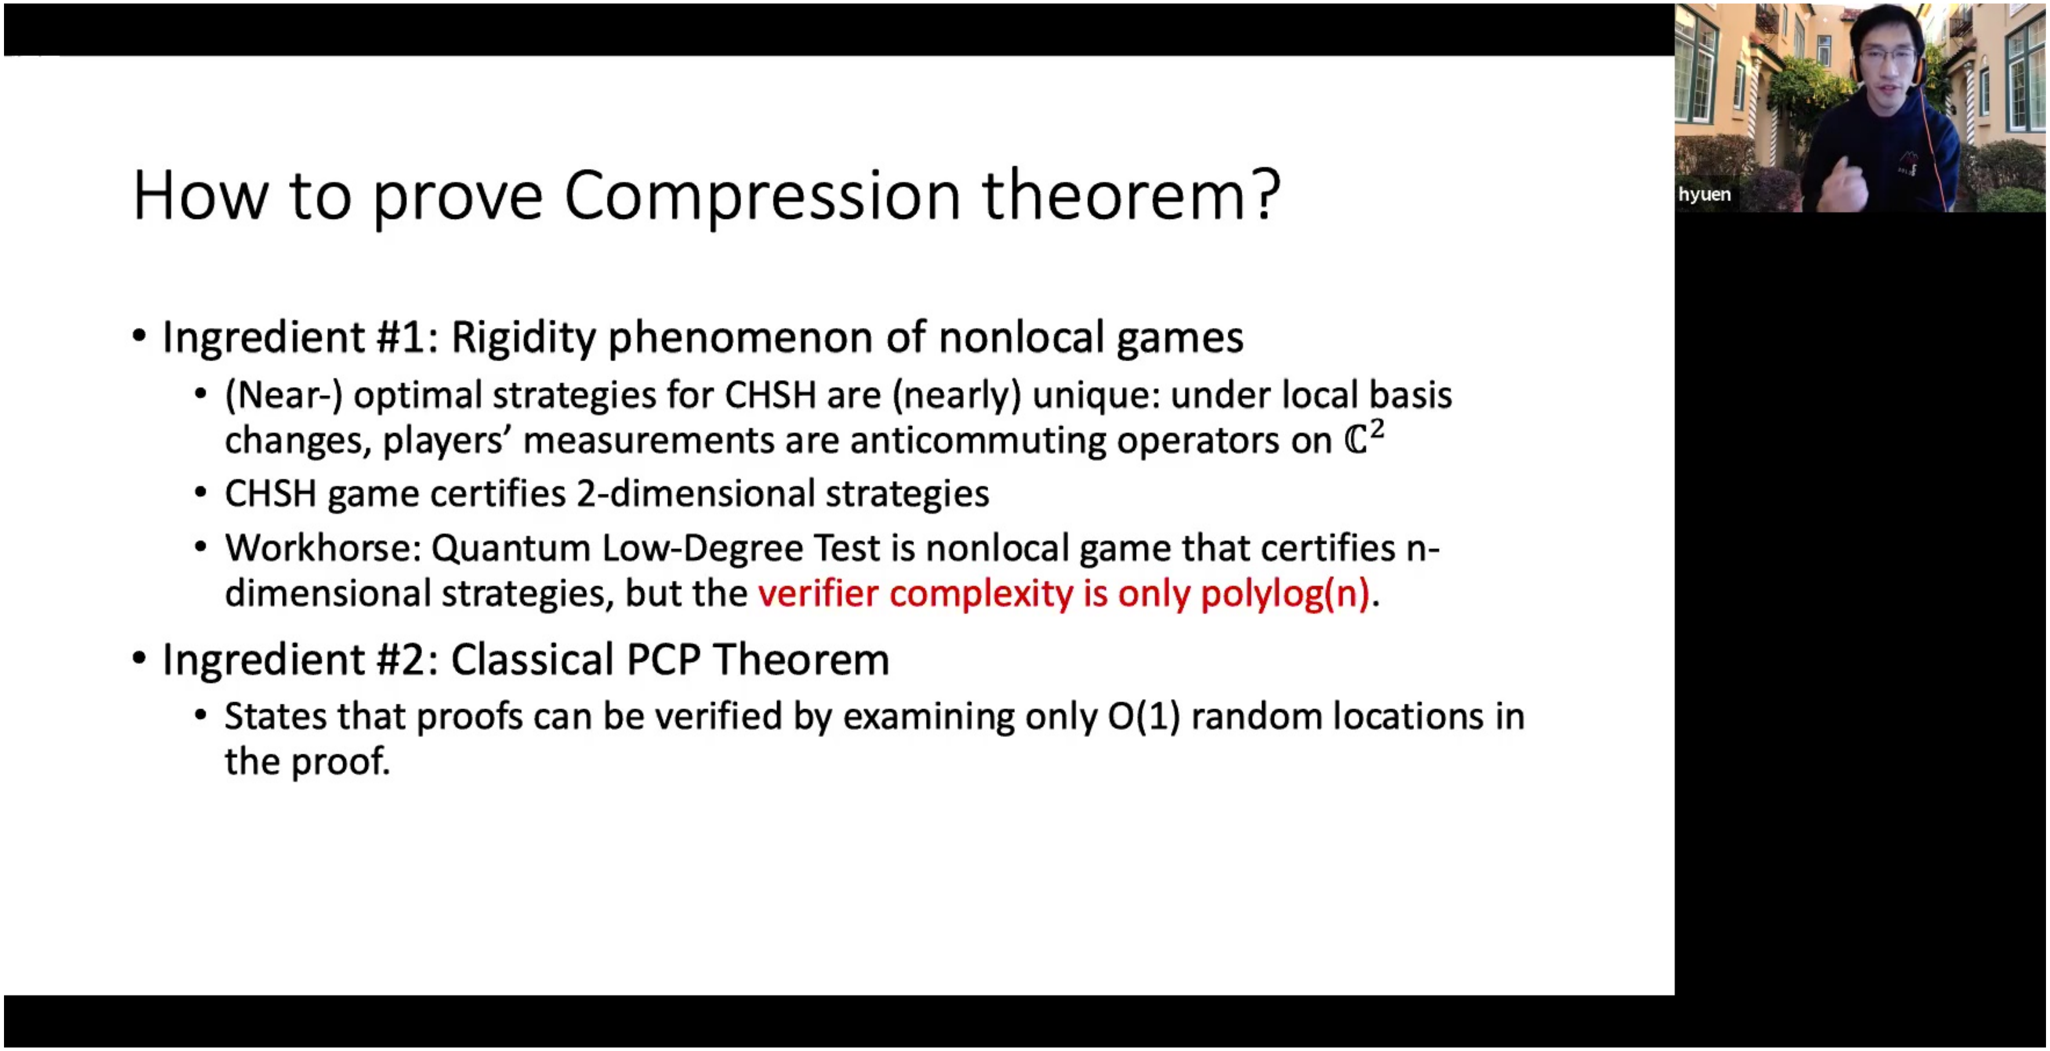

How to prove Compression theorem?

- Ingredient #1: Rigidity phenomenon of nonlocal games
  - (Near-) optimal strategies for CHSH are (nearly) unique: under local basis changes, players' measurements are anticommuting operators on $\mathbb{C}^2$
  - CHSH game certifies 2-dimensional strategies
  - Workhorse: Quantum Low-Degree Test is nonlocal game that certifies n-dimensional strategies, but the verifier complexity is only polylog(n).
- Ingredient #2: Classical PCP Theorem
  - States that proofs can be verified by examining only O(1) random locations in the proof.

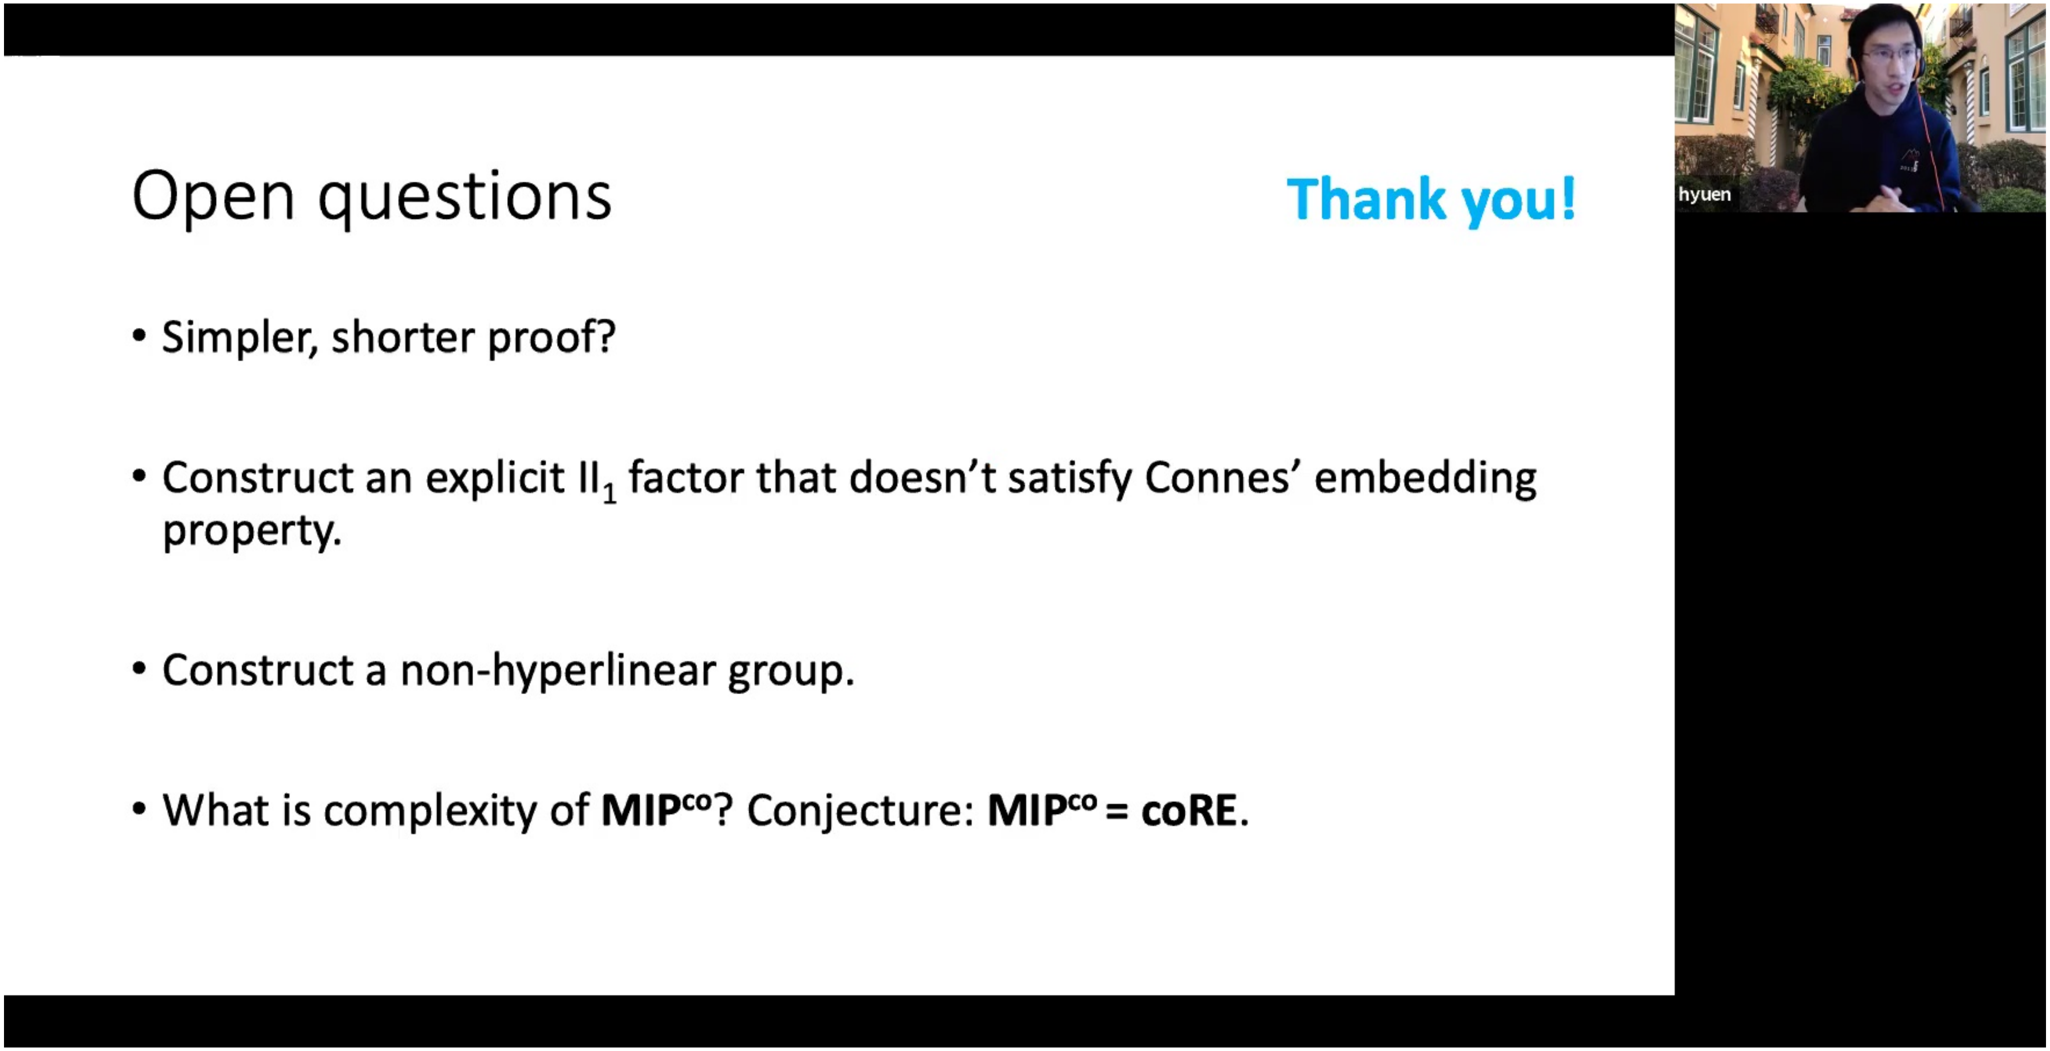

Open questions

Thank you!

- Simpler, shorter proof?
- Construct an explicit II$_1$ factor that doesn't satisfy Connes' embedding property.
- Construct a non-hyperlinear group.
- What is complexity of **MIP^{co}**? Conjecture: **MIP^{co} = coRE**.

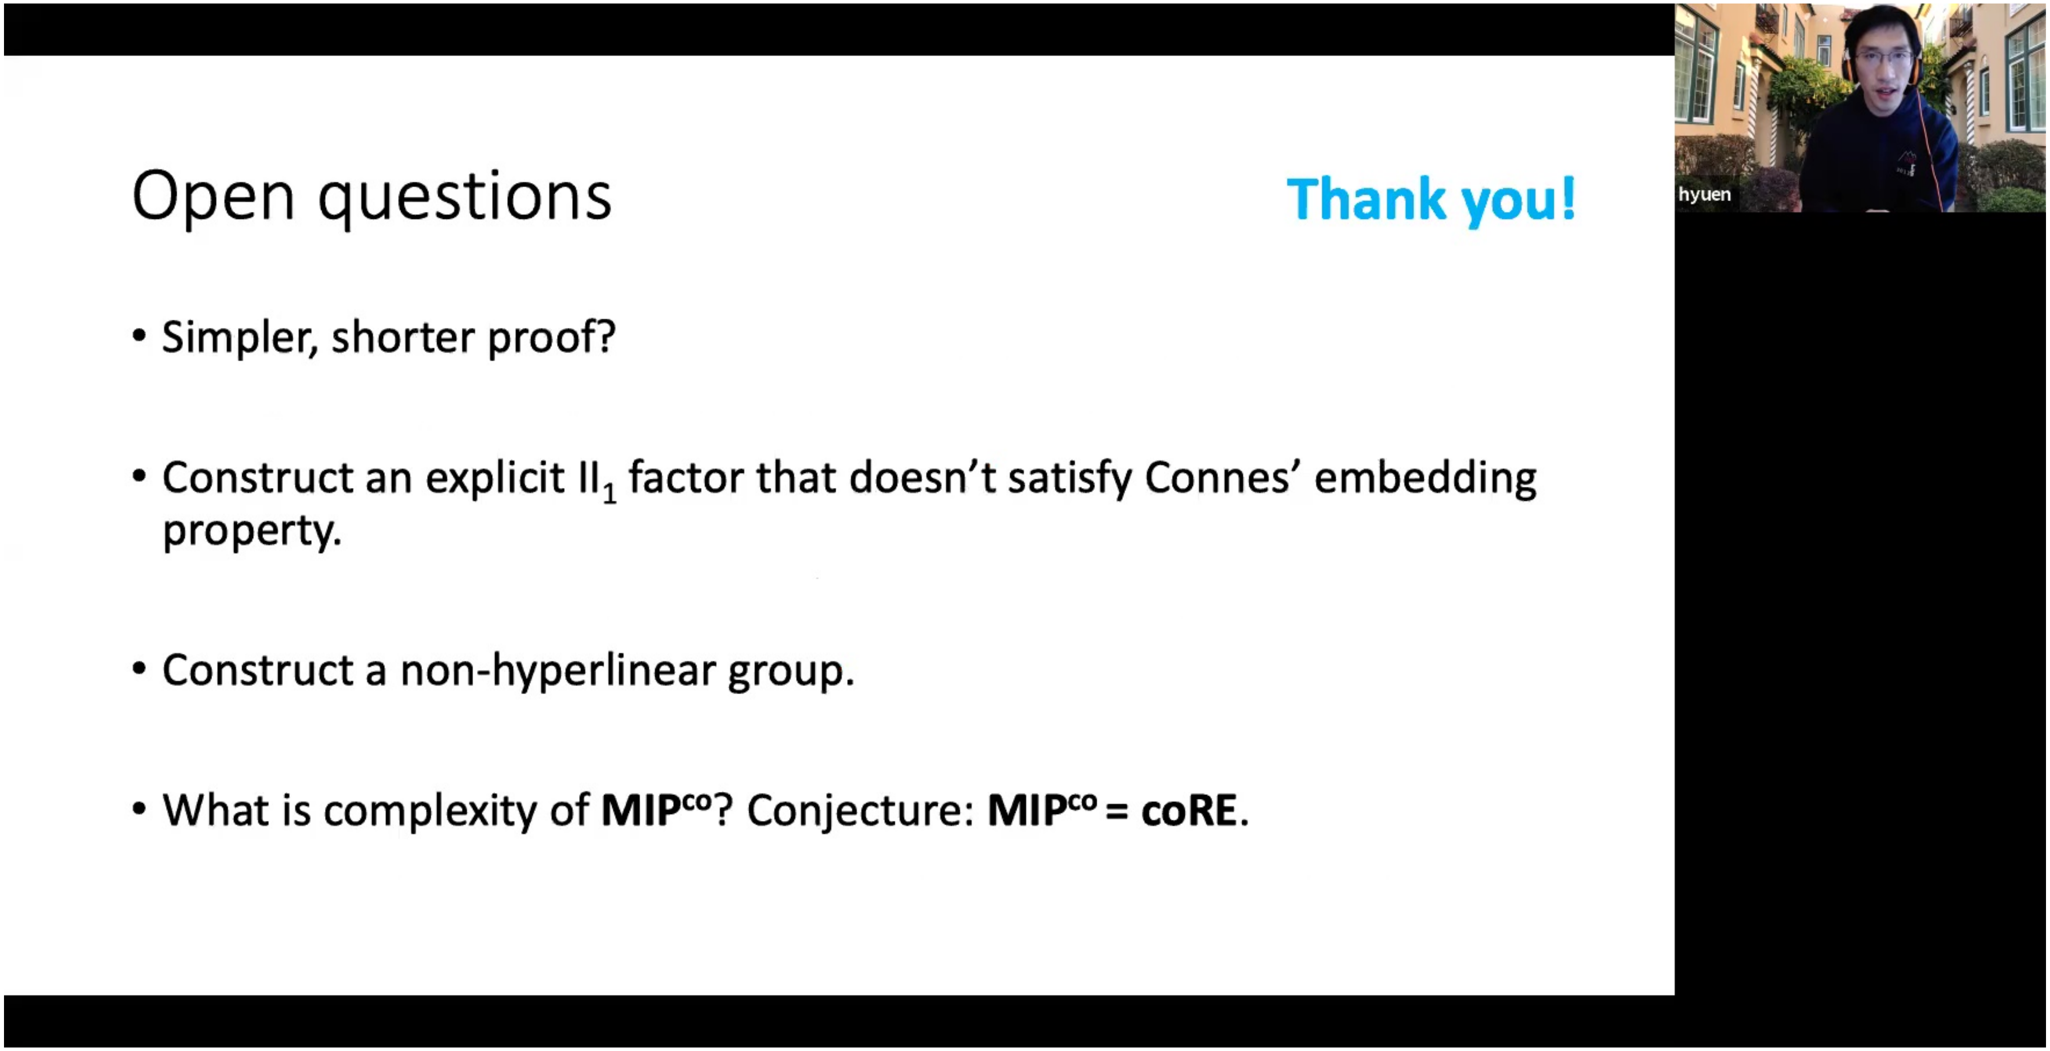

Open questions

Thank you!

- Simpler, shorter proof?
- Construct an explicit II_1 factor that doesn’t satisfy Connes’ embedding property.
- Construct a non-hyperlinear group.
- What is complexity of $\mathbf{MIP}^{\mathbf{co}}$? Conjecture: $\mathbf{MIP}^{\mathbf{co}} = \mathbf{coRE}$.