The Compression Paradigm Part II: Question Reduction

Henry Yuen

Columbia University

Reminder: (Gapless) Compression Theorem

Theorem

There exists poly-time computable $\mathtt{GaplessCompress}$ *where if* D *is decider for UGS* $\mathscr{G} = (G_n)_n$ *with complexity* n^λ *then* $\mathtt{GaplessCompress}(D, \lambda)$ *outputs a decider* D' *for* $\mathscr{G}' = (G'_n)_n$ *where*

1. **(Complexity)** D' *has complexity* $\log^\beta n$ *where* $\beta = \mathrm{poly}(\lambda)$*,*
2. **(Value)** $\omega(G'_n) = 1$ *iff* $\omega(G_n) = 1$

Compression = Question Reduction + Answer Reduction

Question Reduction

Theorem (Question Reduction)

There exists poly-time computable map $\mathtt{QuestionReduce}$ *where if* D *is decider for UGS* $\mathscr{G} = (G_n)_n$ *with complexity* n^λ *then* $\mathtt{QuestionReduce}(D, \lambda)$ *outputs a decider* D' *for* $\mathscr{G}' = (G'_n)_n$ *where*

1. **(Complexity)**

$$\mathrm{complexity}(D') \leq n^\beta$$

$$\textit{question lengths of } G'_n \leq \log^\beta n$$

for $\beta = \mathrm{poly}(\lambda)$,

Theorem (Question Reduction)

There exists poly-time computable map `QuestionReduce` *where if D is decider for UGS $\mathscr{G} = (G_n)_n$ with complexity n^λ then* $\texttt{QuestionReduce}(D, \lambda)$ *outputs a decider D' for $\mathscr{G}' = (G'_n)_n$ where*

1. **(Complexity)**

$$\mathrm{complexity}(D') \leq n^\beta$$

$$\textit{question lengths of } G'_n \leq \log^\beta n$$

   for $\beta = \mathrm{poly}(\lambda)$,

2. **(Value)** $\omega(G'_n) = 1$ *iff* $\omega(G_n) = 1$

Question Reduction: High-level idea

Super High level idea: Suppose $\mathscr{G} = (G_n)_n$ has complexity n (i.e. $\lambda = 1$).

Question Reduction: High-level idea

Super High level idea: Suppose $\mathscr{G} = (G_n)_n$ has complexity n (i.e. $\lambda = 1$).

Instead of sampling questions $(x, y) \sim \mu_n$, the game G'_n plays a random subgame:

1. (*Introspection game*) Ask Alice to sample x herself and respond with answer a, ask Bob to sample y himself and respond with answer b, and compute $D_n(x, y, a, b)$; or

Question Reduction: High-level idea

Super High level idea: Suppose $\mathscr{G} = (G_n)_n$ has complexity n (i.e. $\lambda = 1$).

Instead of sampling questions $(x, y) \sim \mu_n$, the game G'_n plays a random subgame:

1. (*Introspection game*) Ask Alice to sample x herself and respond with answer a, ask Bob to sample y himself and respond with answer b, and compute $D_n(x, y, a, b)$; or
2. (*Rigidity game*) Verify that Alice/Bob sample uniformly random questions, and **Alice does not know Bob's question and vice versa**.

Question Reduction: High-level idea

Fix n. Let $G = G_n$ and let $G' = G'_n$.

Question Reduction: High-level idea

Fix n. Let $G = G_n$ and let $G' = G'_n$.

First, we design an **honest strategy** $\mathcal{S}'$ for the Introspection subgame.

Question Reduction: High-level idea

Fix n. Let $G = G_n$ and let $G' = G'_n$.

First, we design an **honest strategy** $\mathcal{S}'$ for the Introspection subgame.

Then, we design the Rigidity game to “force” *near-optimal* strategies for G' to be *close* to $\mathcal{S}'$.

Introspection game

Introspection game is played as follows:

- Send Alice question label "$\textsc{Introspect}_A$" and get $(x, a) \in \{0, 1\}^{2n}$.
- Send Bob question label "$\textsc{Introspect}_B$" and get $(y, b) \in \{0, 1\}^{2n}$.
- Compute $D(n, x, y, a, b)$. If output is 1 or $\perp$, players win. If output is 0, players lose.

Honest strategy: Introspection game

Let $\mathcal{S} = (A_{x,a})$ be optimal strategy for G with dimension d.

Honest strategy $\mathcal{S}' = (F_{w,c})$ for the Introspection game:

1. Hilbert space: $\underbrace{(\mathbb{C}^2)^{\otimes n}}_{\text{Alice questions}} \otimes \underbrace{(\mathbb{C}^2)^{\otimes n}}_{\text{Bob questions}} \otimes \underbrace{\mathbb{C}^d}_{\text{Answers}}$

Honest strategy: Introspection game

Let $\mathcal{S} = (A_{x,a})$ be optimal strategy for G with dimension d.

Honest strategy $\mathcal{S}' = (F_{w,c})$ for the Introspection game:

1. Hilbert space: $\underbrace{(\mathbb{C}^2)^{\otimes n}}_{\text{Alice questions}} \otimes \underbrace{(\mathbb{C}^2)^{\otimes n}}_{\text{Bob questions}} \otimes \underbrace{\mathbb{C}^d}_{\text{Answers}}$
2. $F_{\textsc{Intro}_A,(x,a)} := |x\rangle\langle x| \otimes I_n \otimes A_{x,a}$

Honest strategy: Introspection game

Let $\mathcal{S} = (A_{x,a})$ be optimal strategy for G with dimension d.

Honest strategy $\mathcal{S}' = (F_{w,c})$ for the Introspection game:

1. Hilbert space: $\underbrace{(\mathbb{C}^2)^{\otimes n}}_{\text{Alice questions}} \otimes \underbrace{(\mathbb{C}^2)^{\otimes n}}_{\text{Bob questions}} \otimes \underbrace{\mathbb{C}^d}_{\text{Answers}}$
2. $F_{\textsc{Intro}_A,(x,a)} := |x\rangle\langle x| \otimes I_n \otimes A_{x,a}$
3. $F_{\textsc{Intro}_B,(y,b)} := I_n \otimes |y\rangle\langle y| \otimes A_{y,b}$

Honest strategy: Introspection game

Let $\mathcal{S} = (A_{x,a})$ be optimal strategy for G with dimension d.

Honest strategy $\mathcal{S}' = (F_{w,c})$ for the Introspection game:

1. Hilbert space: $\underbrace{(\mathbb{C}^2)^{\otimes n}}_{\text{Alice questions}} \otimes \underbrace{(\mathbb{C}^2)^{\otimes n}}_{\text{Bob questions}} \otimes \underbrace{\mathbb{C}^d}_{\text{Answers}}$
2. $F_{\mathrm{INTRO}_A,(x,a)} := |x\rangle\langle x| \otimes I_n \otimes A_{x,a}$
3. $F_{\mathrm{INTRO}_B,(y,b)} := I_n \otimes |y\rangle\langle y| \otimes A_{y,b}$

Claim: Success probability of honest strategy $\mathcal{S}'$ in Introspection game

$$(1-\alpha) + \alpha \cdot \omega(G)$$

where $\alpha = 2^{-2n} \cdot |\mathrm{supp}(\mu_n)|$.

In particular: $\omega(G) = 1$ iff $\mathcal{S}'$ wins Introspection game with probability 1.

That's nice, but why would Alice and Bob follow the honest Introspection strategy?

That's nice, but why would Alice and Bob follow the honest Introspection strategy?

Consider the following **cheating** strategy: suppose for every question pair (x, y) there is a canonical answer pair (a_{xy}, b_{xy}) where $D(n, x, y, a, b) = 1$.

That's nice, but why would Alice and Bob follow the honest Introspection strategy?

Consider the following **cheating** strategy: suppose for every question pair (x, y) there is a canonical answer pair (a_{xy}, b_{xy}) where $D(n, x, y, a, b) = 1$.

In the Introspection Game, Alice and Bob measure *both* question registers to sample (x, y). Alice outputs a_{xy} and Bob outputs b_{xy}.

This evil strategy always wins!

Rigidity game: High level

Goal of Rigidity game: Force (near-)optimal strategies of G' to be (close to) the the honest strategy for the Introspection game.

Rigidity game: High level

Goal of Rigidity game: Force (near-)optimal strategies of G' to be (close to) the the honest strategy for the Introspection game.

Rigidity game consists of three subgames:

- (**Pauli game**) Test for Pauli measurements on $2n$ qubits
- (**Sampling game**) Test $\textsc{Introspect}_A$ is consistent with standard basis measurements on Alice's question register.
- (**Don't Peek game**) Test $\textsc{Introspect}_A$ does not "peek" at Bob's question register.

Interlude: rigidity for many qubits

Magic Square

Yesterday: rigidity/self-testing for CHSH game.

Magic Square

Yesterday: rigidity/self-testing for CHSH game.

It will be more convenient to use the **Magic Square** game. The relevant properties:

- Synchronous game
- Has perfect quantum strategy
- Question set includes two questions labelled X and Z.
- Answers for questions X, Z are binary $\{0, 1\}$.

Theorem (Magic Square rigidity)

Any value-$(1-\epsilon)$ strategy for Magic Square must be $O(\sqrt{\epsilon})$-close to the honest strategy where

- *(**Two qubits**) Hilbert space: $\mathbb{C}^2 \otimes \mathbb{C}^2$*

Theorem (Magic Square rigidity)

Any value-$(1-\epsilon)$ strategy for Magic Square must be $O(\sqrt{\epsilon})$-close to the honest strategy where

- *(**Two qubits**) Hilbert space:* $\mathbb{C}^2 \otimes \mathbb{C}^2$
- *(**Standard basis**) The POVM for question Z*

$$M_{Z,0} := |0\rangle\langle 0| \otimes I, \qquad M_{Z,1} := |1\rangle\langle 1| \otimes I$$

Theorem (Magic Square rigidity)

Any value-$(1-\epsilon)$ strategy for Magic Square must be $O(\sqrt{\epsilon})$-close to the honest strategy where

- *(**Two qubits**) Hilbert space:* $\mathbb{C}^2 \otimes \mathbb{C}^2$
- *(**Standard basis**) The POVM for question Z*

$$M_{Z,0} := |0\rangle\langle 0| \otimes I, \qquad M_{Z,1} := |1\rangle\langle 1| \otimes I$$

- *(**Hadamard basis**) The POVM for question X are*

$$M_{X,0} = (H|0\rangle\langle 0|H) \otimes I \ , \qquad M_{X,1} = (H|1\rangle\langle 1|H) \otimes I$$

where $H = \frac{1}{\sqrt{2}}\begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}$.

Getting more qubits

2-out-of-n Magic Square:

1. Sample random distinct pair $1 \leq i < j \leq n$.

Getting more qubits

2-out-of-n Magic Square:

1. Sample random distinct pair $1 \leq i < j \leq n$.
2. Sample Magic Square questions (x_i, y_i) and (x_j, y_j).

Getting more qubits

2-out-of-n Magic Square:

1. Sample random distinct pair $1 \leq i < j \leq n$.
2. Sample Magic Square questions (x_i, y_i) and (x_j, y_j).
3. Send Alice (i, x_i) and (j, x_j). Get answers (a_i, a_j).
4. Sample $k \in \{i, j\}$ uniformly at random. Send Bob (k, y_k) and get answer b.

Getting more qubits

2-out-of-n Magic Square:

1. Sample random distinct pair $1 \leq i < j \leq n$.
2. Sample Magic Square questions (x_i, y_i) and (x_j, y_j).
3. Send Alice (i, x_i) and (j, x_j). Get answers (a_i, a_j).
4. Sample $k \in \{i, j\}$ uniformly at random. Send Bob (k, y_k) and get answer b.
5. Players win iff (x_k, y_k, a_k, b_k) wins Magic Square and $b = b_k$.

Question length: $O(\log n)$ **Answer length**: $O(1)$

Theorem (CRSV17, MNY17)

Any value-$(1-\epsilon)$ strategy for 2-of-n Magic Square must be $O(\mathrm{poly}(n) \cdot \sqrt{\epsilon})$-close to the honest strategy where

- *Hilbert space:* $(\mathbb{C}^2)^{\otimes 2n}$
- *Measurement for question* (k, Z) *is standard basis measurement on qubit* $2k$
- *Measurement for question* (k, X) *is Hadamard basis measurement on qubit* $2k$

Back to the Rigidity game

Pauli game

Question set includes $\{\textsc{Sample}_A, \textsc{Erase}_A, \textsc{Sample}_B, \textsc{Erase}_B\}$.

The honest strategy:

- Hilbert space: $\underbrace{(\mathbb{C}^2)^{\otimes n}}_{H_A} \otimes \underbrace{(\mathbb{C}^2)^{\otimes n}}_{H_B}$
- $\textsc{Sample}_A$ (resp. $\textsc{Sample}_B$) measures the first (resp. second) block of n qubits in standard basis.
- $\textsc{Erase}_A$ (resp. $\textsc{Erase}_B$) measures the first (resp. second) block of n qubits in the Hadamard basis.

Pauli game

POVMs for the honest strategy: for every $a \in \{0,1\}^n$,

$$F_{\textsc{Sample}_A, a} = |a\rangle\langle a| \otimes I_n,$$
$$F_{\textsc{Sample}_B, a} = I_n \otimes |a\rangle\langle a|$$

$$F_{\textsc{Erase}_A, a} = (H^{\otimes n}|a\rangle\langle a|H^{\otimes n}) \otimes I_n$$
$$F_{\textsc{Erase}_B, a} = I_n \otimes (H^{\otimes n}|a\rangle\langle a|H^{\otimes n}).$$

Properties of the Pauli game

Pauli game consists of

- 2-of-n Magic Square
- Consistency checks between Magic Square questions and $\textsc{Sample}, \textsc{Erase}$ questions.

Theorem

Any strategy with value $1-\epsilon$ in the Pauli game must be $\mathrm{poly}(n) \cdot \sqrt{\epsilon}$-close to the honest Pauli game strategy.

Sampling game

Goal of **Sampling game**: test consistency between INTROSPECT with SAMPLE.

Sampling game

Goal of **Sampling game**: test consistency between $\textsc{Introspect}$ with $\textsc{Sample}$.

To test consistency between $\textsc{Introspect}_A$ and $\textsc{Sample}_A$:

- Send $\textsc{Introspect}_A$ to Alice, get $(x, a) \in \{0, 1\}^{2n}$.
- Send $\textsc{Sample}_A$ to Bob, get $x' \in \{0, 1\}^n$.
- Accept iff $x = x'$.

Sampling game

Goal of **Sampling game**: test consistency between $\textsc{Introspect}$ with $\textsc{Sample}$.

To test consistency between $\textsc{Introspect}_A$ and $\textsc{Sample}_A$:

- Send $\textsc{Introspect}_A$ to Alice, get $(x, a) \in \{0, 1\}^{2n}$.
- Send $\textsc{Sample}_A$ to Bob, get $x' \in \{0, 1\}^n$.
- Accept iff $x = x'$.

Passing Sampling game whp means

$$F_{\textsc{Intro}_A,(x,a)} \approx \underbrace{|x\rangle\langle x|}_{\text{Alice's question}} \otimes M_{x,a}$$

for some other POVM $\{M_{x,a}\}_a$ that could act on Bob's question register.

Don't Peek game

Goal of **Don't Peek game**: test that $M_{x,a}$ does not act on Bob's question register.

Don't Peek game

Goal of **Don't Peek game**: test that $M_{x,a}$ does not act on Bob's question register.

Idea: Test that $\textsc{Introspect}_A$ (approx.) commutes with $\textsc{Sample}_B$ and $\textsc{Erase}_B$.

This implies that in fact

$$F_{\textsc{Intro}_A,(x,a)} \approx \underbrace{|x\rangle\langle x|}_{\text{Alice's question}} \otimes \underbrace{I_n}_{\text{Bob's question}} \otimes A_{x,a}$$

for some POVM $\{A_{x,a}\}_a$.

Don't Peek game

Testing that $\textsc{Introspect}_A$ (approx.) commutes with $\textsc{Erase}_B$.

- Send to Alice either $\textsc{Introspect}_A$ (getting (x, a)) or $\textsc{Erase}_B$ (getting z).
- Send $(\textsc{Introspect}_A, \textsc{Erase}_B)$ to Bob, get $(x', a', z') \in \{0, 1\}^{3n}$.
- Perform consistency check between Alice and Bob.

Putting everything together

Question Reduction

Theorem

There exists poly-time computable $\texttt{QuestionReduce}$ *where if* D *is decider for UGS* $\mathscr{G} = (G_n)_n$ *with complexity* n^λ *then* $\texttt{QuestionReduce}(D, \lambda)$ *outputs a decider* D' *for UGS* $\mathscr{G}' = (G'_n)_n$ *where*

1. **(Complexity)** *For* $\beta = \mathrm{poly}(\lambda)$*,*

$$\mathrm{complexity}(D') \leq n^\beta$$

$$\textit{question lengths of } G'_n \leq \log^\beta n$$

2. **(Value)** $\omega(G'_n) = 1$ *iff* $\omega(G_n) = 1$

`QuestionReduce`(D, λ):

Output following TM code of $D'(n, x', y', a', b')$:

If $x' = \textsc{Introspect}_A, y' = \textsc{Introspect}_B$:

1. Parse a', b' as (x, a) and (y, b), respectively.
2. Output $D(n, x, y, a, b)$.

If $x' = \textsc{Introspect}_A, y' = \textsc{Sample}_A$:

$\vdots$

`QuestionReduce`(D, λ) clearly runs in polynomial time, because it is outputs a string representing the Turing machine D', and `QuestionReduce` just has to "paste" the description of D as well as λ into the description of D'.

Question Reduction

The complexity of the question-reduced game G' satisfies:

- $\text{complexity}(D') = n^{O(\lambda)}$, because D' has to run
    1. The original decider D which has complexity n^{λ}, and
    2. The Rigidity game, which has complexity $n^{O(\lambda)}$.

Question Reduction

The complexity of the question-reduced game G' satisfies:

- $\mathrm{complexity}(D') = n^{O(\lambda)}$, because D' has to run
    1. The original decider D which has complexity n^{λ}, and
    2. The Rigidity game, which has complexity $n^{O(\lambda)}$.
- Question lengths: there are $O(1)$ questions like $\textsc{Introspect}, \textsc{Sample}, \textsc{Erase}$, and there are Pauli game questions of length $O(\log n^{\lambda})$.

Question Reduction

If $\omega(G) = 1$, then $\omega(G') = 1$ due to honest Introspection and Rigidity strategy.

Question Reduction

If $\omega(G) = 1$, then $\omega(G') = 1$ due to honest Introspection and Rigidity strategy.

If $\omega(G') = 1$, then

- Rigidity game is passed with probability 1, implying that Introspection POVMs are, up to isometry, equal to

$$F_{\text{INTRO}_A,(x,a)} \equiv |x\rangle\langle x| \otimes I_n \otimes A_{x,a}$$
$$F_{\text{INTRO}_B,(y,b)} \equiv I_n \otimes |y\rangle\langle y| \otimes A_{y,b}$$

Question Reduction

If $\omega(G) = 1$, then $\omega(G') = 1$ due to honest Introspection and Rigidity strategy.

If $\omega(G') = 1$, then

- Rigidity game is passed with probability 1, implying that Introspection POVMs are, up to isometry, equal to

$$F_{\text{INTRO}_A,(x,a)} \equiv |x\rangle\langle x| \otimes I_n \otimes A_{x,a}$$
$$F_{\text{INTRO}_B,(y,b)} \equiv I_n \otimes |y\rangle\langle y| \otimes A_{y,b}$$

- Introspection game is passed with probability 1, implying that $(A_{x,a})$ is value-1 strategy for G.

Getting a gap

For $\text{MIP}^* = \text{RE}$ we need a **gapped** Compression procedure: in addition to compressing game complexity, the procedure also preserves a gap in the game values:

- If $\omega(G) = 1$, then $\omega(G') = 1$.
- If $\omega(G) \leq \frac{1}{2}$, then $\omega(G') \leq \frac{1}{2}$.

For $\mathrm{MIP}^* = \mathrm{RE}$ we need a **gapped** Compression procedure: in addition to compressing game complexity, the procedure also preserves a gap in the game values:

- If $\omega(G) = 1$, then $\omega(G') = 1$.
- If $\omega(G) \leq \frac{1}{2}$, then $\omega(G') \leq \frac{1}{2}$.

This requires Question Reduction to preserve the gap also! (Or at least, not ruin it so much).

The non-gap-preserving Question Reduction procedure has the following effect on game value:

$$\omega(G) = 1 - \epsilon \quad \implies \quad \omega(G') \leq 1 - \frac{\epsilon}{\exp(n^c)} \ .$$

There are two main sources of “gaplessness” in today’s Question Reduction procedure:

1. The rigidity of the Pauli game gets worse as n grows large.

There are two main sources of “gaplessness” in today’s Question Reduction procedure:

1. The rigidity of the Pauli game gets worse as n grows large.
   - **Solution:** Rigidity games for many qubits with better robustness.

There are two main sources of "gaplessness" in today's Question Reduction procedure:

1. The rigidity of the Pauli game gets worse as n grows large.
   - **Solution:** Rigidity games for many qubits with better robustness.
2. The question-reduced game G' automatically wins if Alice and Bob introspect a pair of questions (x, y) not in the support of G (and support may be a vanishing fraction of $\mathcal{X} \times \mathcal{X}$).

There are two main sources of "gaplessness" in today's Question Reduction procedure:

1. The rigidity of the Pauli game gets worse as n grows large.
    - **Solution:** Rigidity games for many qubits with better robustness.
2. The question-reduced game G' automatically wins if Alice and Bob introspect a pair of questions (x, y) not in the support of G (and support may be a vanishing fraction of $\mathcal{X} \times \mathcal{X}$).
    - **Solution:** Design Introspection games to sample from larger class of question distributions.

Next time

Tomorrow (Anand): Getting a better gap for Question Reduction.

Thursday (Part 3): Answer Reduction.

Next time

Tomorrow (Anand): Getting a better gap for Question Reduction.

Thursday (Part 3): Answer Reduction.

Thank you

Better Rigidity

Recall: Today's Rigidity game has guarantee that any value-$(1-\epsilon)$ strategy must be $\mathrm{poly}(n) \cdot \sqrt{\epsilon}$-close to an n-qubit strategy.

- We only get nontrivial guarantees when the success probability is at least $1 - \frac{1}{\mathrm{poly}(n)}$.

Better Rigidity

Recall: Today's Rigidity game has guarantee that any value-$(1-\epsilon)$ strategy must be $\mathrm{poly}(n) \cdot \sqrt{\epsilon}$-close to an n-qubit strategy.

- We only get nontrivial guarantees when the success probability is at least $1 - \frac{1}{\mathrm{poly}(n)}$.

To get a **gap**, we need a better Rigidity game.

Better Rigidity

Dream Rigidity Game:

- **Low complexity**: $\log^\beta n$
- **High robustness**: any value-$(1-\epsilon)$ strategy must be $\delta(\epsilon)$-close to an n-qubit strategy, where $\delta(\cdot)$ does not depend on n!

Better Rigidity

Dream Rigidity Game:

- **Low complexity**: $\log^{\beta} n$
- **High robustness**: any value-$(1-\epsilon)$ strategy must be $\delta(\epsilon)$-close to an n-qubit strategy, where $\delta(\cdot)$ does not depend on n!

Unfortunately, don't know (yet) whether this is possible!

Better Rigidity

Theorem (Ji-Natarajan-Vidick-Wright-Y. '22)

There exists a UGS $\mathscr{R} = (R_n)_n$ where

- **Low complexity**: $\mathrm{complexity}(R_n) = \log^\beta n$.
- **High robustness**: *any value-$(1-\epsilon)$ strategy for R_n must be $\delta(\epsilon, n)$-close to an honest n-qubit strategy involving Pauli measurements, where*

$$\delta(\epsilon, n) = \mathrm{poly}\log(n) \cdot \epsilon^c \ .$$

Better Rigidity

Theorem (Ji-Natarajan-Vidick-Wright-Y. '22)

There exists a UGS $\mathscr{R} = (R_n)_n$ *where*

- **Low complexity**: $\mathrm{complexity}(R_n) = \log^{\beta} n$.
- **High robustness**: *any value-*$(1-\epsilon)$ *strategy for* R_n *must be* $\delta(\epsilon, n)$*-close to an honest* n*-qubit strategy involving Pauli measurements, where*

$$\delta(\epsilon, n) = \mathrm{poly}\log(n) \cdot \epsilon^{c} \ .$$

Milder dependence on n, and sufficient to get Question Reduction with better gap!

Introspecting complex distributions

Today's Introspection game is a way to force Alice and Bob to sample uniform, independent strings.

However, the question distributions in games arising from Question and Answer Reduction are far from uniform!

Introspecting complex distributions

Today's Introspection game is a way to force Alice and Bob to sample uniform, independent strings.

However, the question distributions in games arising from Question and Answer Reduction are far from uniform!

Goal: Design Introspection game to force Alice and Bob to sample correlated questions (x, y) from those distributions?

Introspecting complex distributions

Today's Introspection game is a way to force Alice and Bob to sample uniform, independent strings.

However, the question distributions in games arising from Question and Answer Reduction are far from uniform!

Goal: Design Introspection game to force Alice and Bob to sample correlated questions (x, y) from those distributions?

Alternate perspective: design games with "introspectable" question distributions!

The proof of $\mathsf{MIP}^* = \mathsf{RE}$ identifies a class of distributions called **conditionally linear distributions**, and shows:

- Such distributions can be robustly introspected with few questions.
- All games from Question and Answer Reduction procedures can be designed to use conditionally linear distributions.