Quantum Advantage from Sampling Shallow Circuits: Beyond Hardness of Marginals

Daniel Grier
UC San Diego

Daniel M. Kane
UC San Diego

Jackson Morris
UC San Diego

Anthony Ostuni
UC San Diego

Kewen Wu
IAS

What does quantum advantage even mean?

There is a problem that can be solved by a family of quantum circuits that cannot be solved by a similar family of classical circuits.

"Problem"

Classical inputs, classical outputs

Doesn't have to be useful

Nice-to-have

Implementable in the near term

Verifiable in polynomial time

Requires zero conjectures

Complexity theoretic view of quantum advantage

Traditional Goal:
Find a problem in BQP that is not in P

Barrier: Hard to find lower bounds for P

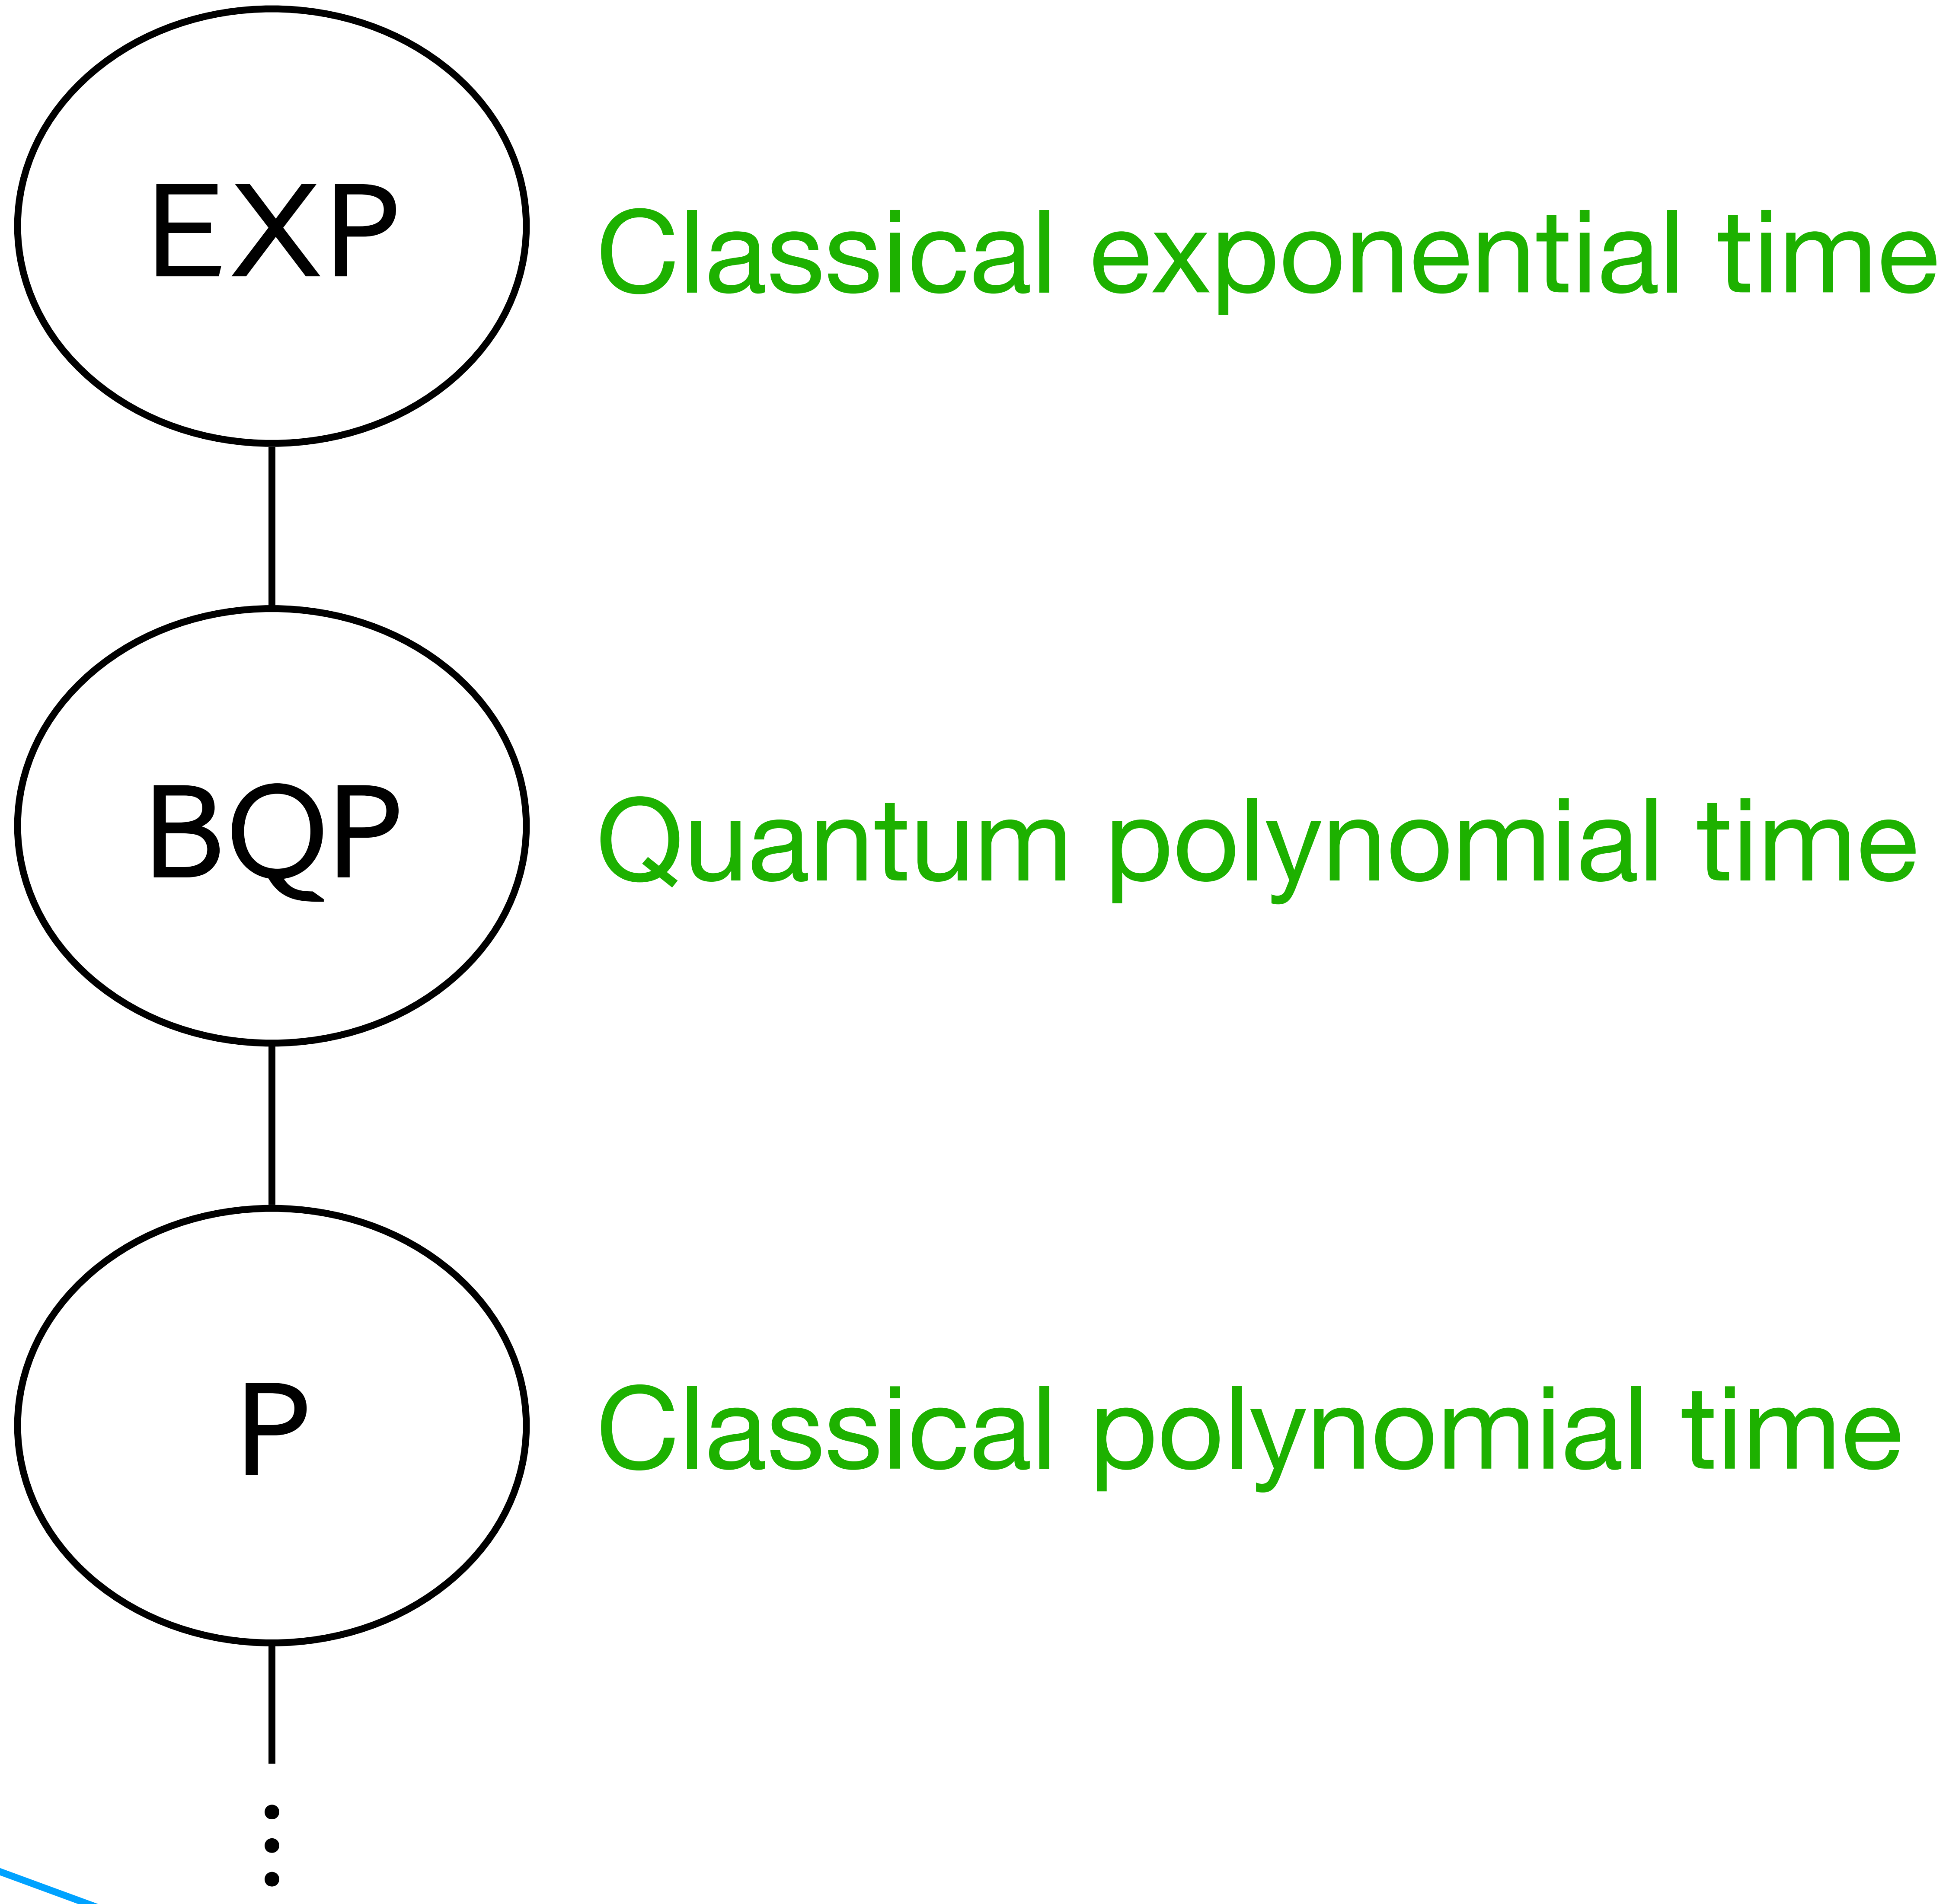

But what's down here?

Diagram of low-depth complexity classes

Hooray: Possible to prove shallow classical circuits can't solve certain problems

→ Can we find shallow quantum circuits to solve those problems?

Constant-depth classical circuits →

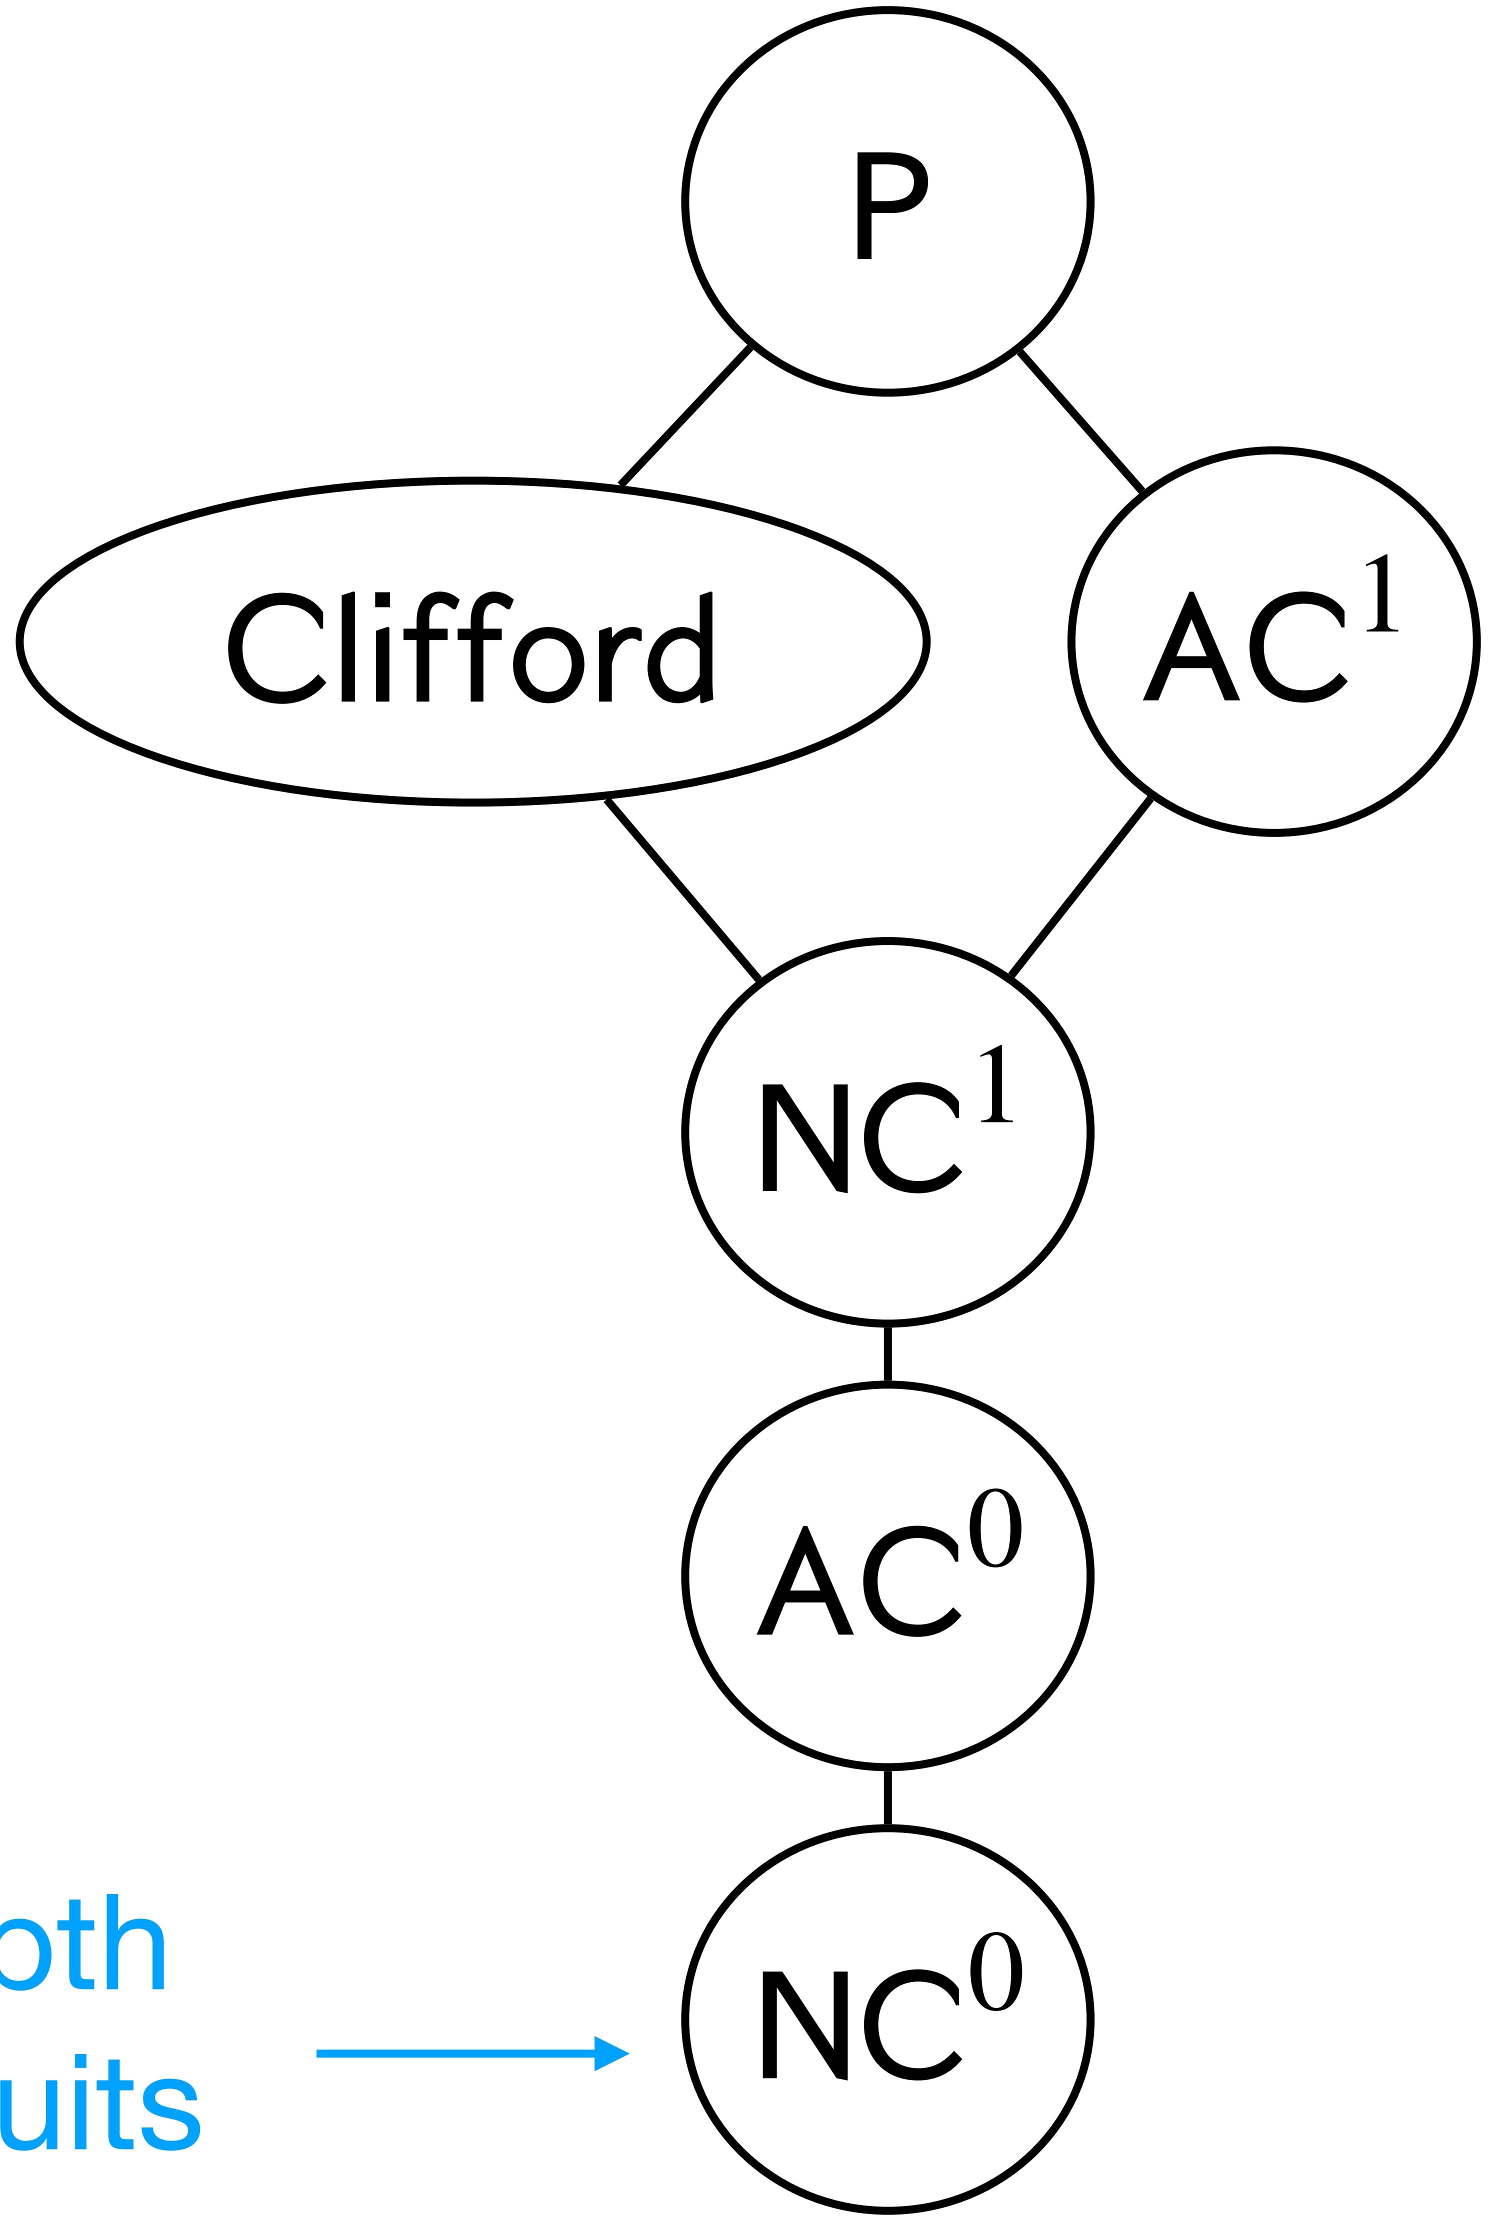

Example: Separating quantum from classical

Constant depth

No large gates → NC^0

Quantum → QNC^0

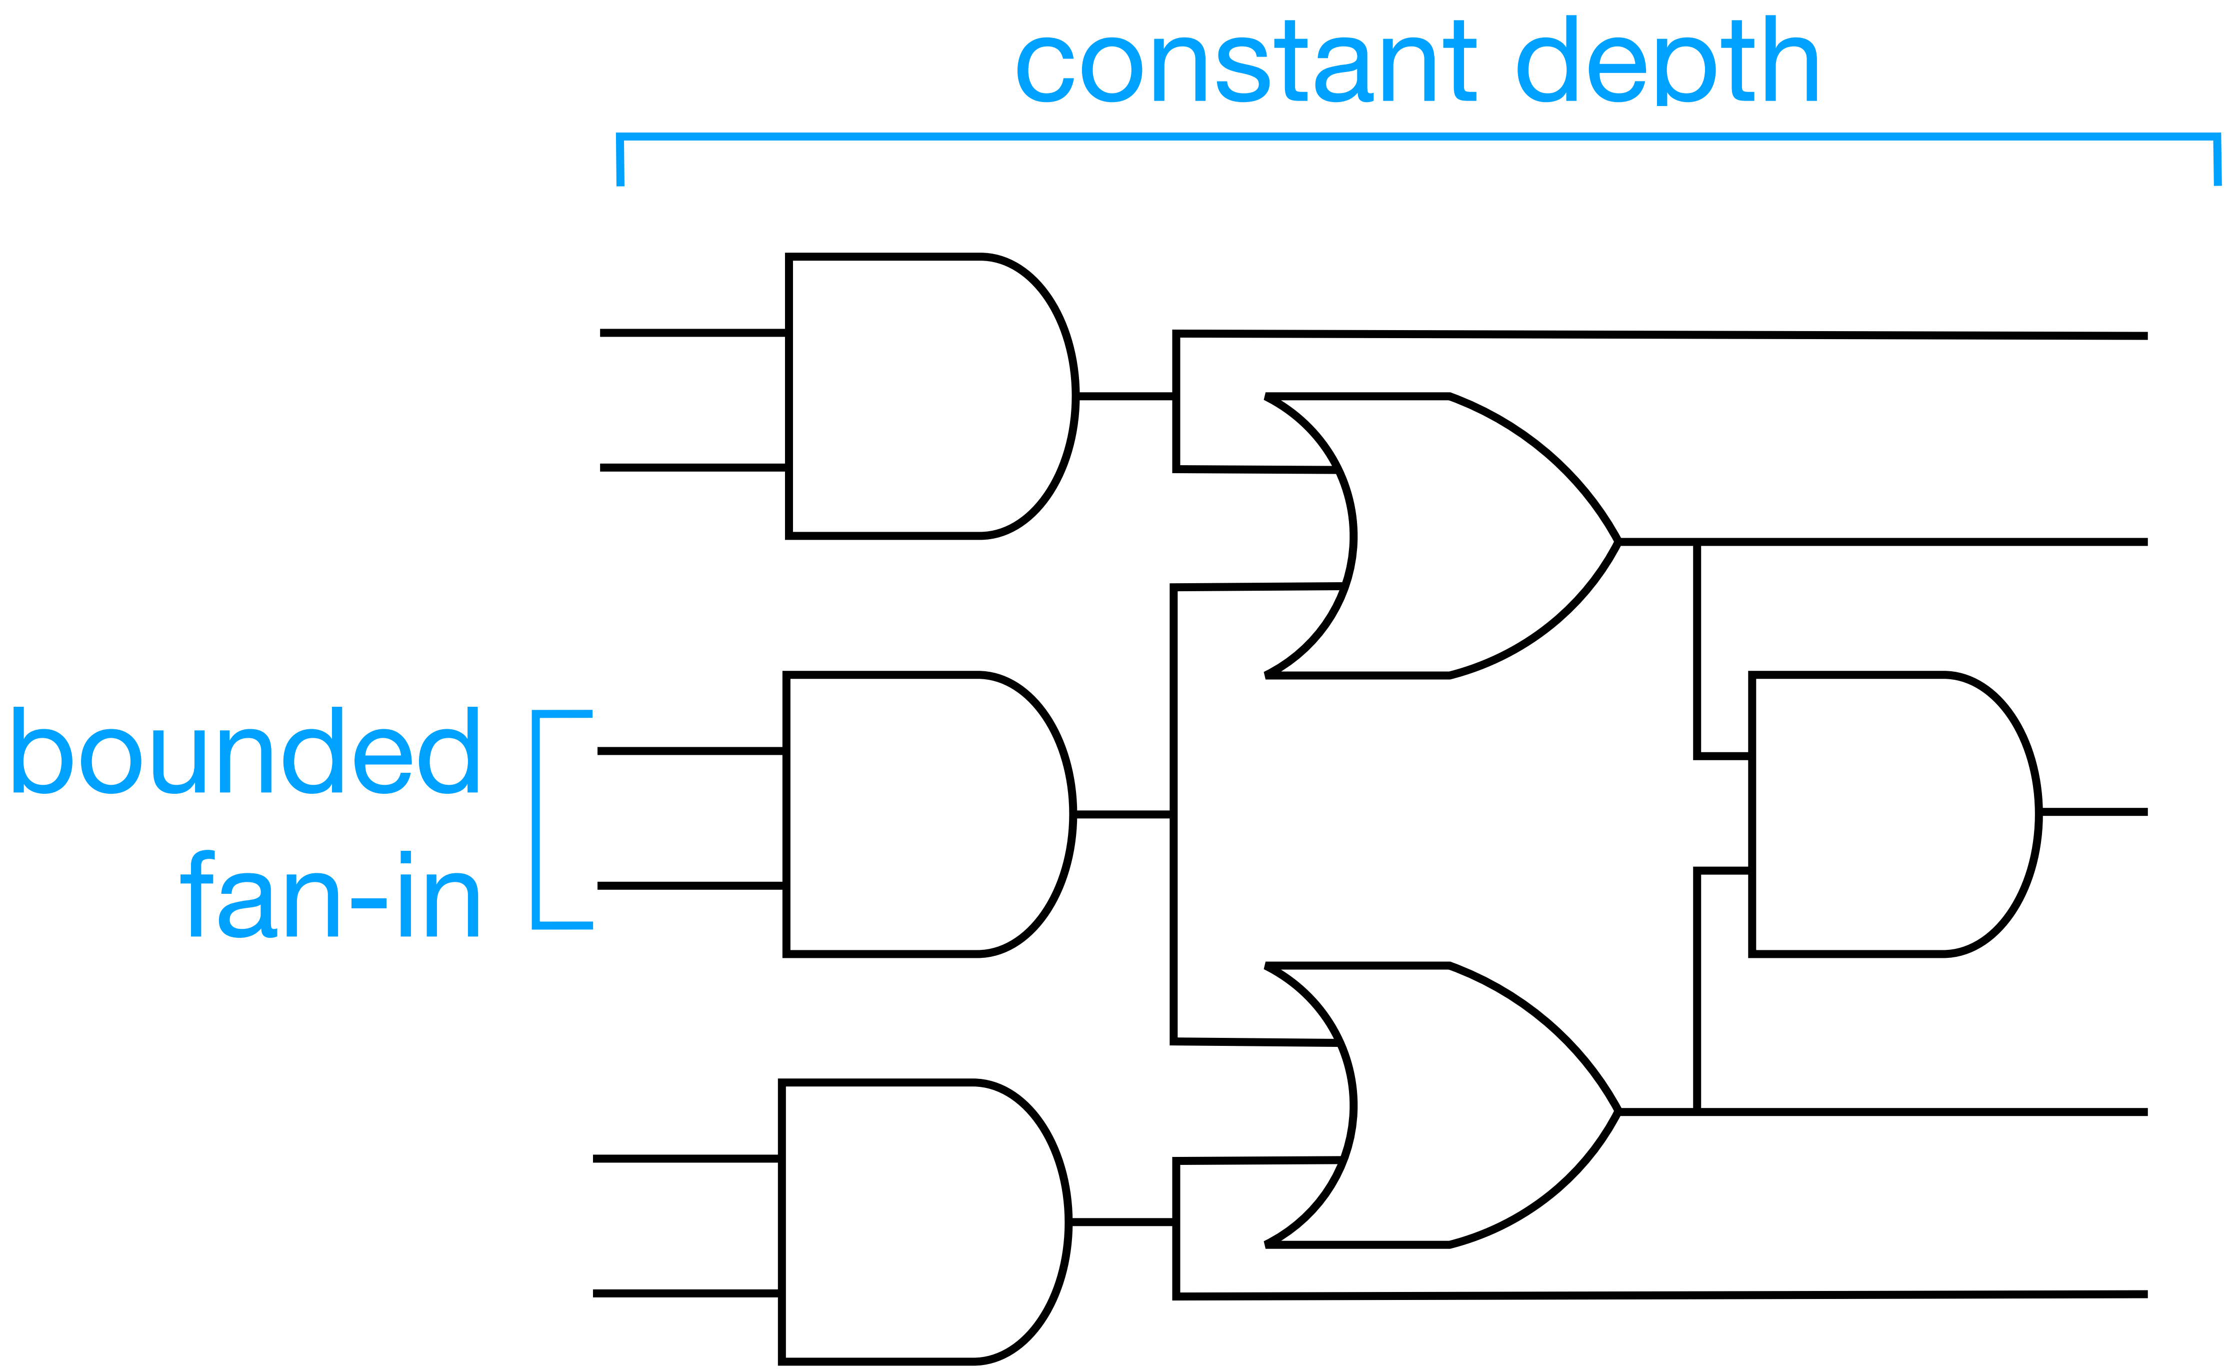

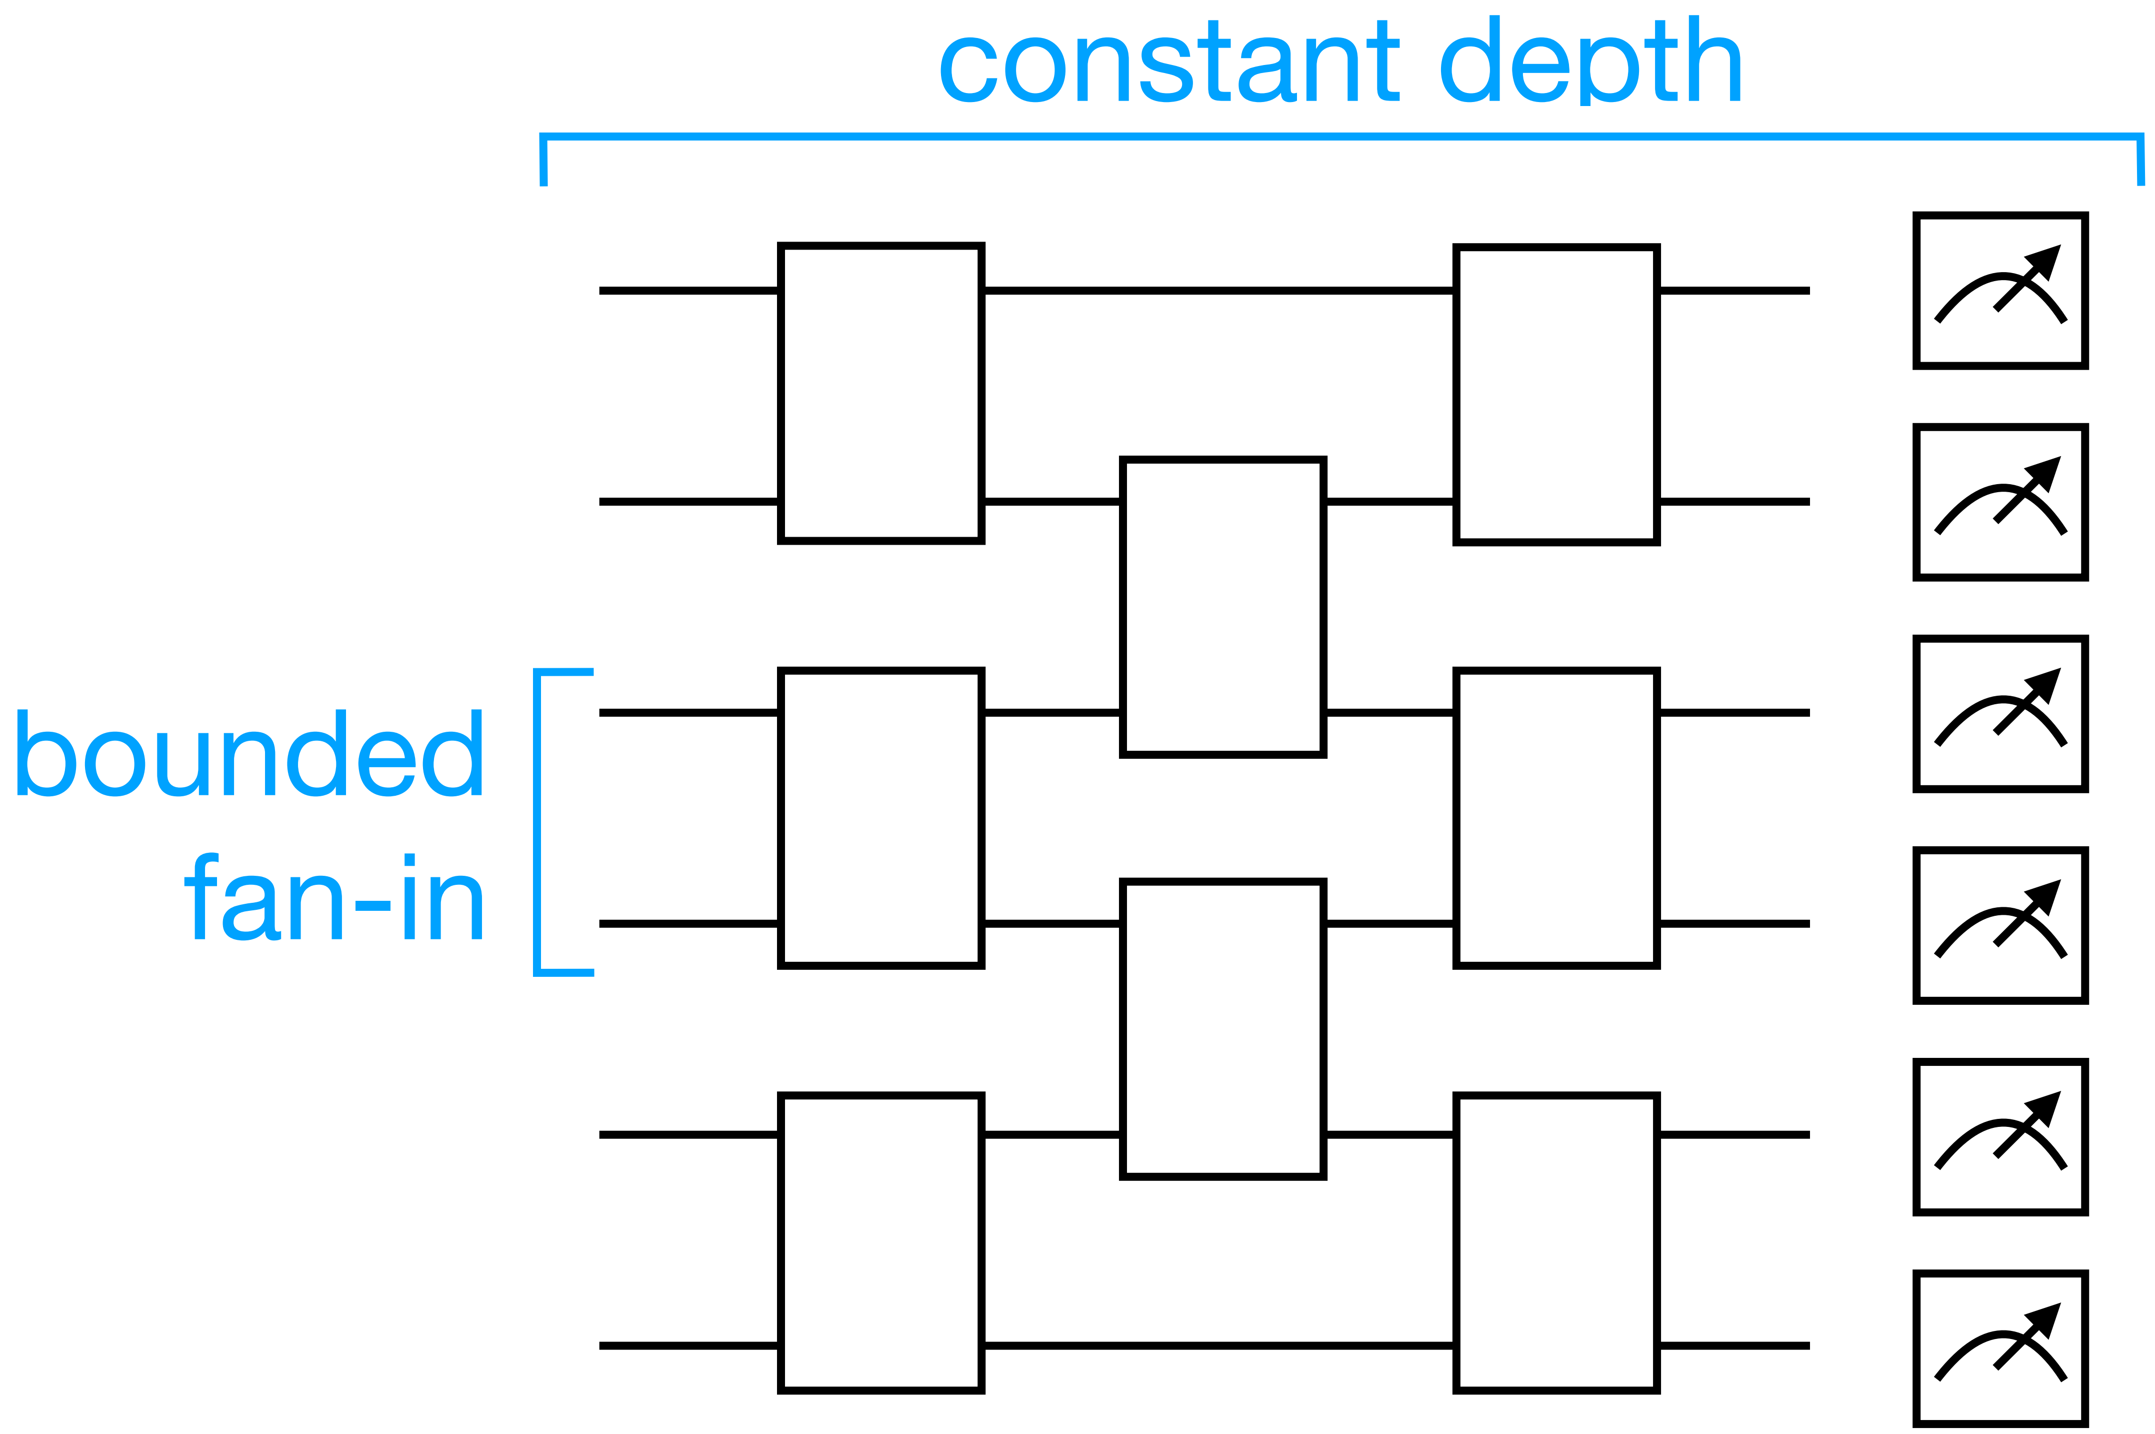

Constant-depth circuit separations

Theorem [Bravyi, Gosset, König 18]**:** Constant-depth quantum circuits can solve a problem that cannot be solved by **bounded fan-in** constant-depth circuits with AND, OR, and NOT gates.

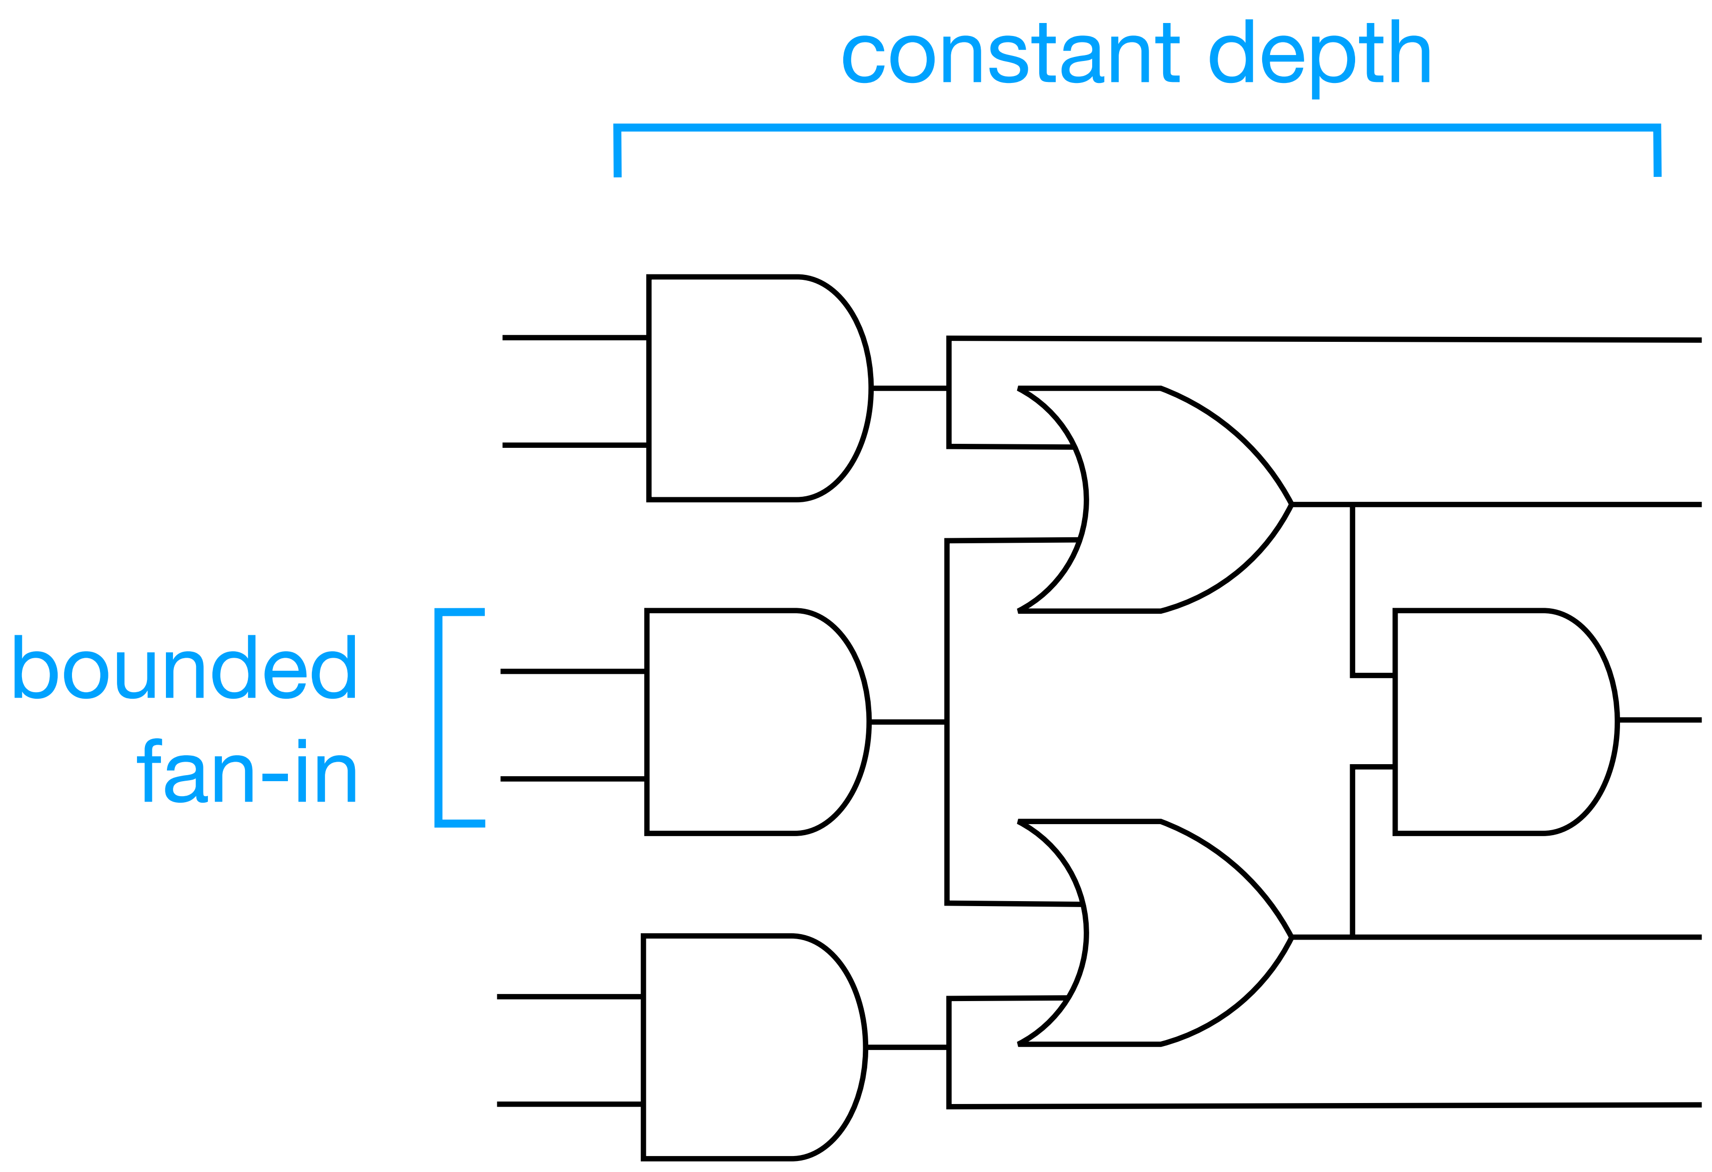

NC^0

$\not\supseteq$

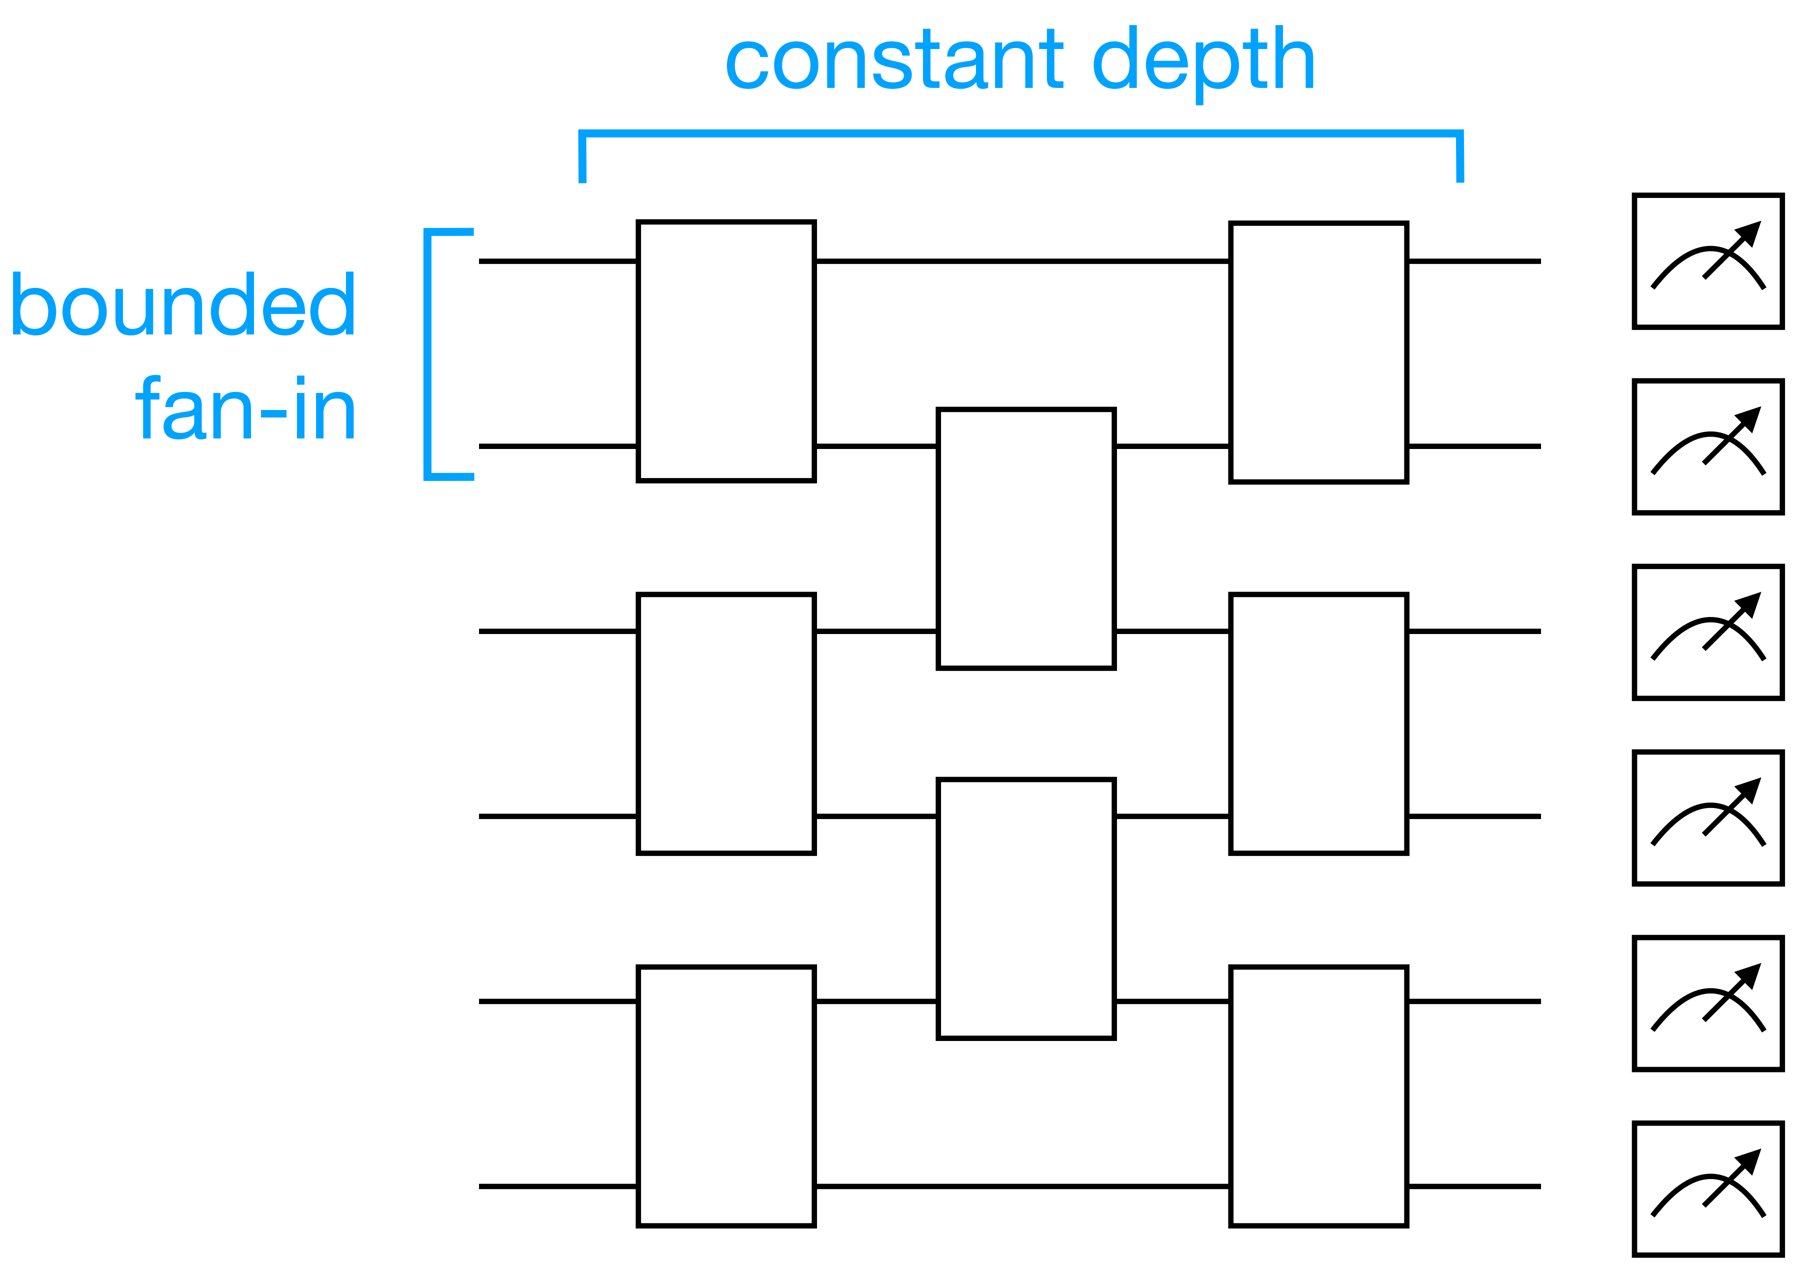

QNC^0

Constant-depth circuit separations

Theorem [Bene Watts, Kothari, Schaeffer, Tal 19]**:** Constant-depth quantum circuits solve a problem that cannot be solved by **unbounded fan-in** constant-depth circuits with AND, OR, and NOT gates.

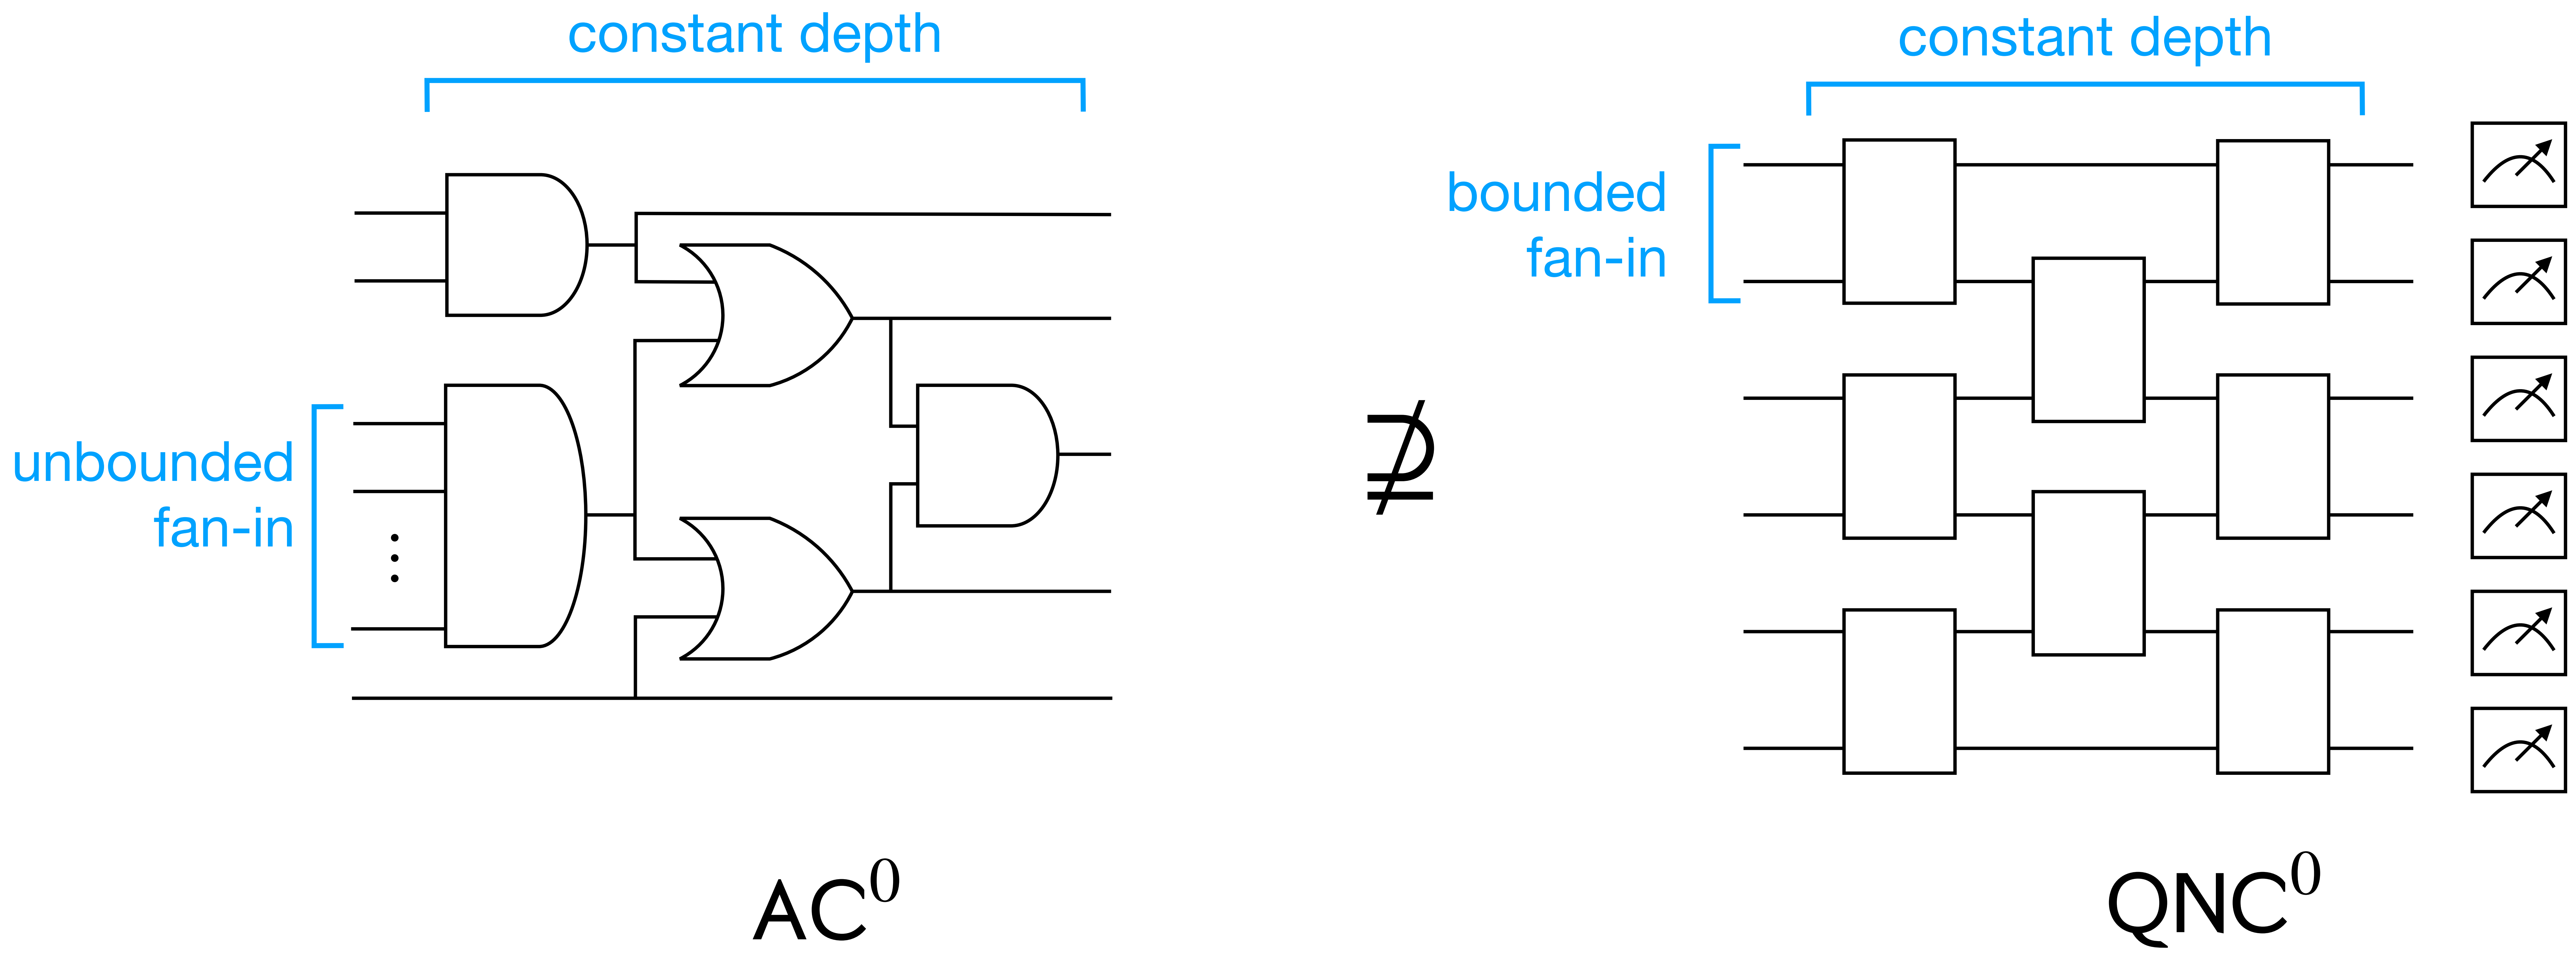

Most common types of problems

Sampling

Input: $x \in \{0,1\}^n$

Output: $y \sim \mathscr{D}_x$

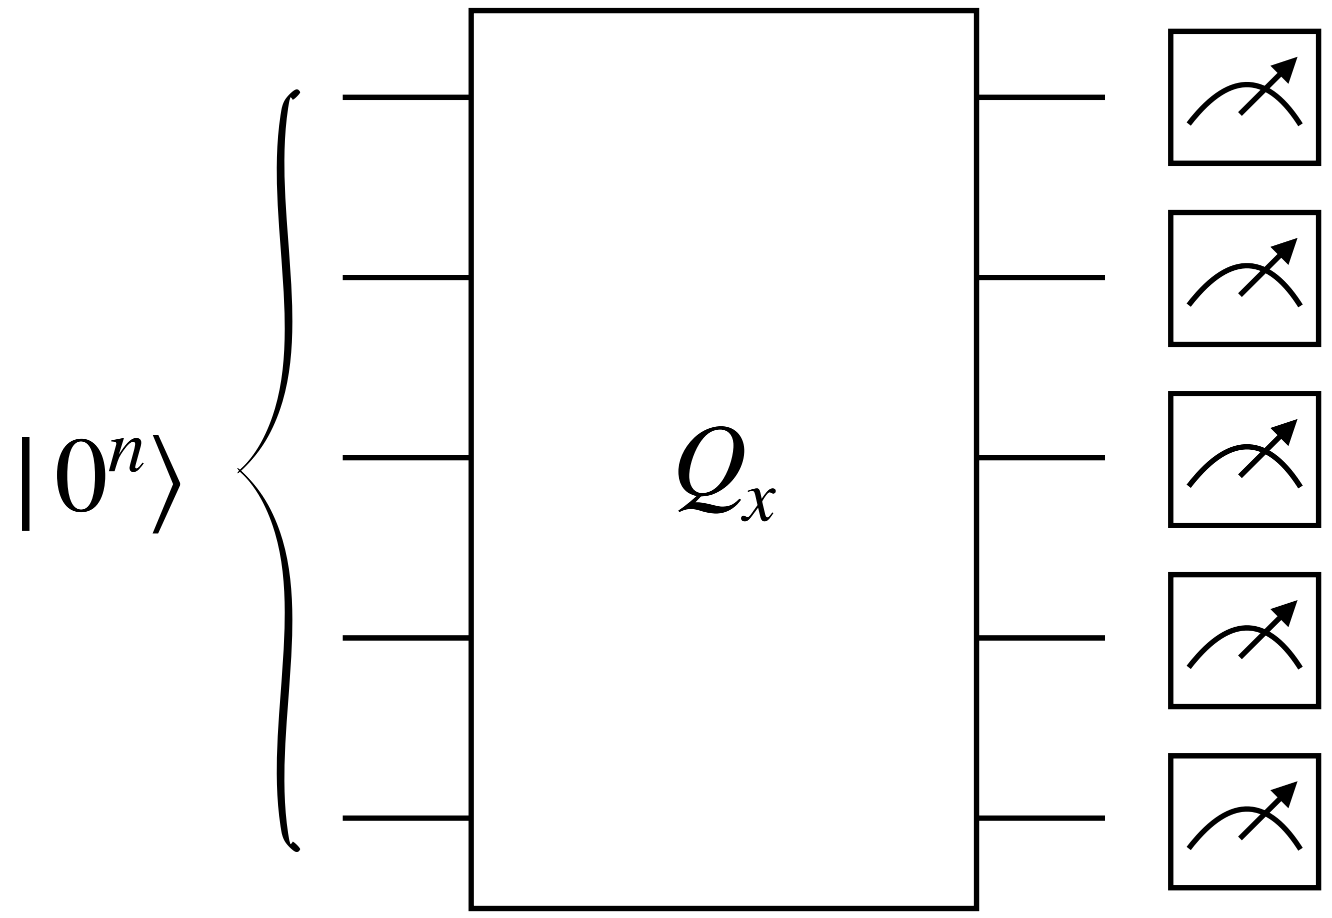

Quantum supremacy using a programmable superconducting processor

[Arute, et al. Nature 2019]

Relation

Input: $x \in \{0,1\}^n$

Output: $y \in \mathrm{Support}(\mathscr{D}_x)$

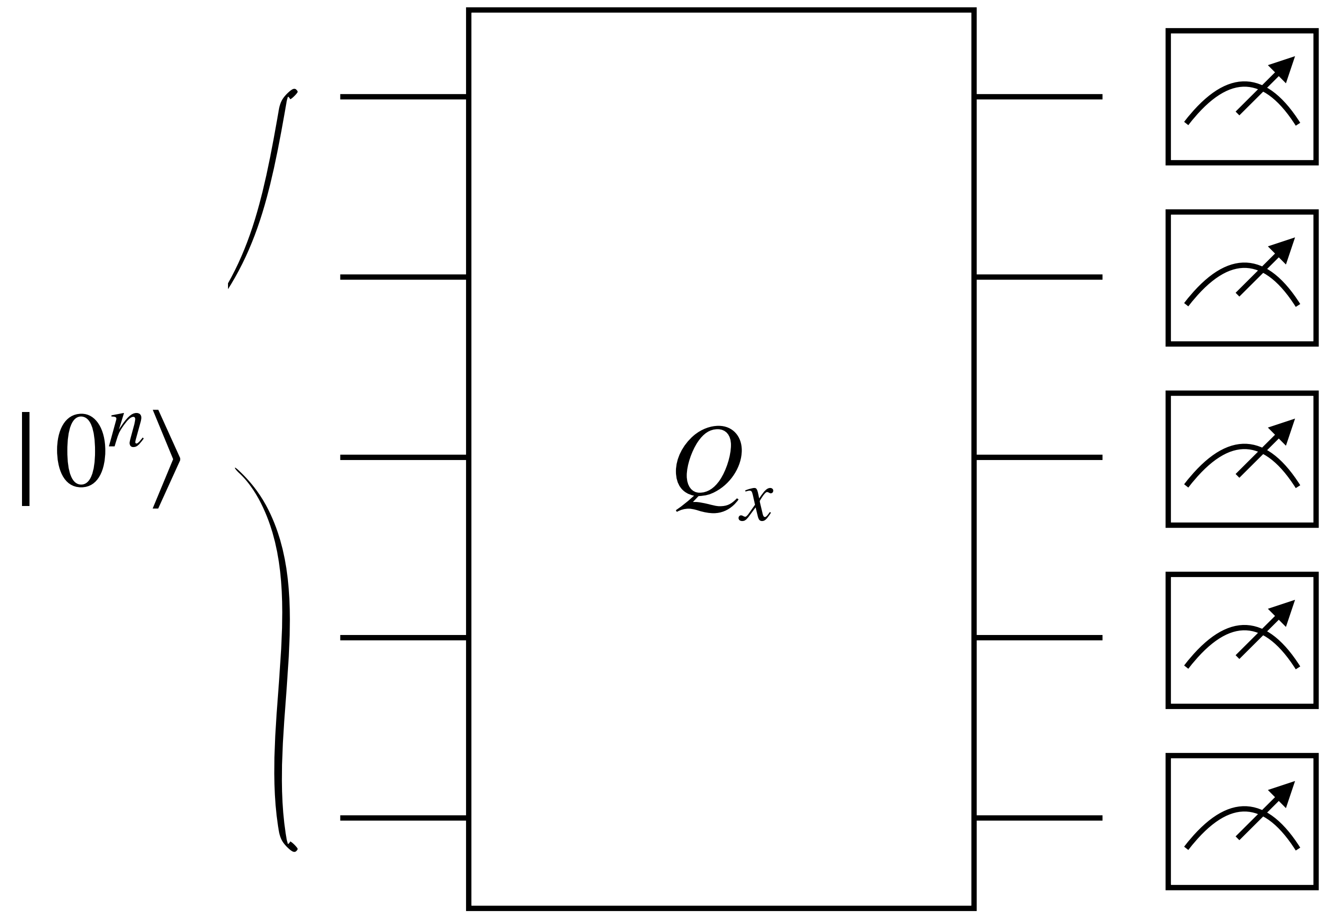

Quantum advantage with shallow circuits

[BGK. Science 2018]

Quantum circuits that depend on the input

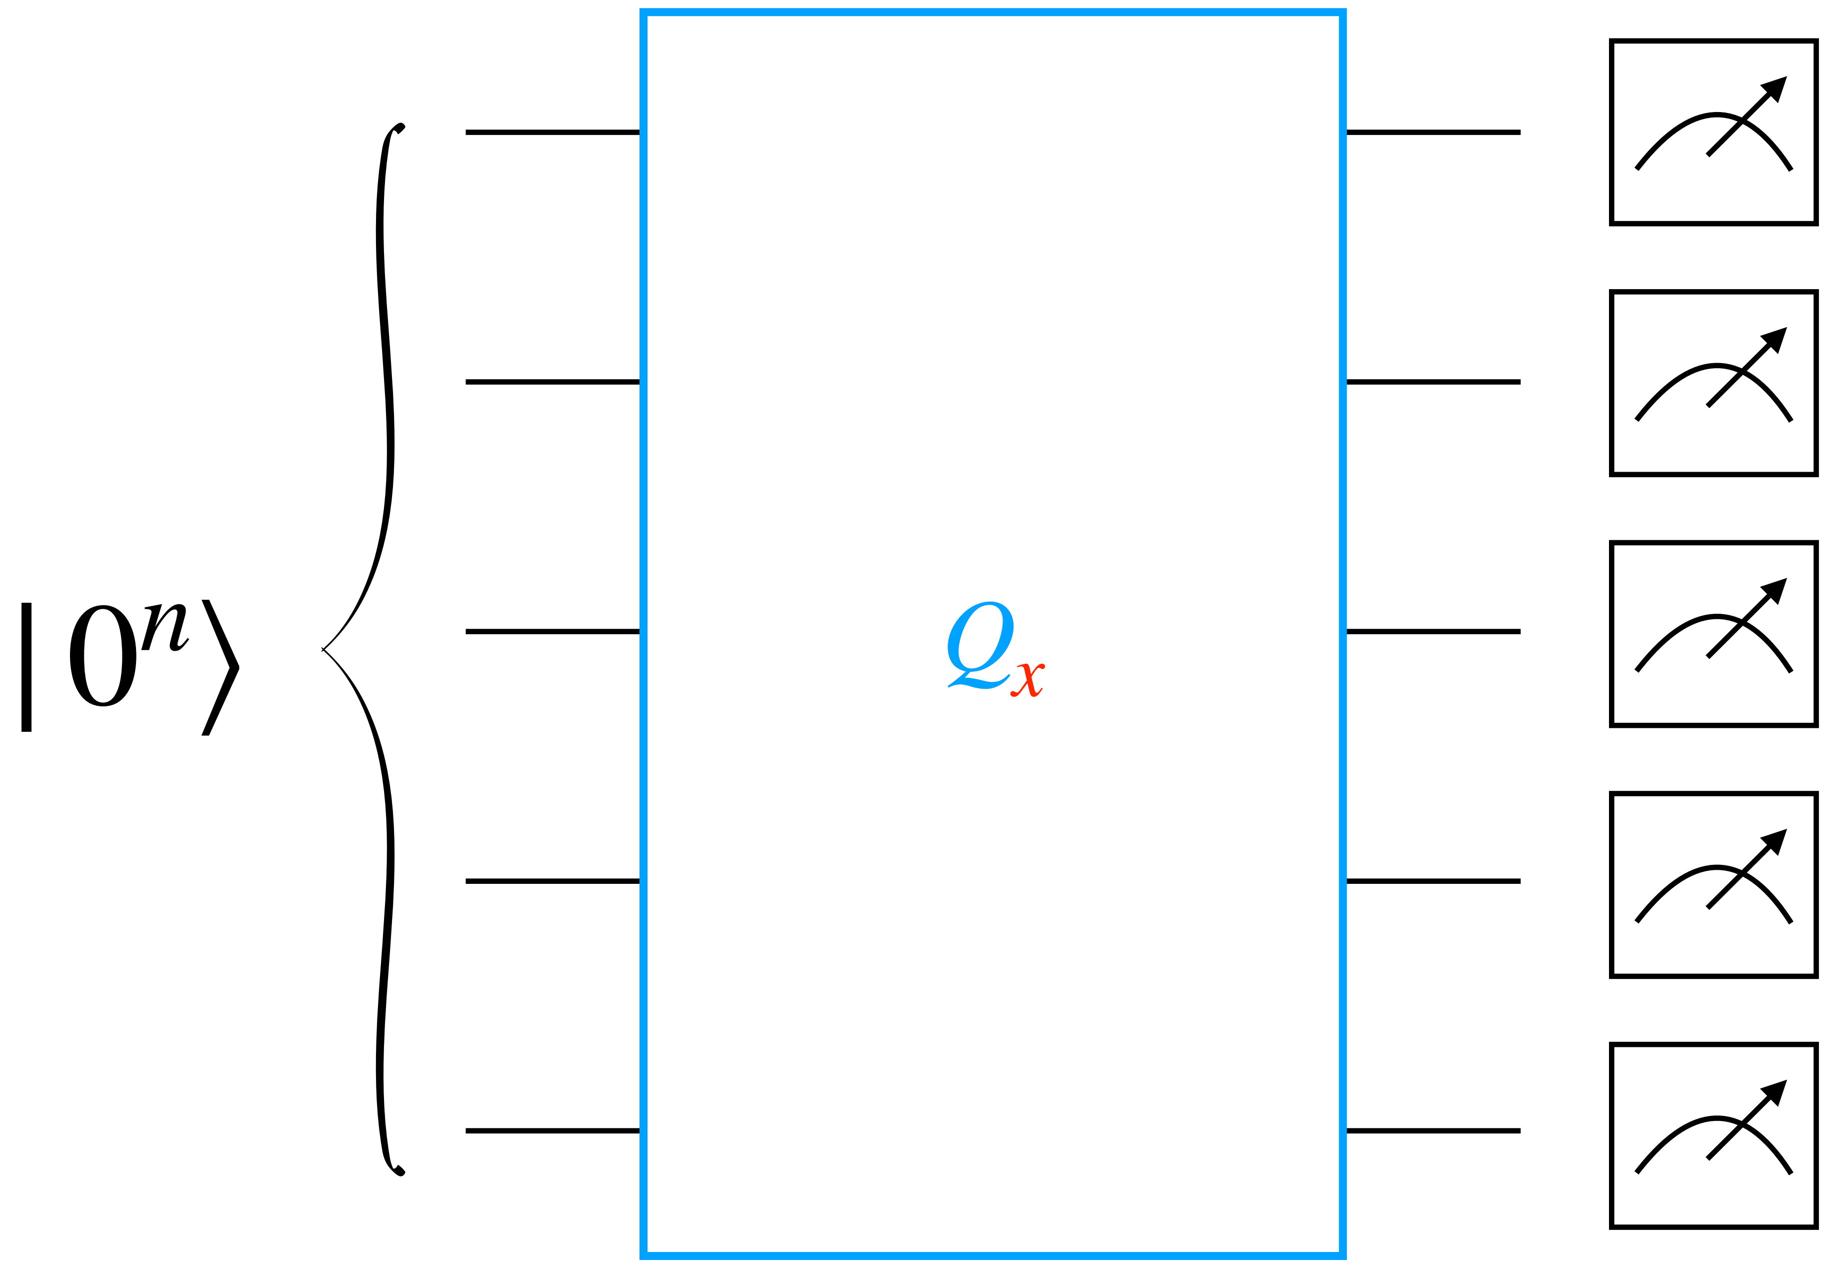

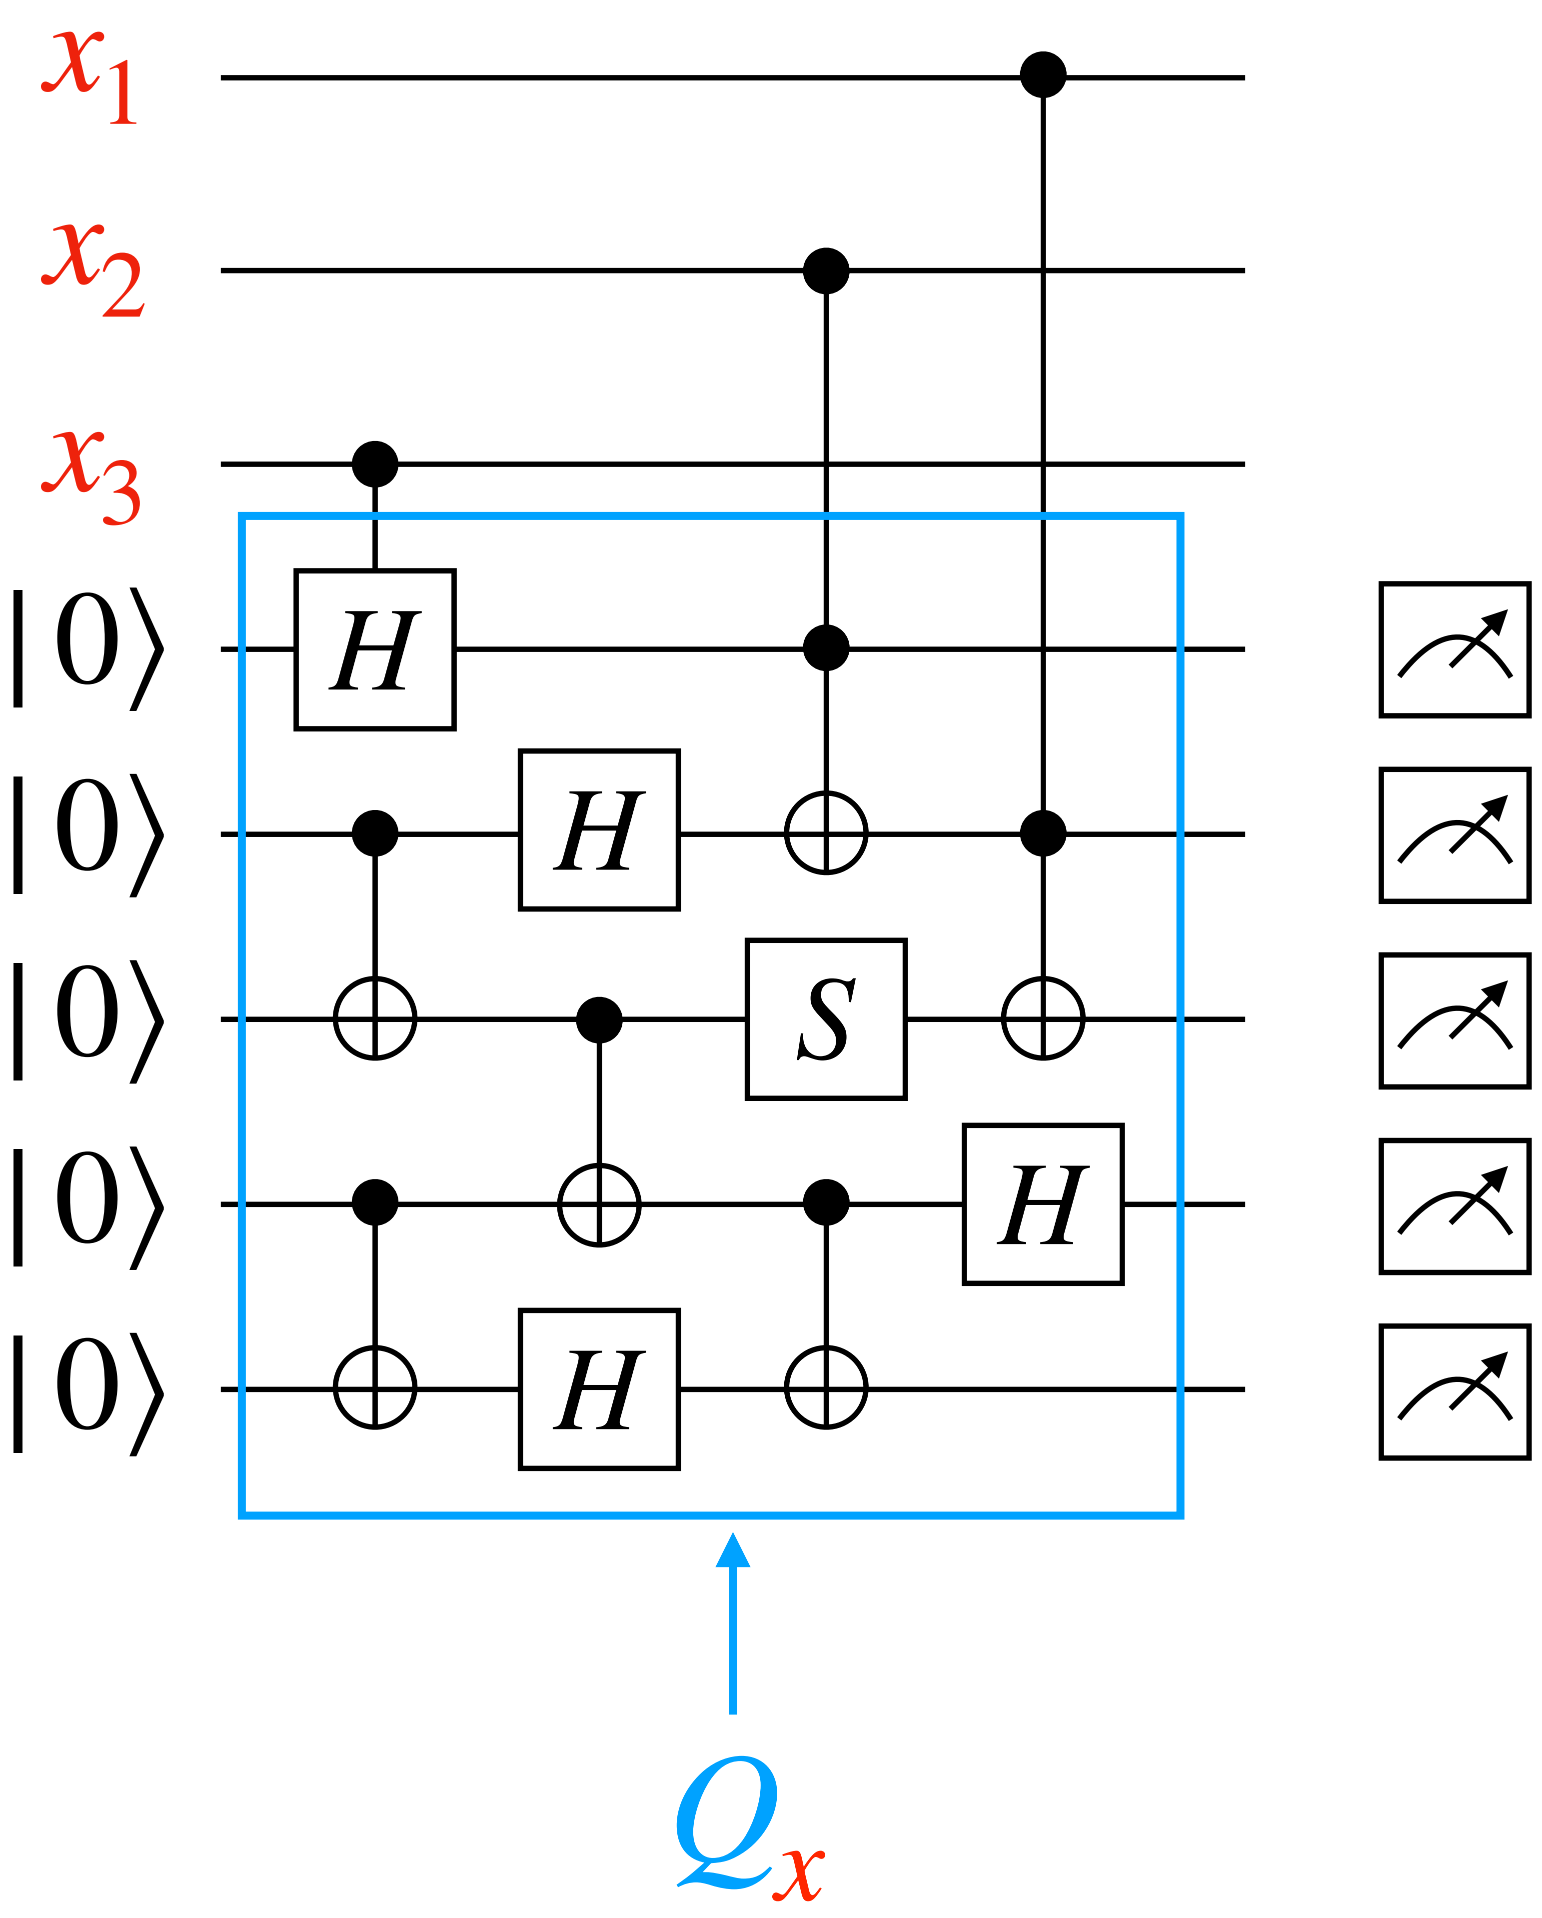

How do relation and sampling problems compare?

Observation: Relation problems are "easier" than sampling problems

→ Every circuit to sample immediately solves the corresponding relation problem

Theorem: Relation = Sampling for constant-depth Clifford circuits

Distribution problems

Distribution

Input: 1^n for some $n \in \mathbb{N}$

Output: $y \sim \mathscr{D}_n$

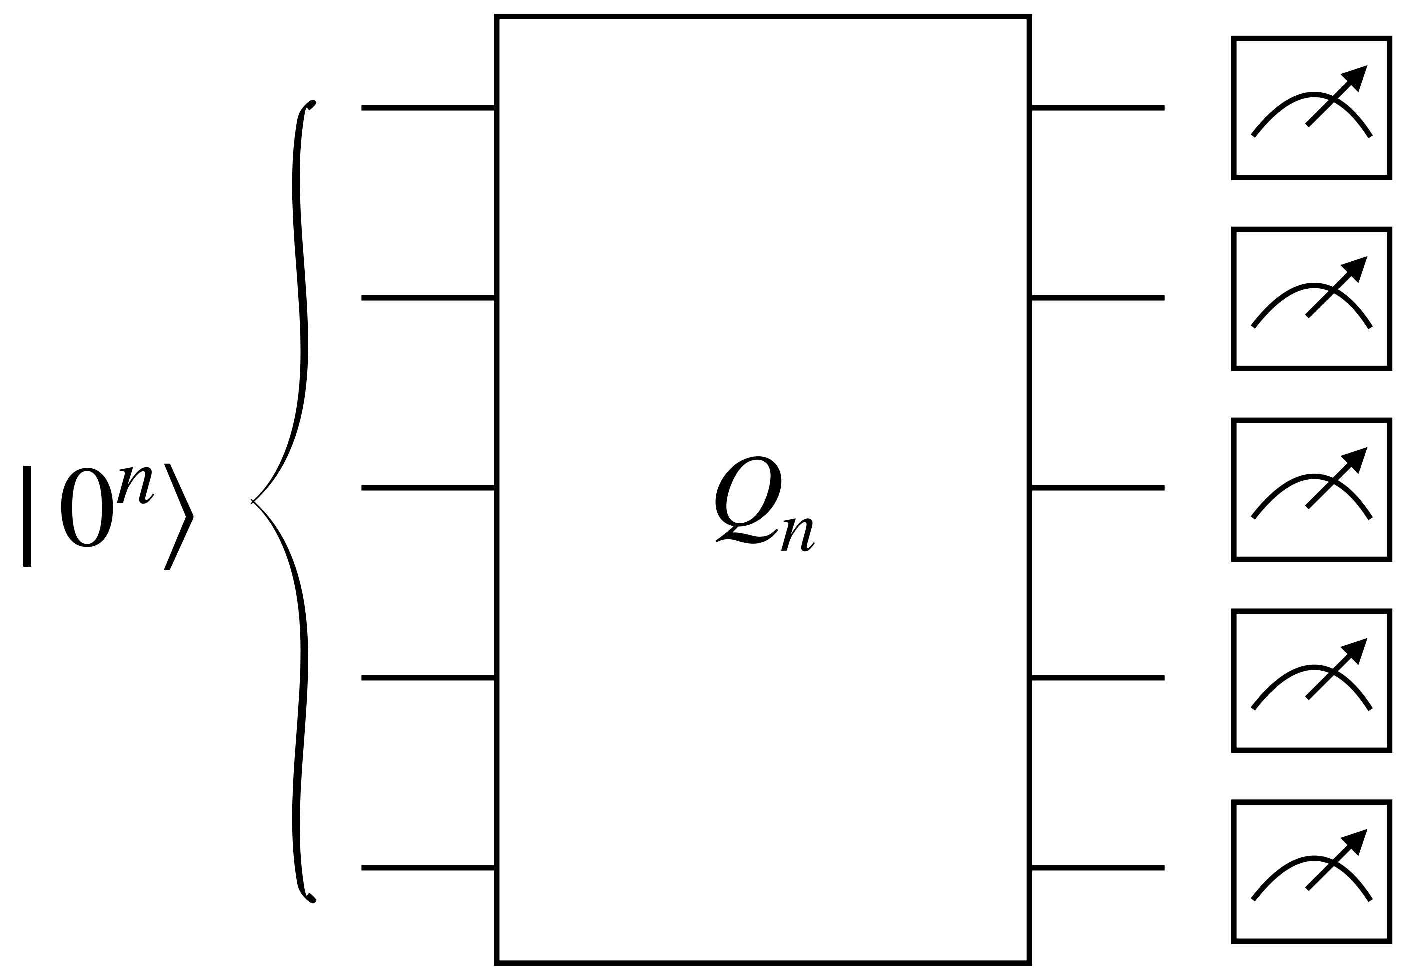

Classical Distribution

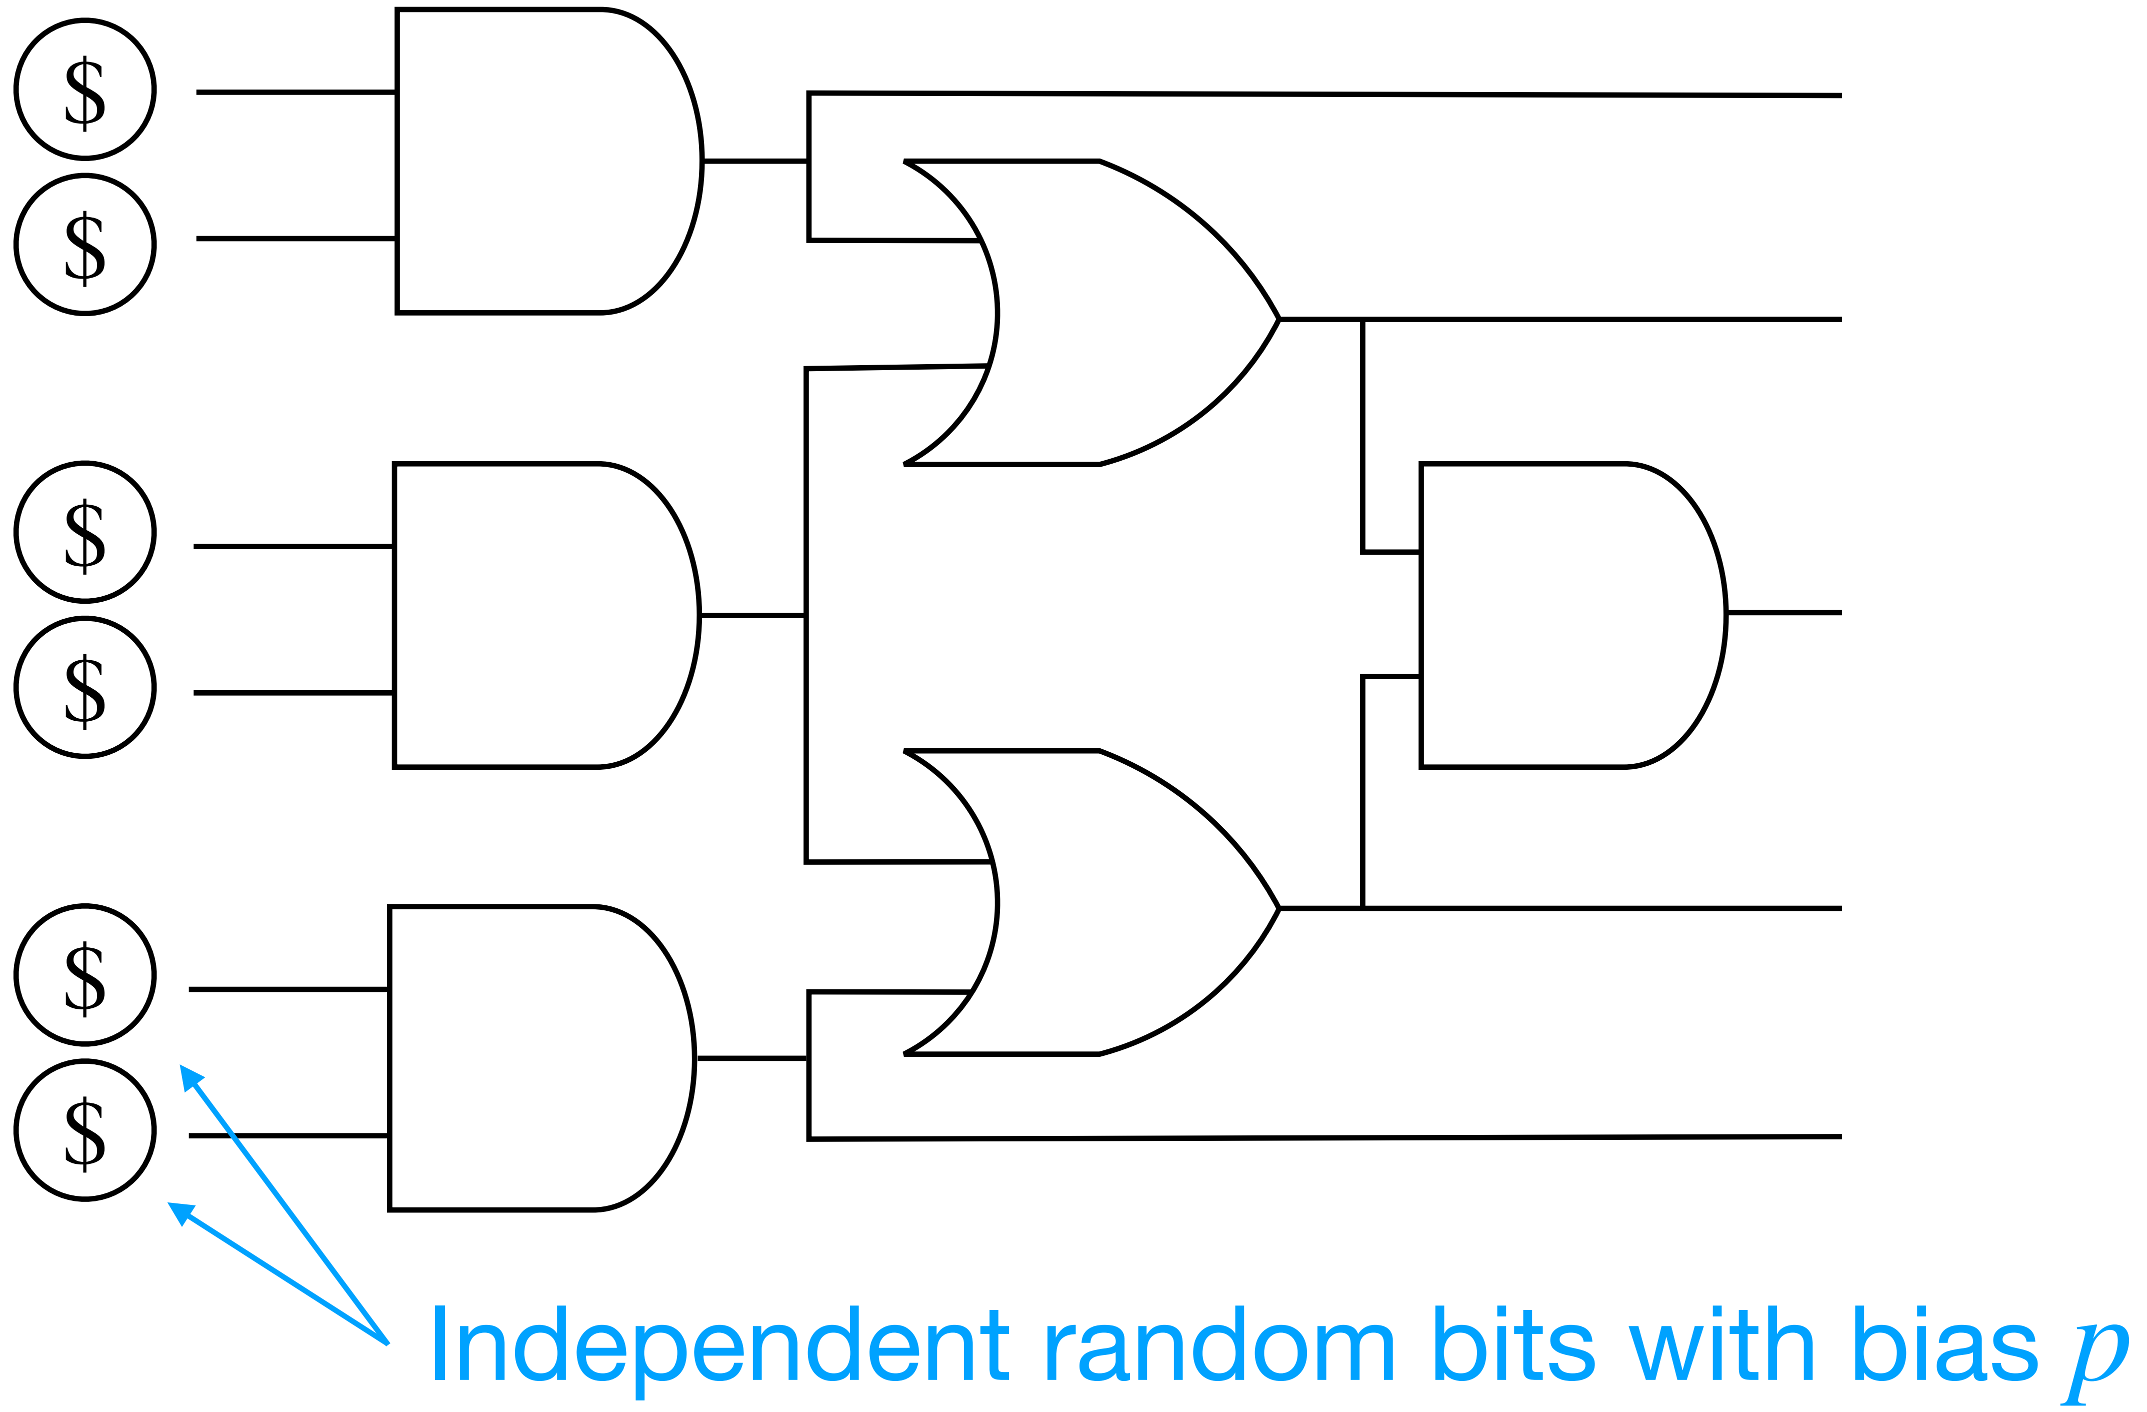

Question: Can we obtain quantum advantage for a distribution problem?

Prior work on distributional separations

Theorem [Parham, Bene Watts 23]**:** $\mathsf{distQNC}^0 \not\subseteq \mathsf{distNC}^0$

- Caveat 1: classical circuit needs $\mathcal{O}(n)$ bound on the number of ancillas
- Caveat 2: Requires a more-or-less arbitrary quantum gate set

Theorem [Viola 23, KOW 24]**:** $\mathsf{distQNC}^0 \not\subseteq \mathsf{distNC}^0$

- Hard Distribution: The (1/3)-biased distribution
- Caveat: only hard if your classical circuit doesn't get biased coins

Main theorem

Theorem [GKMOW 25]**:** $\text{distQNC}^0 \not\subseteq \text{distNC}^0$ (but hopefully better)

- Discrete gate set: Hadamard, controlled-Phase, Toffoli

  Implication: Single-qubit marginals are sampleable with NC^0 circuits

- Geometrically local

  Implication: Could implement the quantum circuit on current hardware

- Negligible overlap: $1 - e^{\Omega(n)}$

  Implication: Parallel repetition works as you expect

Theorem ingredients

Lower bound: Find distribution that cannot be sampled in NC^0

Upper bound: Show that distribution *can* be sampled in QNC^0

Simplification for this talk: Allow certain "quantum advice" states

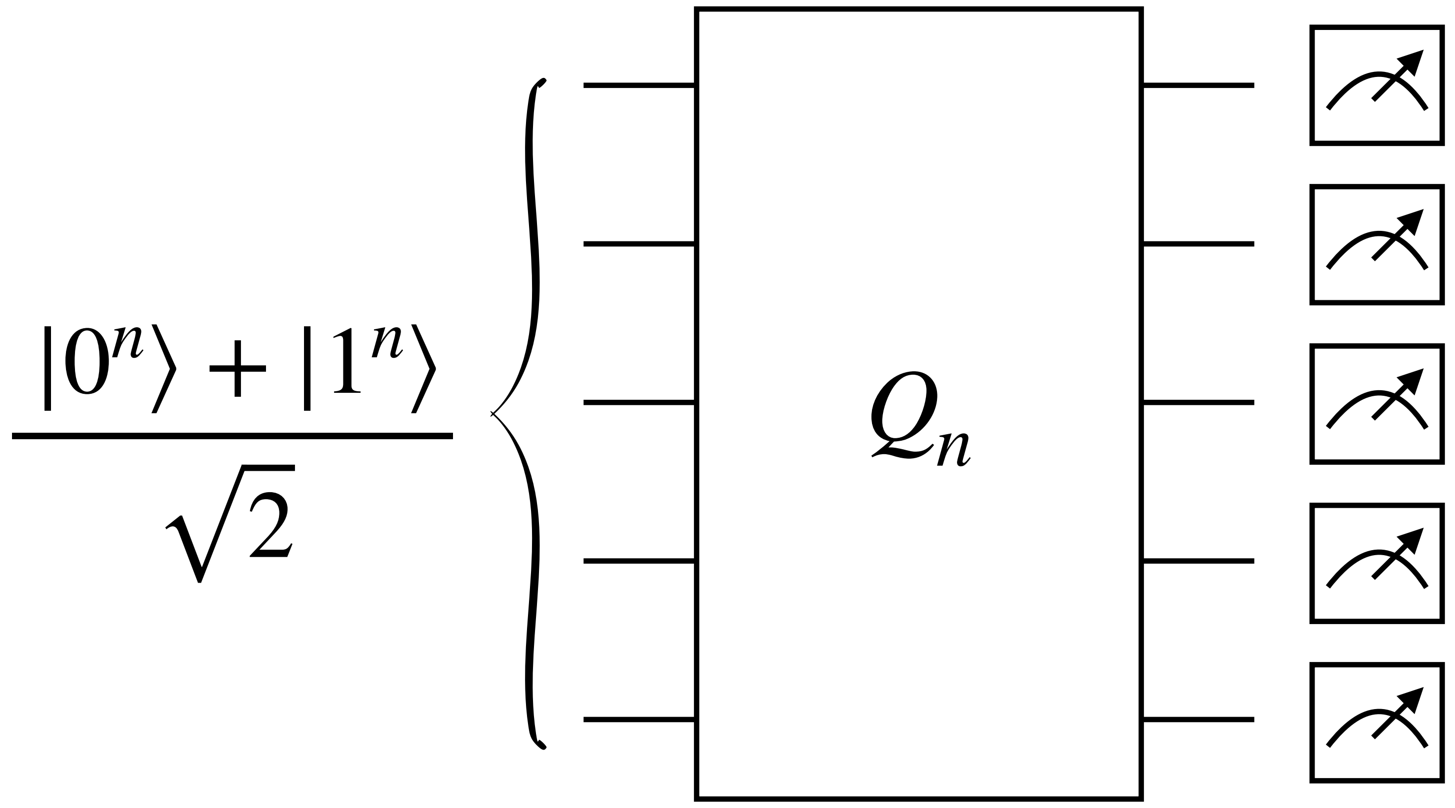

Theorem: $\mathrm{distQNC}^0/$ 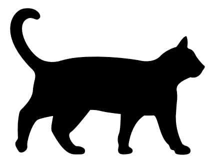 $\not\subseteq \mathrm{distNC}^0$

Creating hard distributions in shallow quantum depth

Why are distributional separations hard to prove?

Reasonable idea: Start with a function $f\colon \{0,1\}^n \to \{0,1\}$ which is hard to compute, and consider the distribution of pairs $(x, f(x))$ where x is a uniformly random string.

→ Quintessential hard function: $\text{Parity}(x) = x_1 \oplus x_2 \oplus \cdots \oplus x_n$

Theorem: $\text{Parity} \notin \text{AC}^0$

→ More than we need!

Quantum circuits can sample from even parity strings

$$\frac{|0^n\rangle + |1^n\rangle}{\sqrt{2}}$$

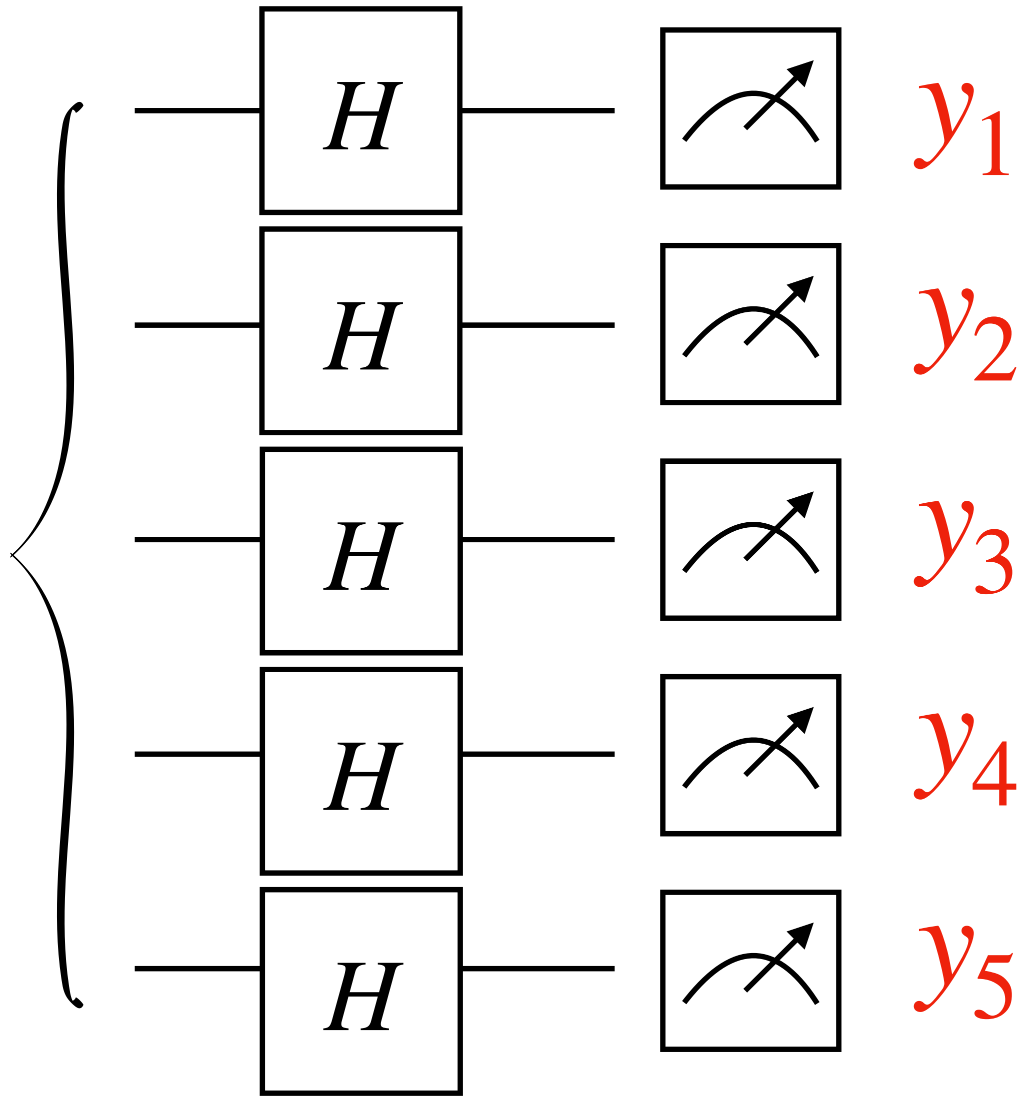

$$y_1 \oplus y_2 \oplus y_3 \oplus y_4 \oplus y_5 = 0$$

$$y_1 \oplus y_2 \oplus y_3 \oplus y_4 = y_5$$

Think of y_5 as the parity of the other bits

Takeaway: QNC^0/ 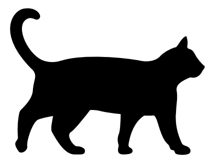 circuit to prepare $\frac{1}{\sqrt{2^n}} \sum_{x \in \{0,1\}^n} |x, \text{Parity}(x)\rangle$.

Hard function problem ≠ Hard distribution problem

Fact: The $(x, \text{Parity}(x))$ distribution is sampleable in NC^0.

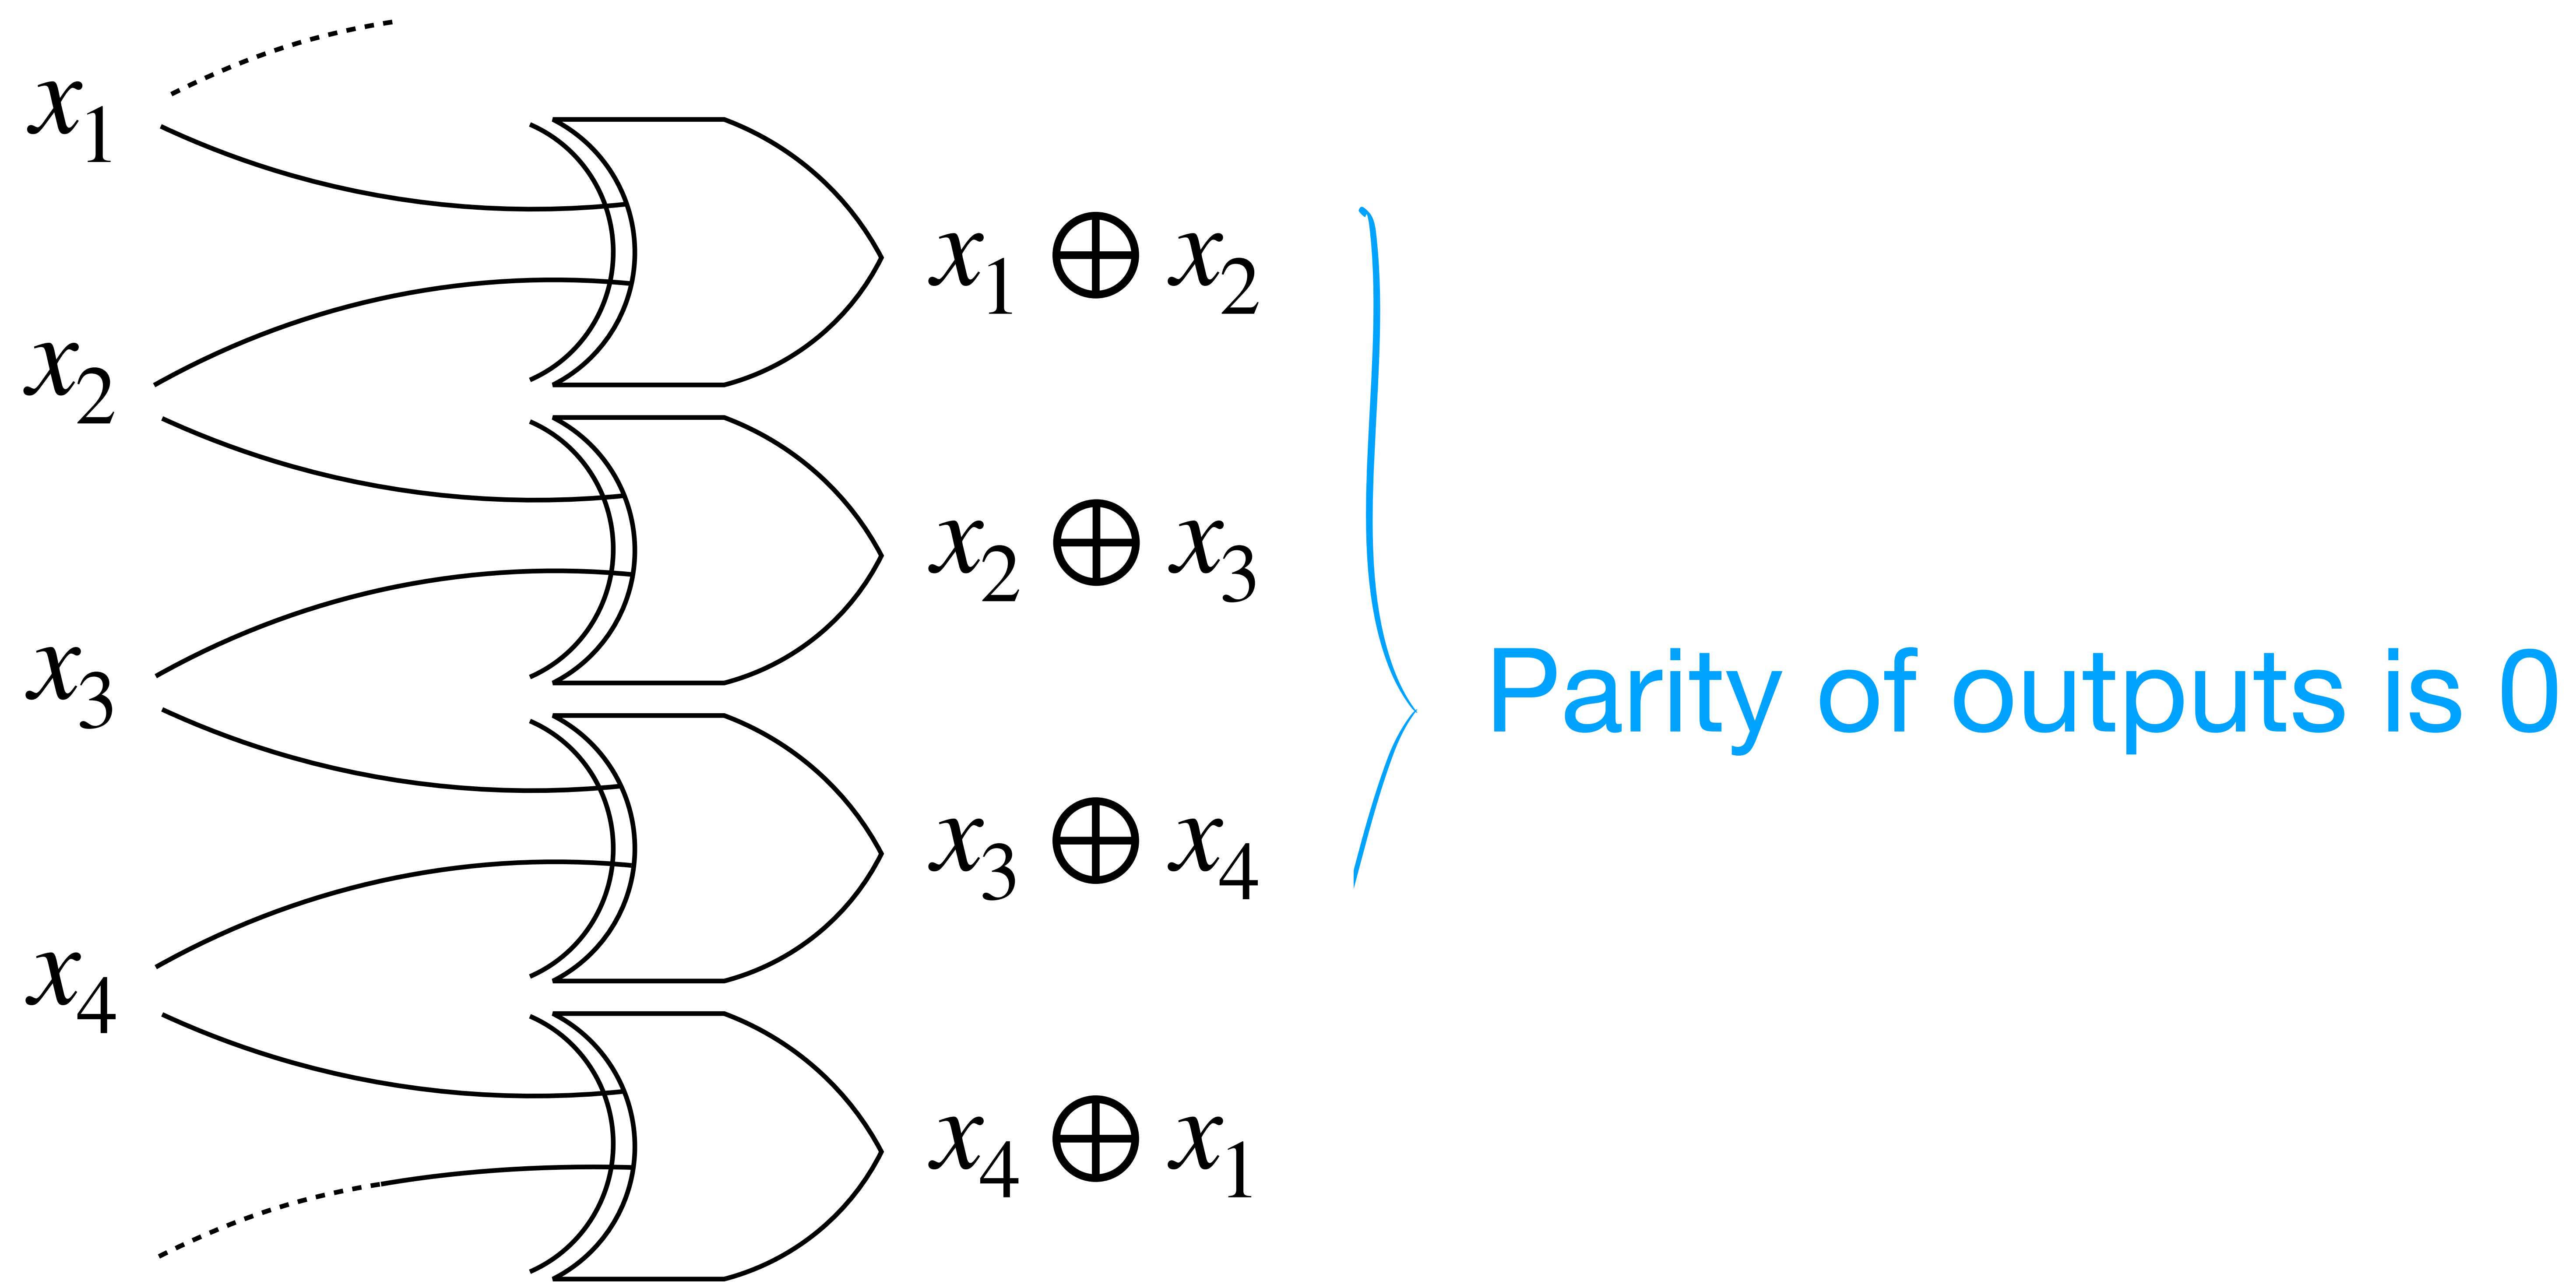

What went wrong?

Key fact: Flipping a bit in the $(x, \text{Parity}(x))$ distribution didn't change the distribution

→ Follows from the fact that x is uniform

Modified reasonable idea: Consider the distribution of pairs $(x, \text{Parity}(x))$ where x is random *but not uniform*.

→ For example… x_i is drawn from the $(1/4)$-biased distribution

→ NC^0 circuits can't sample from this distribution!

→ But neither can QNC^0 circuits… ☹

Parity-Halving to the rescue

Parity-Halving Problem [WKST 18]**:**

Input: $x \in \{0,1\}^n$ such that Parity$(x) = 0$

Output: $y \in \{0,1\}^m$ such that Parity$(y) = \begin{cases} 0 & \text{if } |x| \equiv 0 \pmod{4} \\ 1 & \text{if } |x| \equiv 2 \pmod{4} \end{cases}$

- Specially designed to be solved by low-depth quantum circuits!
- The hardness for classical circuits depends on m

  If $m = \Omega(n^2)$, then Parity-Halving is in NC^0

  If $m = o(n^2)$, then Parity-Halving is not in NC^0 (or even AC^0)

QNC^0/ 🐈 circuit for Parity-Halving Problem

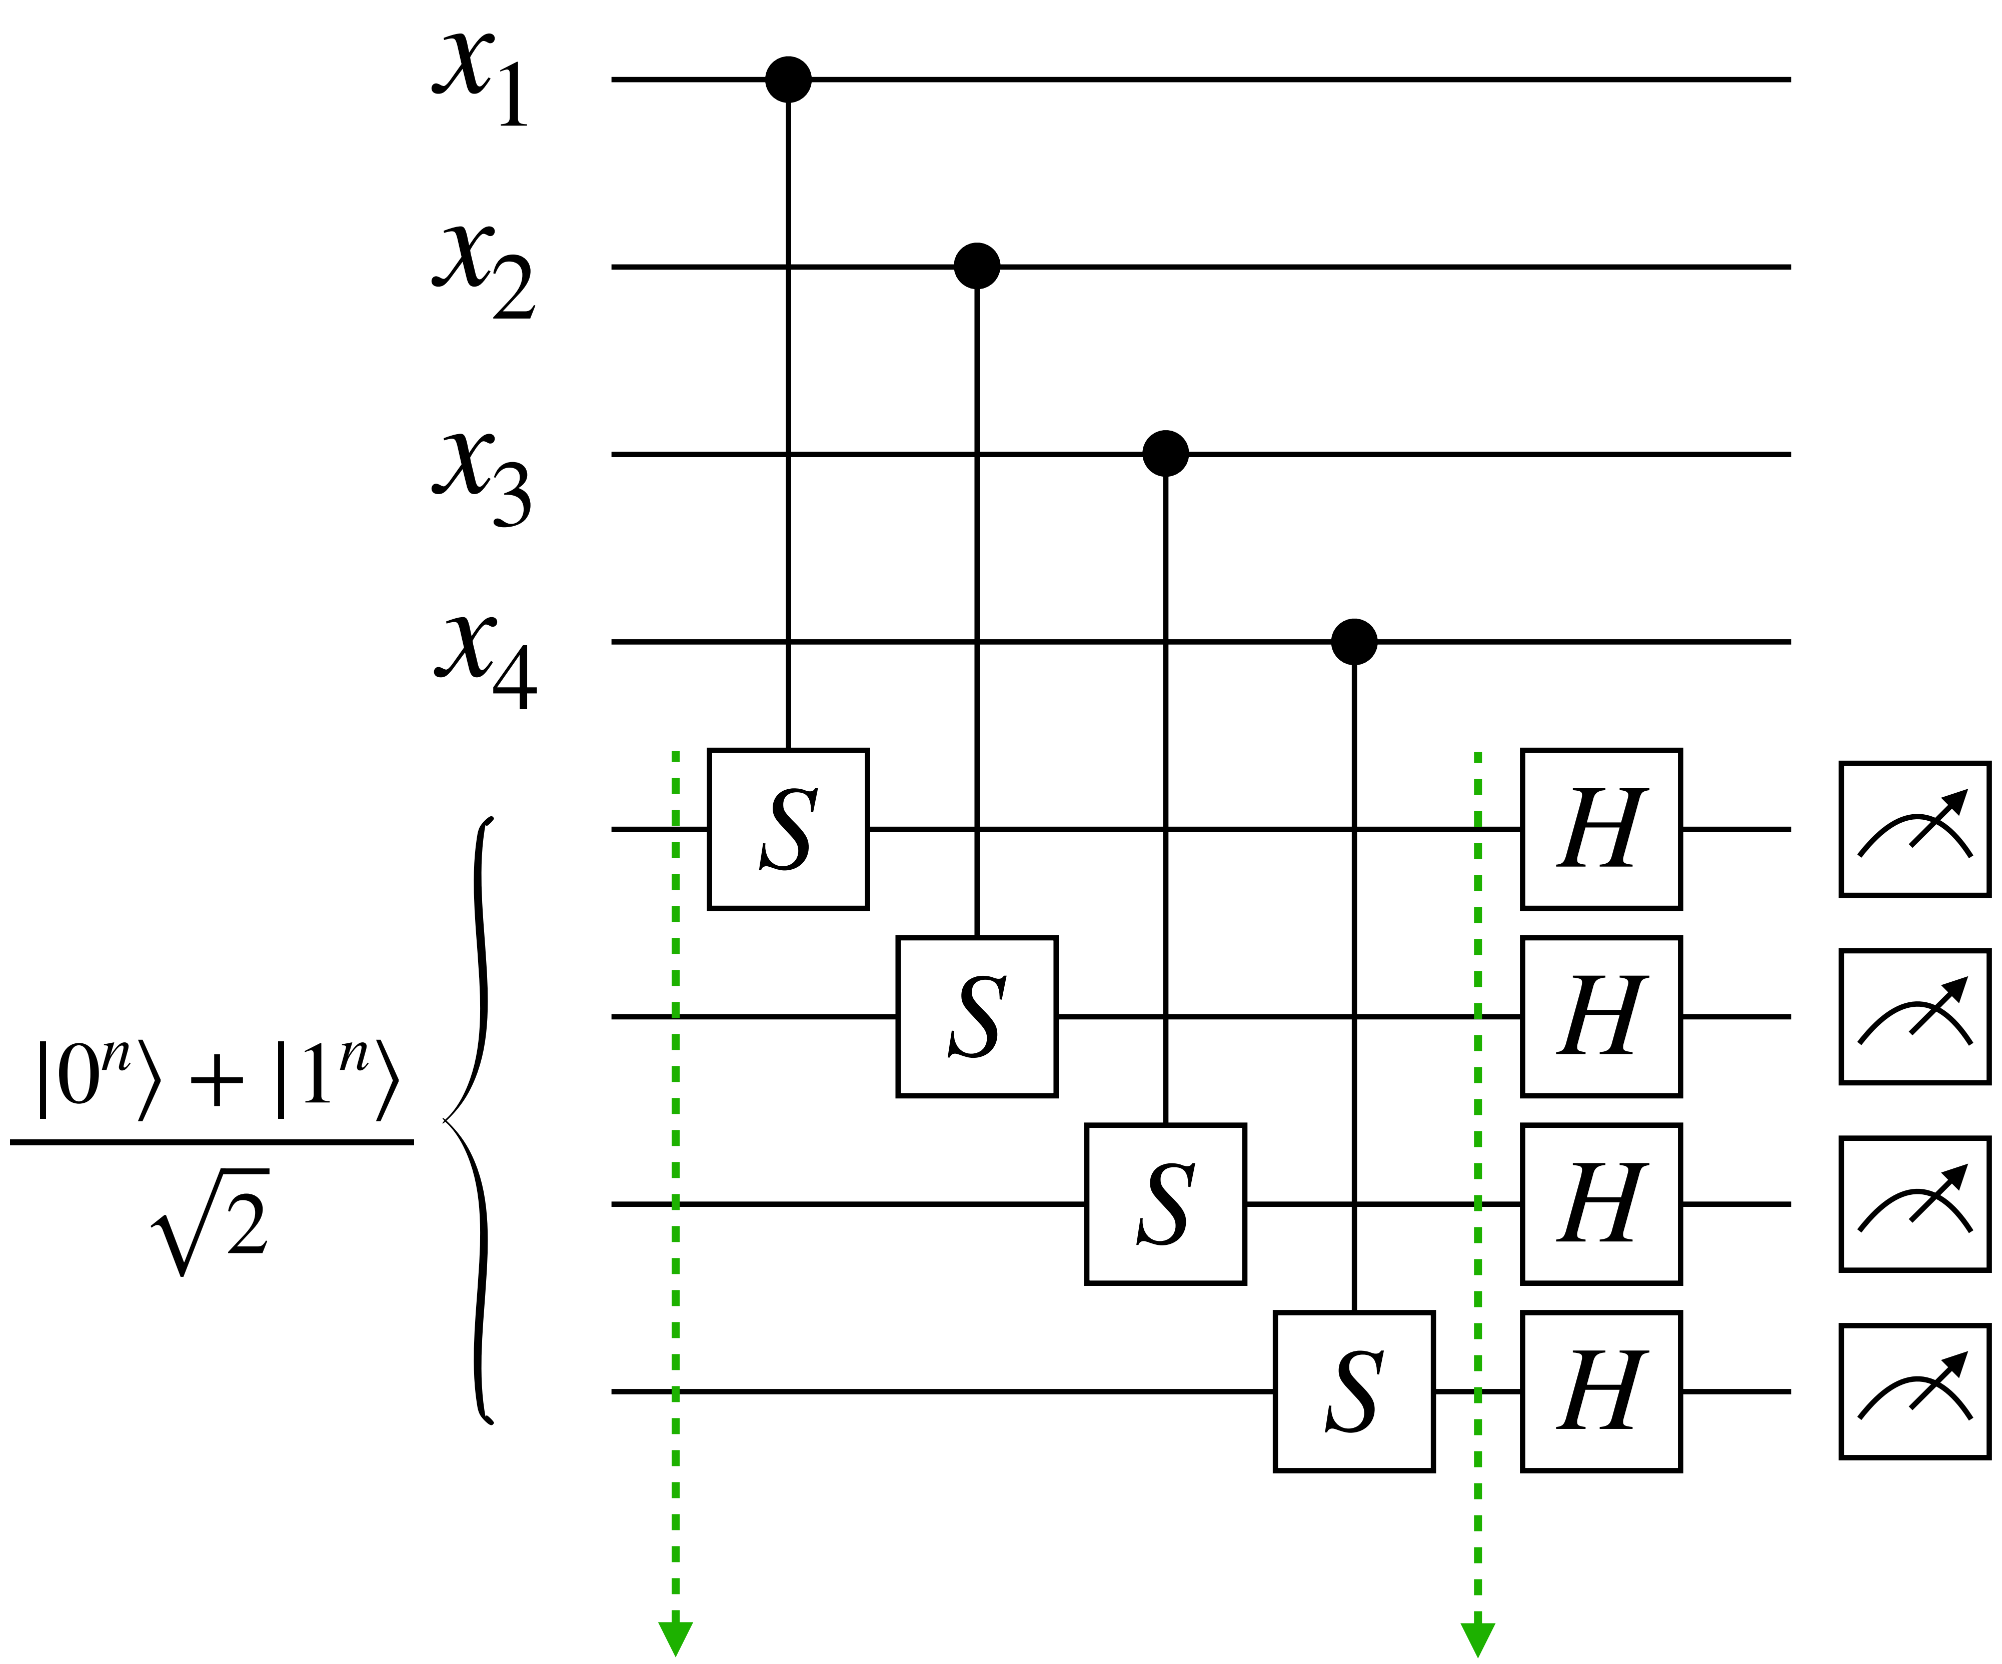

Recall:

$$H^{\otimes n}|🐈\rangle = \frac{1}{\sqrt{2^{n-1}}} \sum_{\text{Parity}(x)=0} |x\rangle$$

$$H^{\otimes n}|-🐈\rangle = \frac{1}{\sqrt{2^{n-1}}} \sum_{\text{Parity}(x)=1} |x\rangle$$

$$\frac{|0^n\rangle + |1^n\rangle}{\sqrt{2}} \rightarrow \frac{|0^n\rangle + i^{|x|}|1^n\rangle}{\sqrt{2}} = \frac{|0^n\rangle + (-1)^{|x|/2}|1^n\rangle}{\sqrt{2}} = \begin{cases} |🐈\rangle & \text{if } |x|/2 \equiv 0 \pmod 2 \\ |-🐈\rangle & \text{if } |x|/2 \equiv 1 \pmod 2 \end{cases}$$

Putting it all together

Most reasonable modified idea: Consider the distribution of pairs $(x, \mathsf{ParityHalving}(x))$ where each bit of x is $(1/4)$-biased, and $\mathsf{ParityHalving}(x)$ is uniform amongst valid solutions

- Are we done yet? Yes, but…

  Recall: $\mathsf{Parity}(x) = 0$ in the promise of the Parity-Halving problem

  Solution: Just run the quantum circuit on those inputs too!

- Also need to be able to generate $(1/4)$-biased bits

  Solution: Use Hadamard + Toffoli gates

Proving classical circuit lower bounds

Independent output bits imply low correlation

Intuitive (oversimplified) idea: Parity is sensitive to all of the input bits, so we shouldn't be able to independently toggle inputs

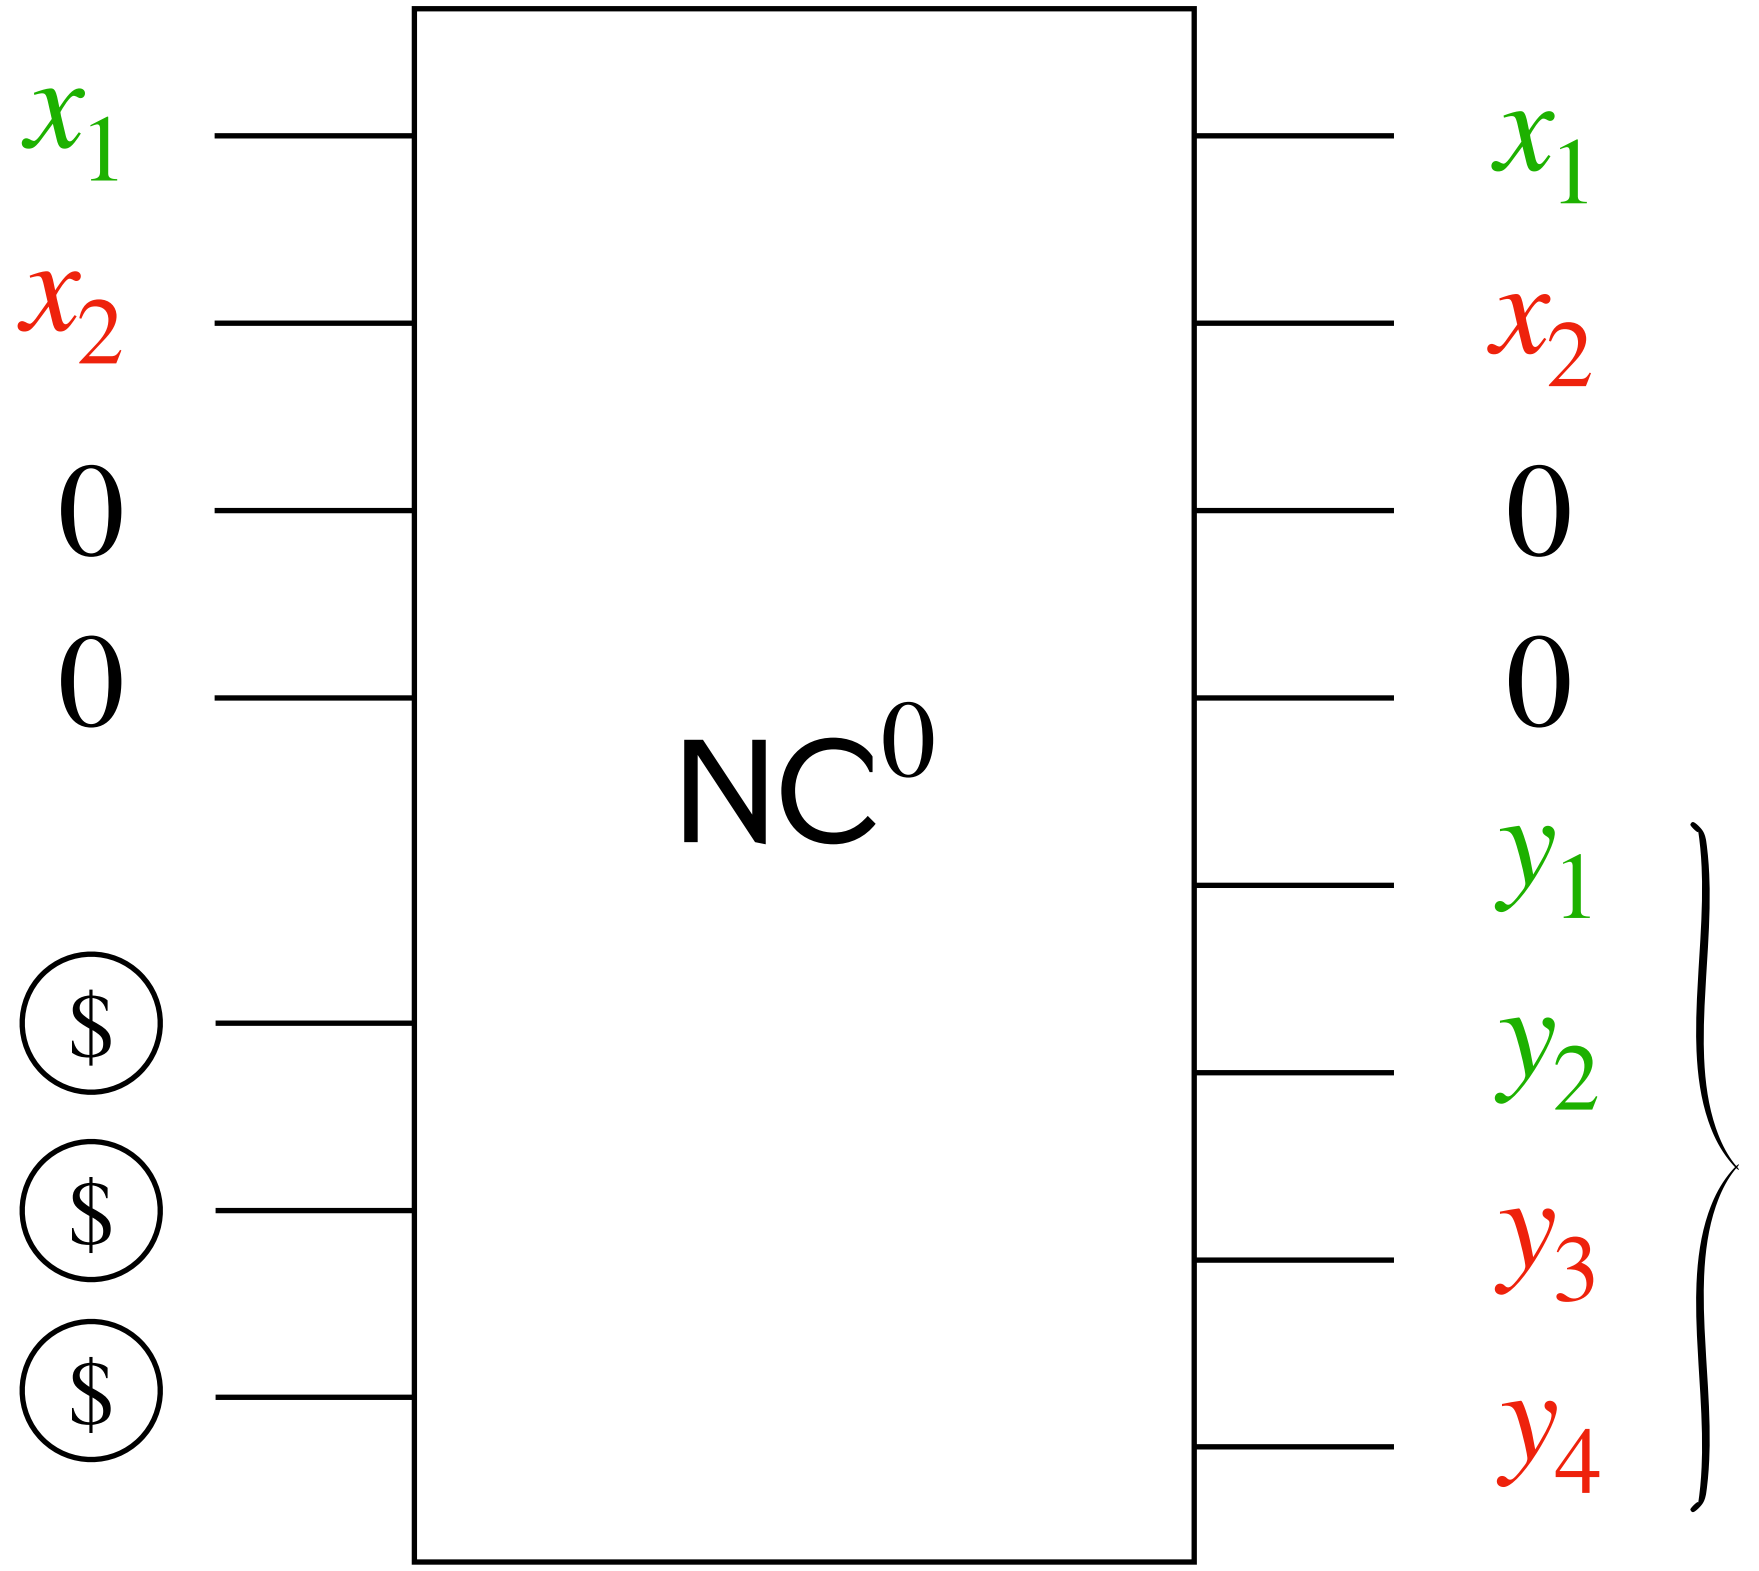

Formal: Consider the potential function

$$\phi(x, y) = i^{|x|+2|y|}$$

$$\text{Parity}(y) = \begin{cases} 0 & \text{if } |x| \equiv 0 \pmod 4 \\ 1 & \text{if } |x| \equiv 2 \pmod 4 \end{cases}$$

y depends on *all* input bits

Potential function under the Parity-Halving distribution

Theorem: $\mathbb{E}[\phi(x,y)] \approx 1/2$ for the Parity-Halving problem

$i^{|x|+2|y|}$

Parity$(x) = 0$:

$$|x| \equiv 0 \pmod 4 \longrightarrow |y| \equiv 0 \pmod 2 \longrightarrow \phi(x,y) = 1$$

$$|x| \equiv 2 \pmod 4 \longrightarrow |y| \equiv 1 \pmod 2 \longrightarrow \phi(x,y) = 1$$

Parity$(x) = 1$:

$$\mathbb{E}_y[i^{2y}] = \mathbb{E}_y[(-1)^{|y|}] = 0$$

Expectation follows since $\Pr[\text{Parity}(x) = 0] \approx \Pr[\text{Parity}(x) = 1] \approx 1/2$

Meanwhile…

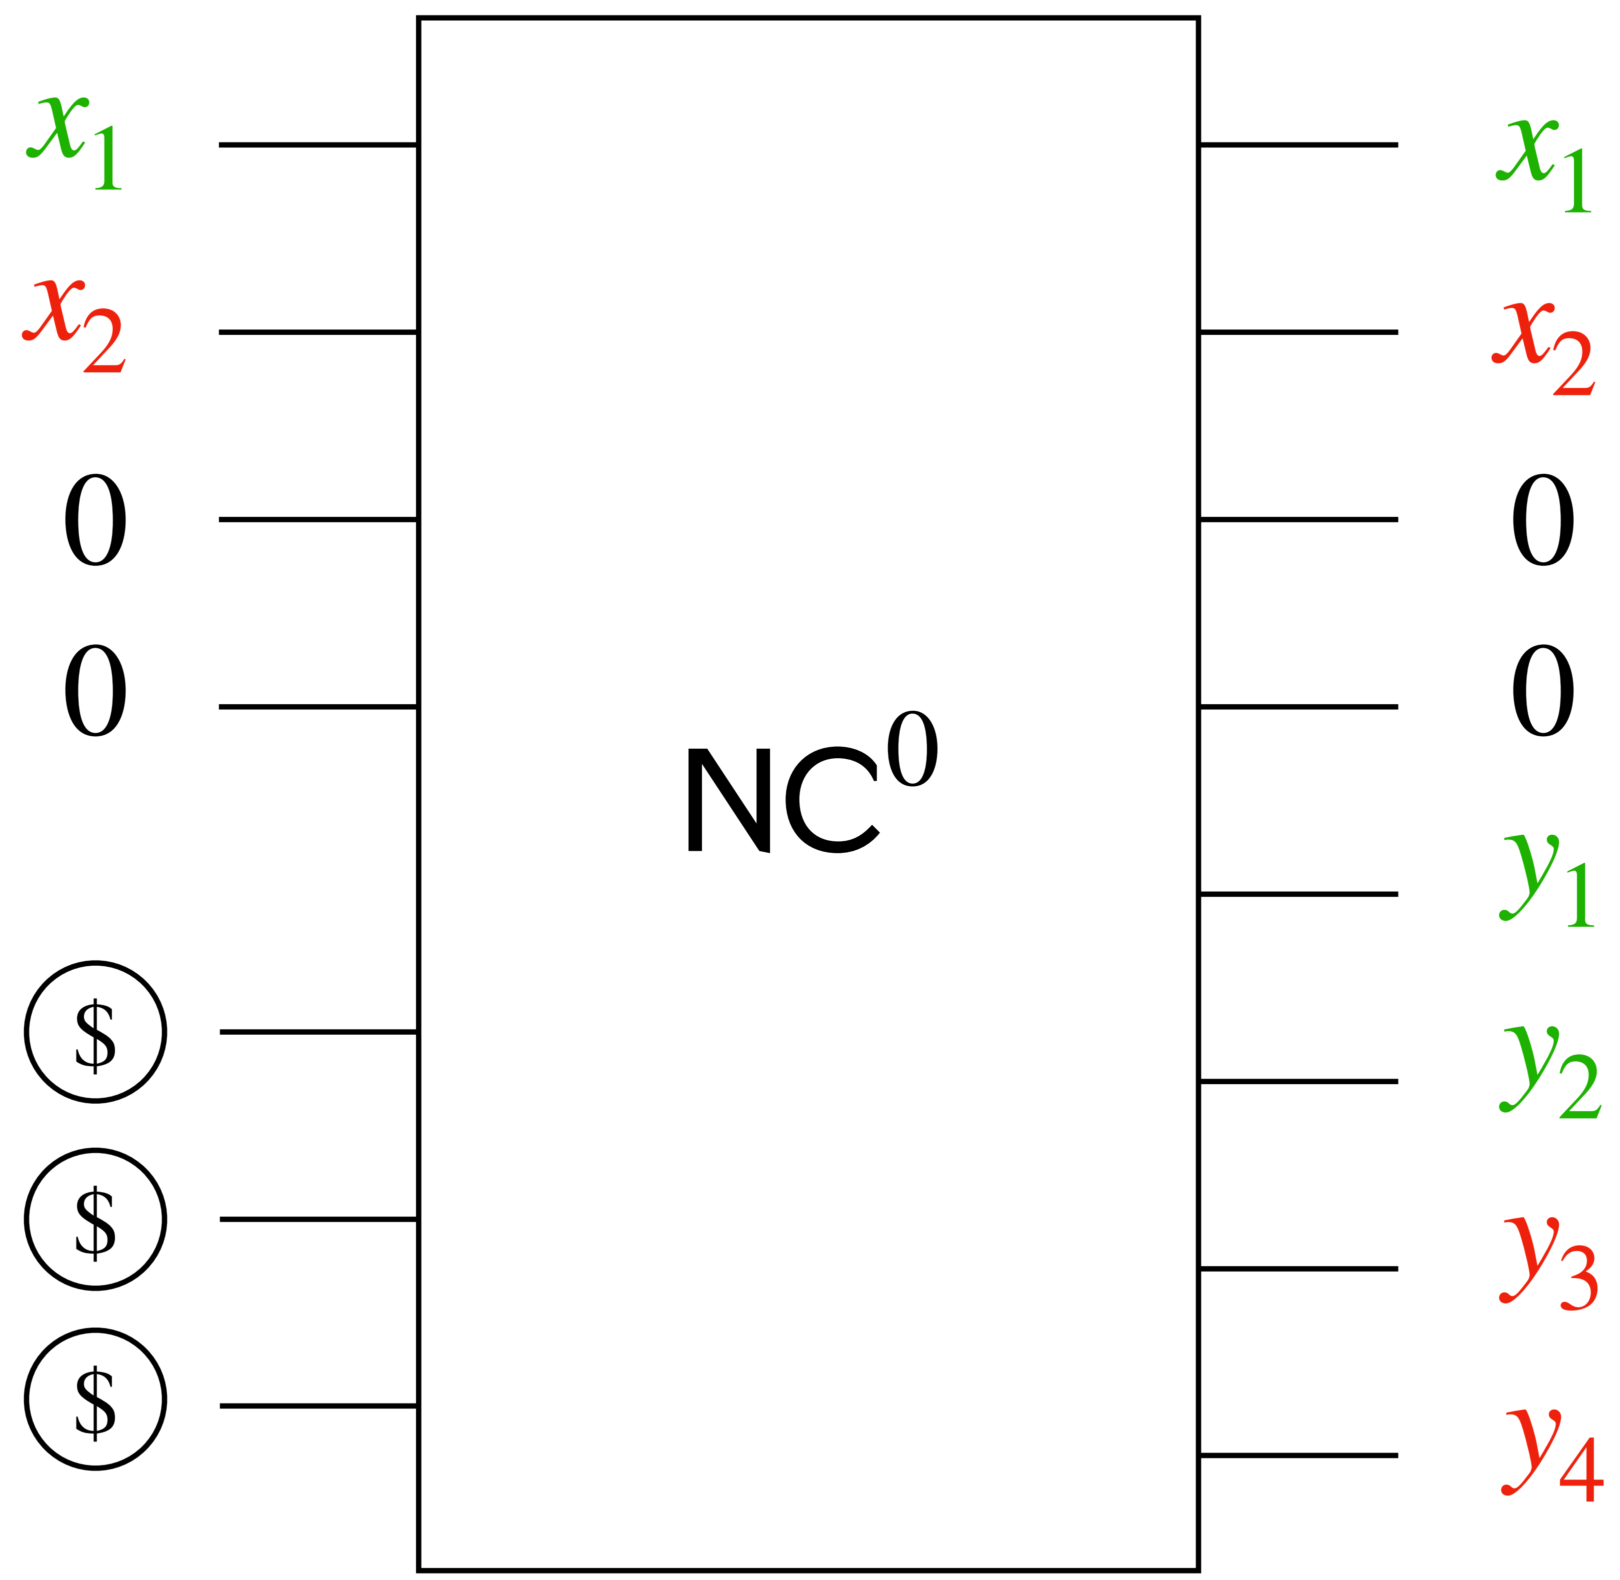

$$\phi(x, y) = i^{|x|+2|y|}$$

$$= i^{x_1+x_2+2(y_1+y_2+y_3+y_4)}$$

$$= i^{x_1+2(y_1+y_2)} \cdot i^{x_2+2(y_3+y_4)}$$

$$= i^{x_1+2(y_1+y_2)} \cdot i^{x_2+2(y_3+y_4)}$$

$$\mathbb{E}[\phi(x, y) \mid \textcircled{\$}\,] = \mathbb{E}[i^{x_1+2(y_1+y_2)}] \cdot \mathbb{E}[i^{x_2+2(y_3+y_4)}]$$

These terms are each $\ll 1$

Lightcones constrain correlations in classical circuits

Backwards lightcone: The set of inputs that affect an output

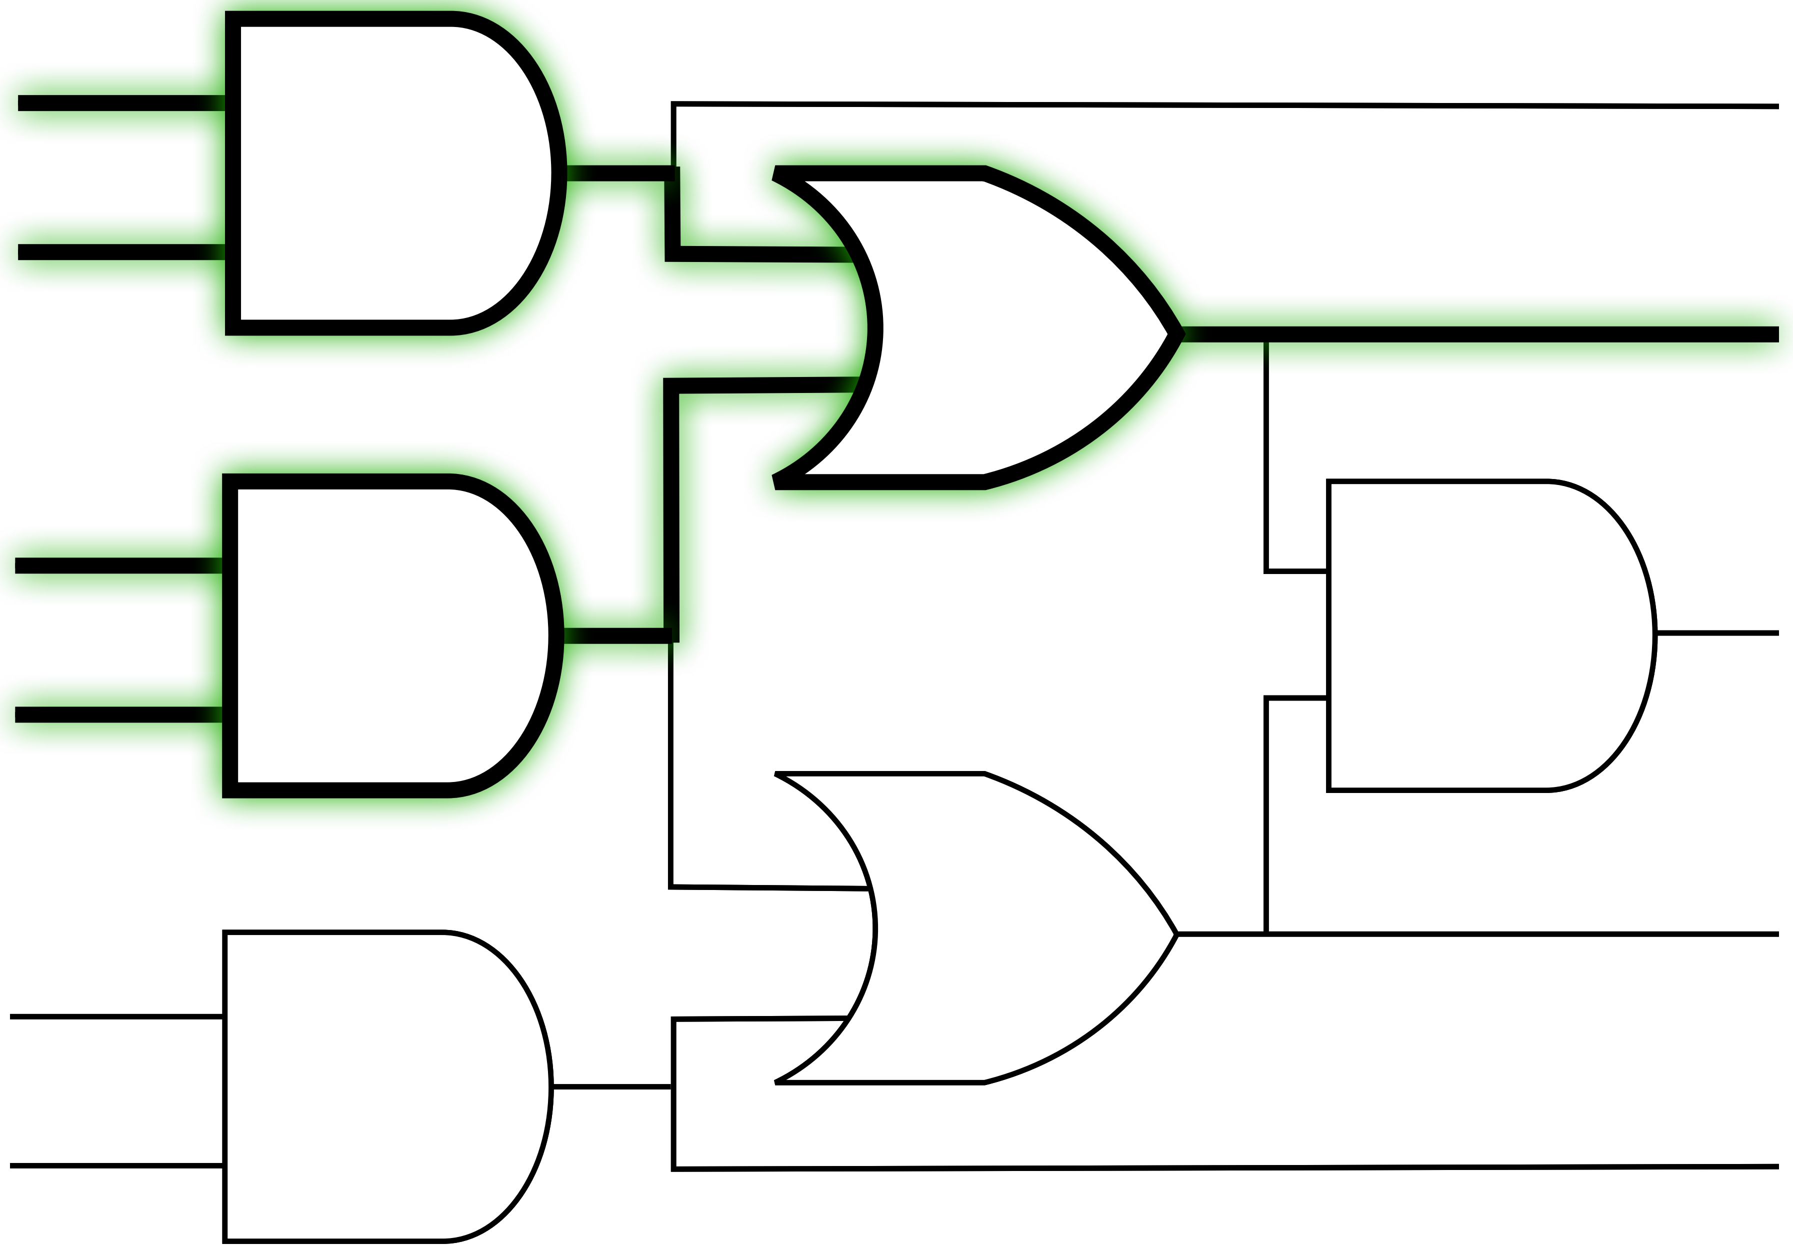

Key fact: Backwards lightcones in NC^0 circuit are of size $2^{O(\text{depth})}$

Why can't all my lightcones be huge?

Forward lightcone: The set of outputs affected by an input

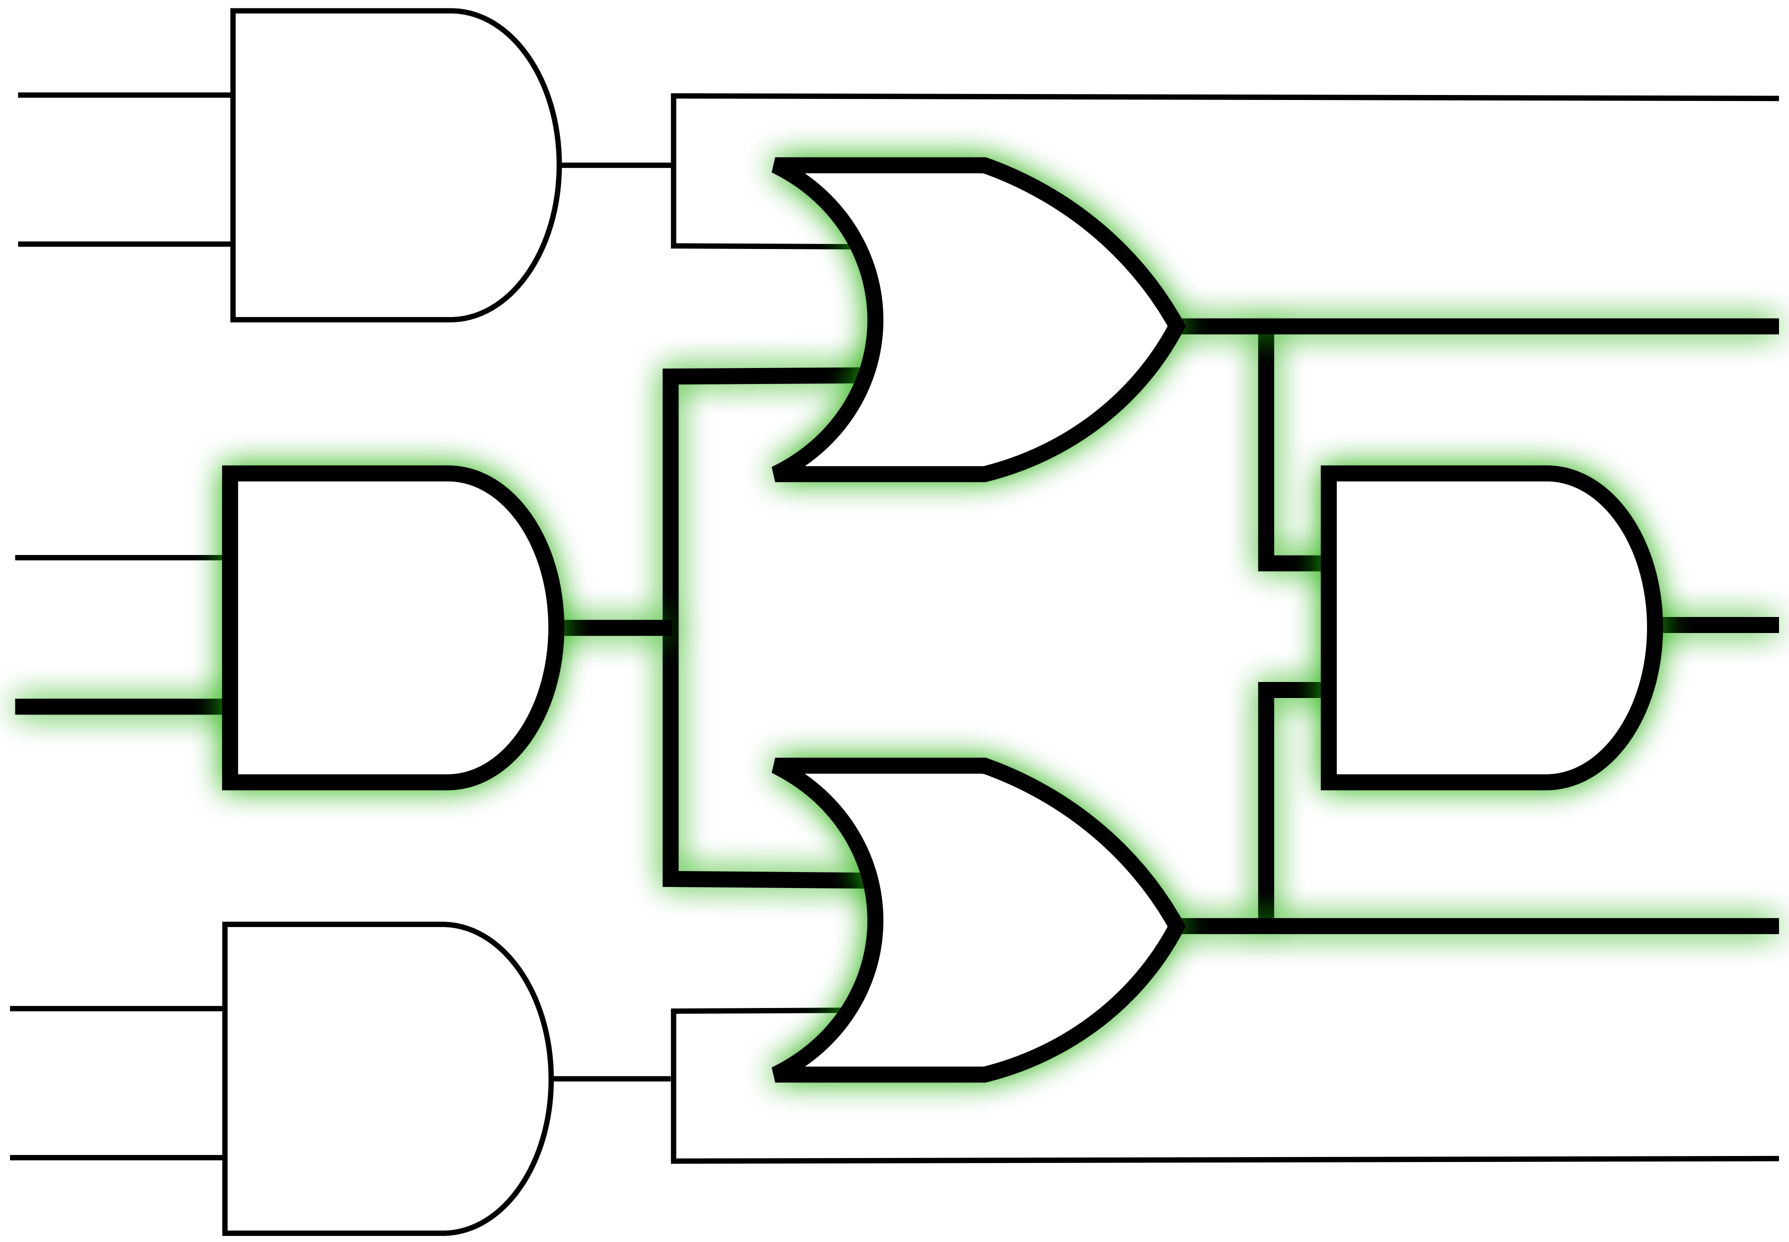

Why can’t all my lightcones be huge?

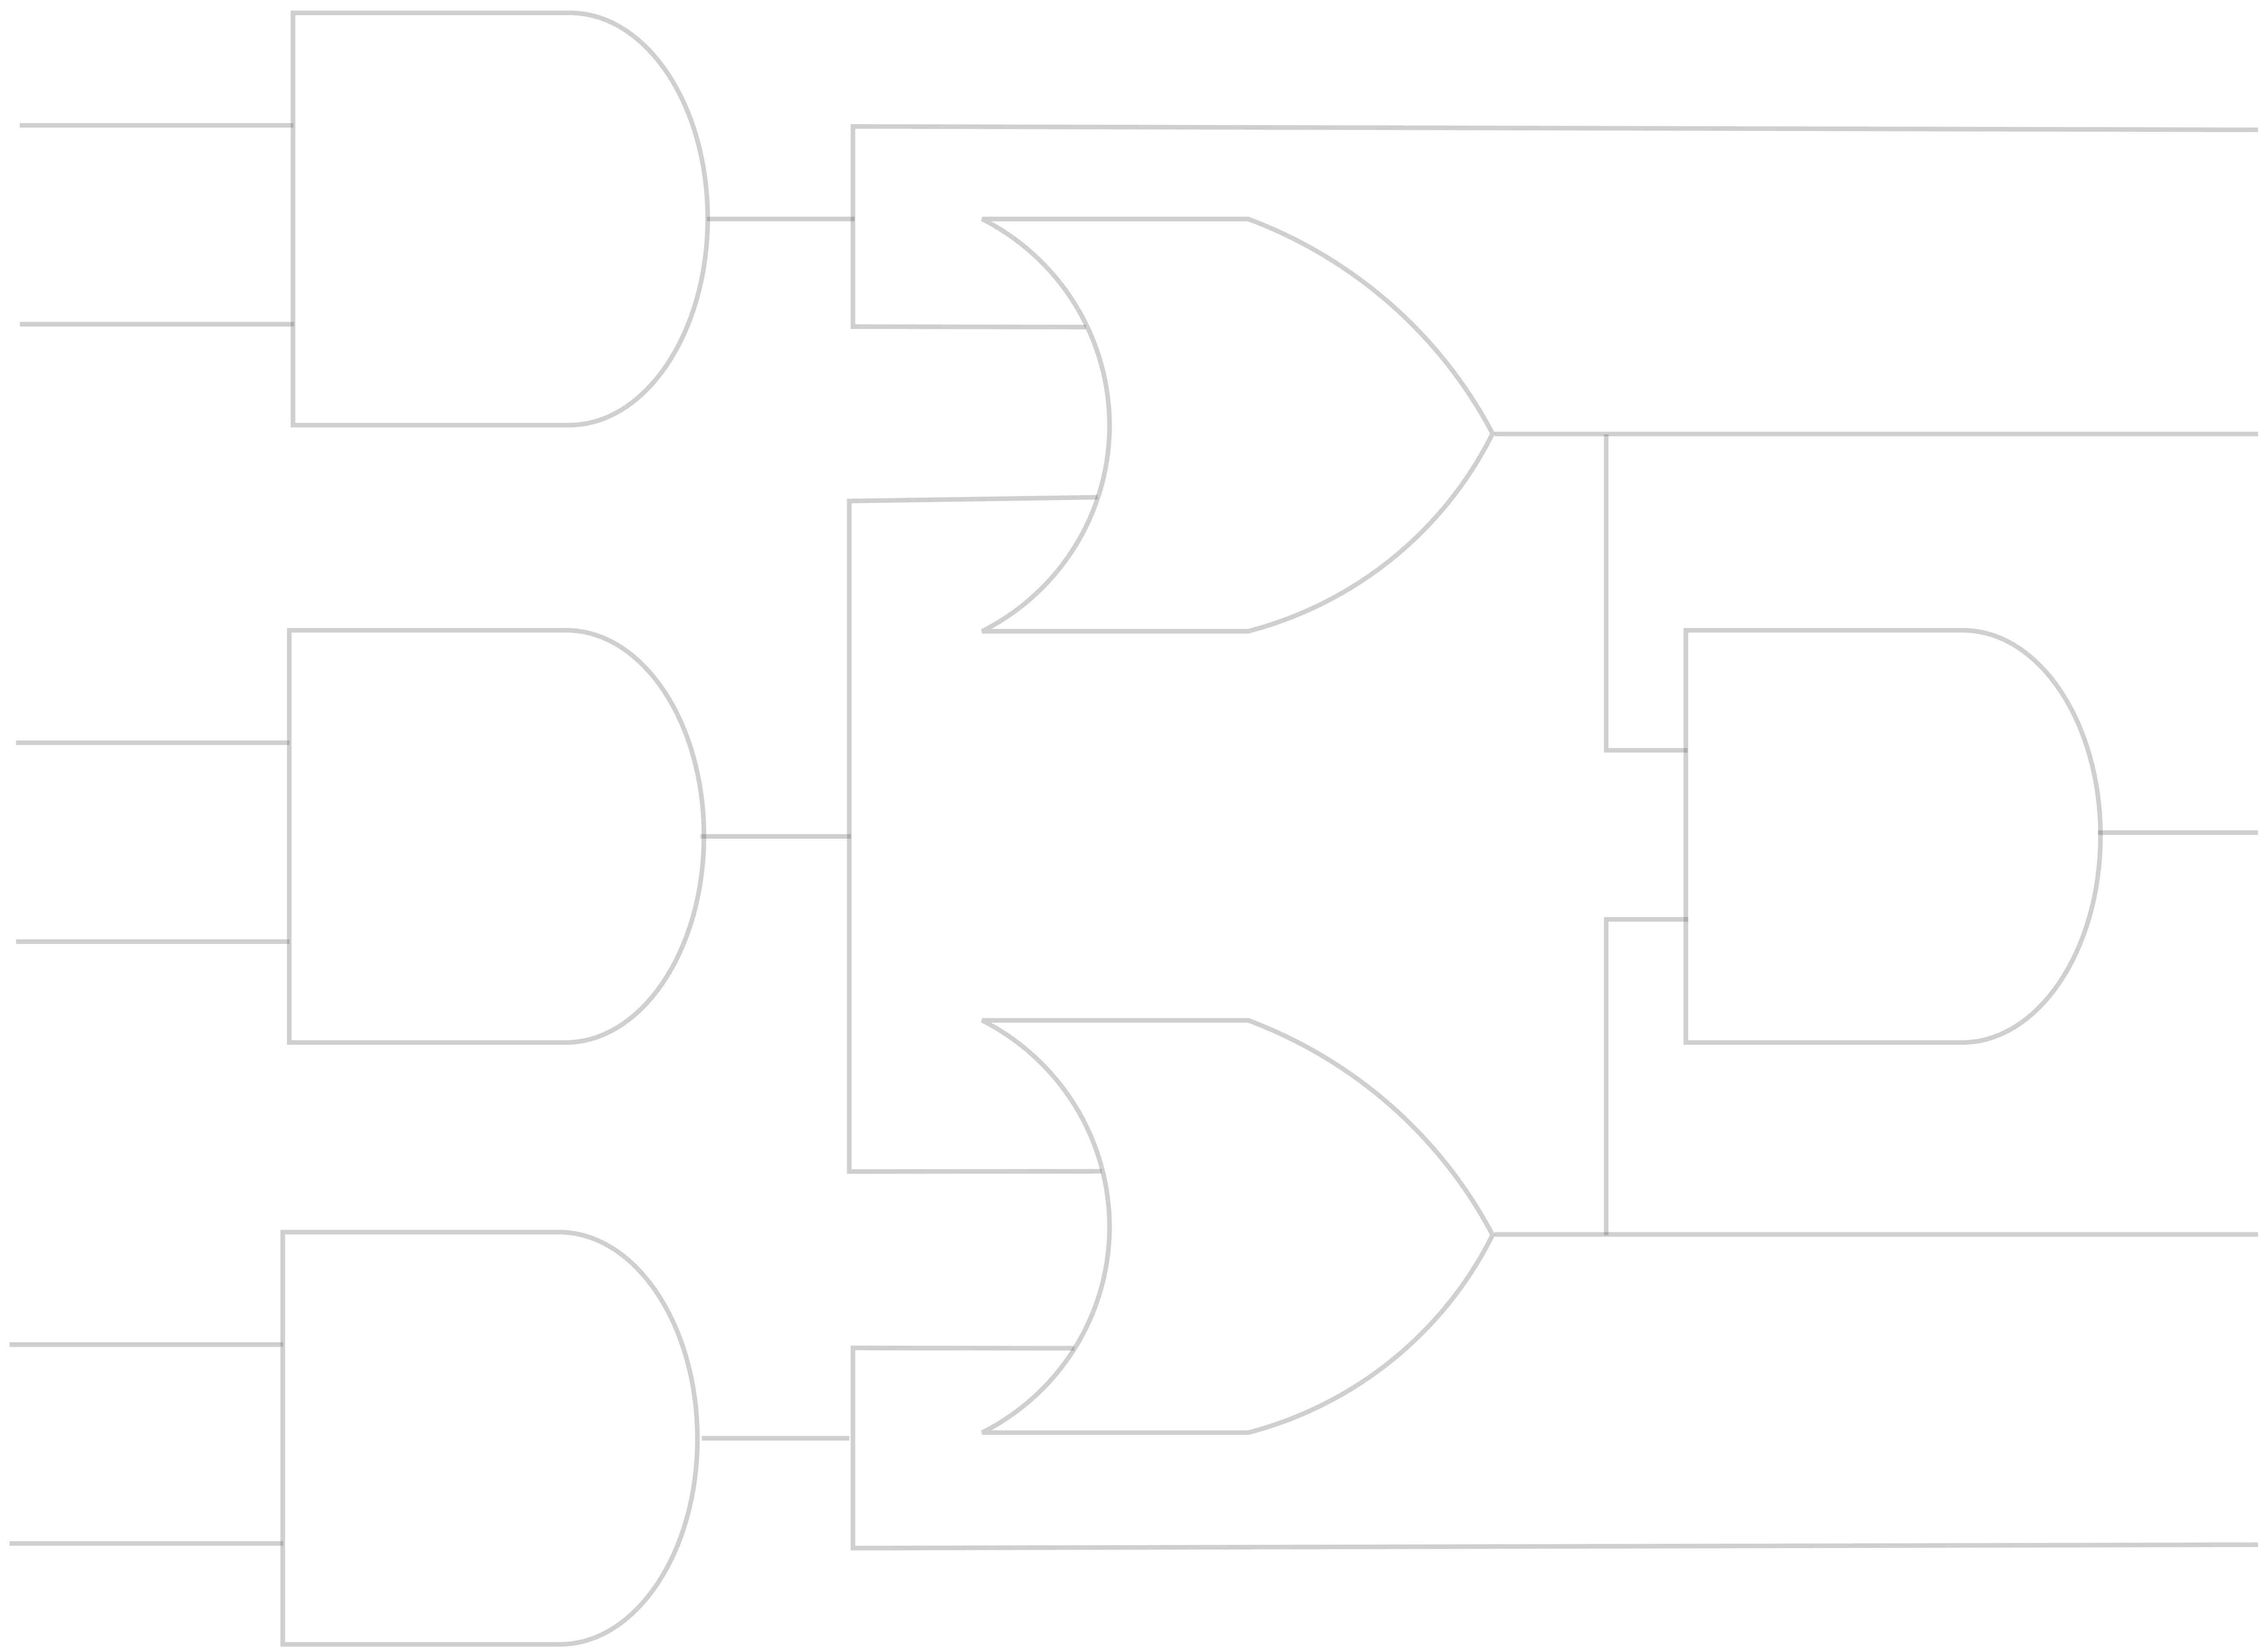

Why can’t all my lightcones be huge?

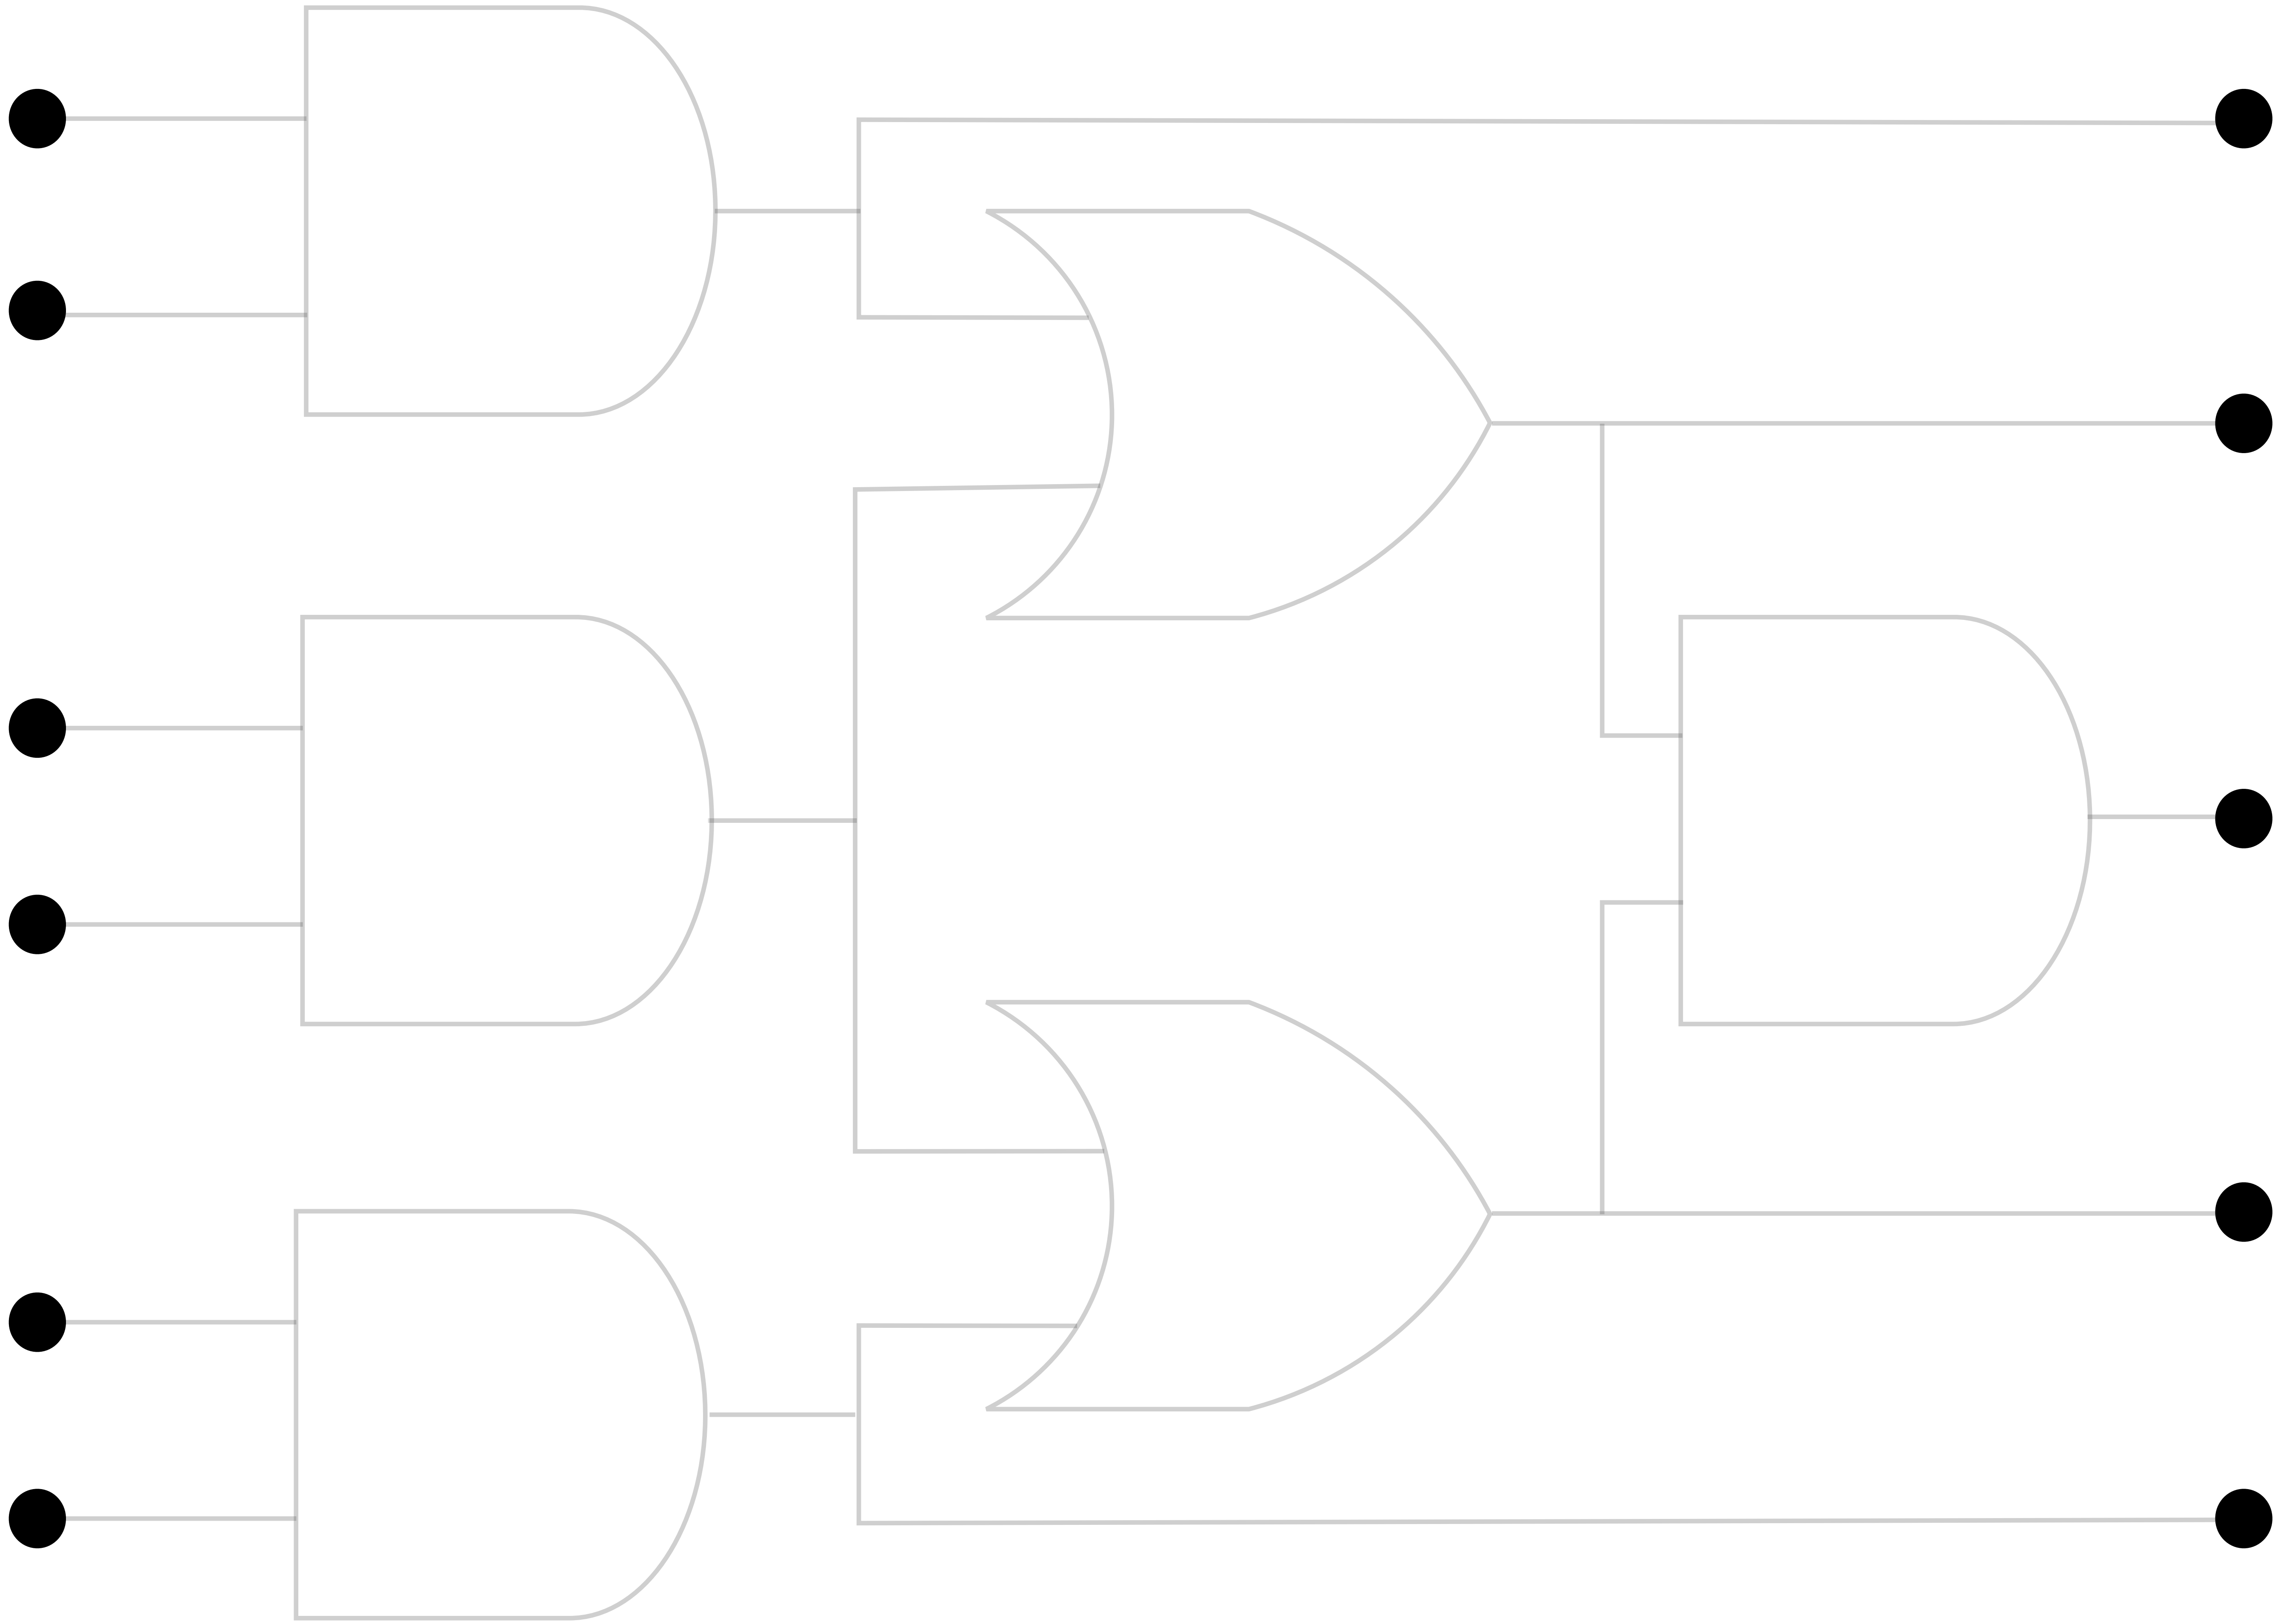

Why can't all my lightcones be huge?

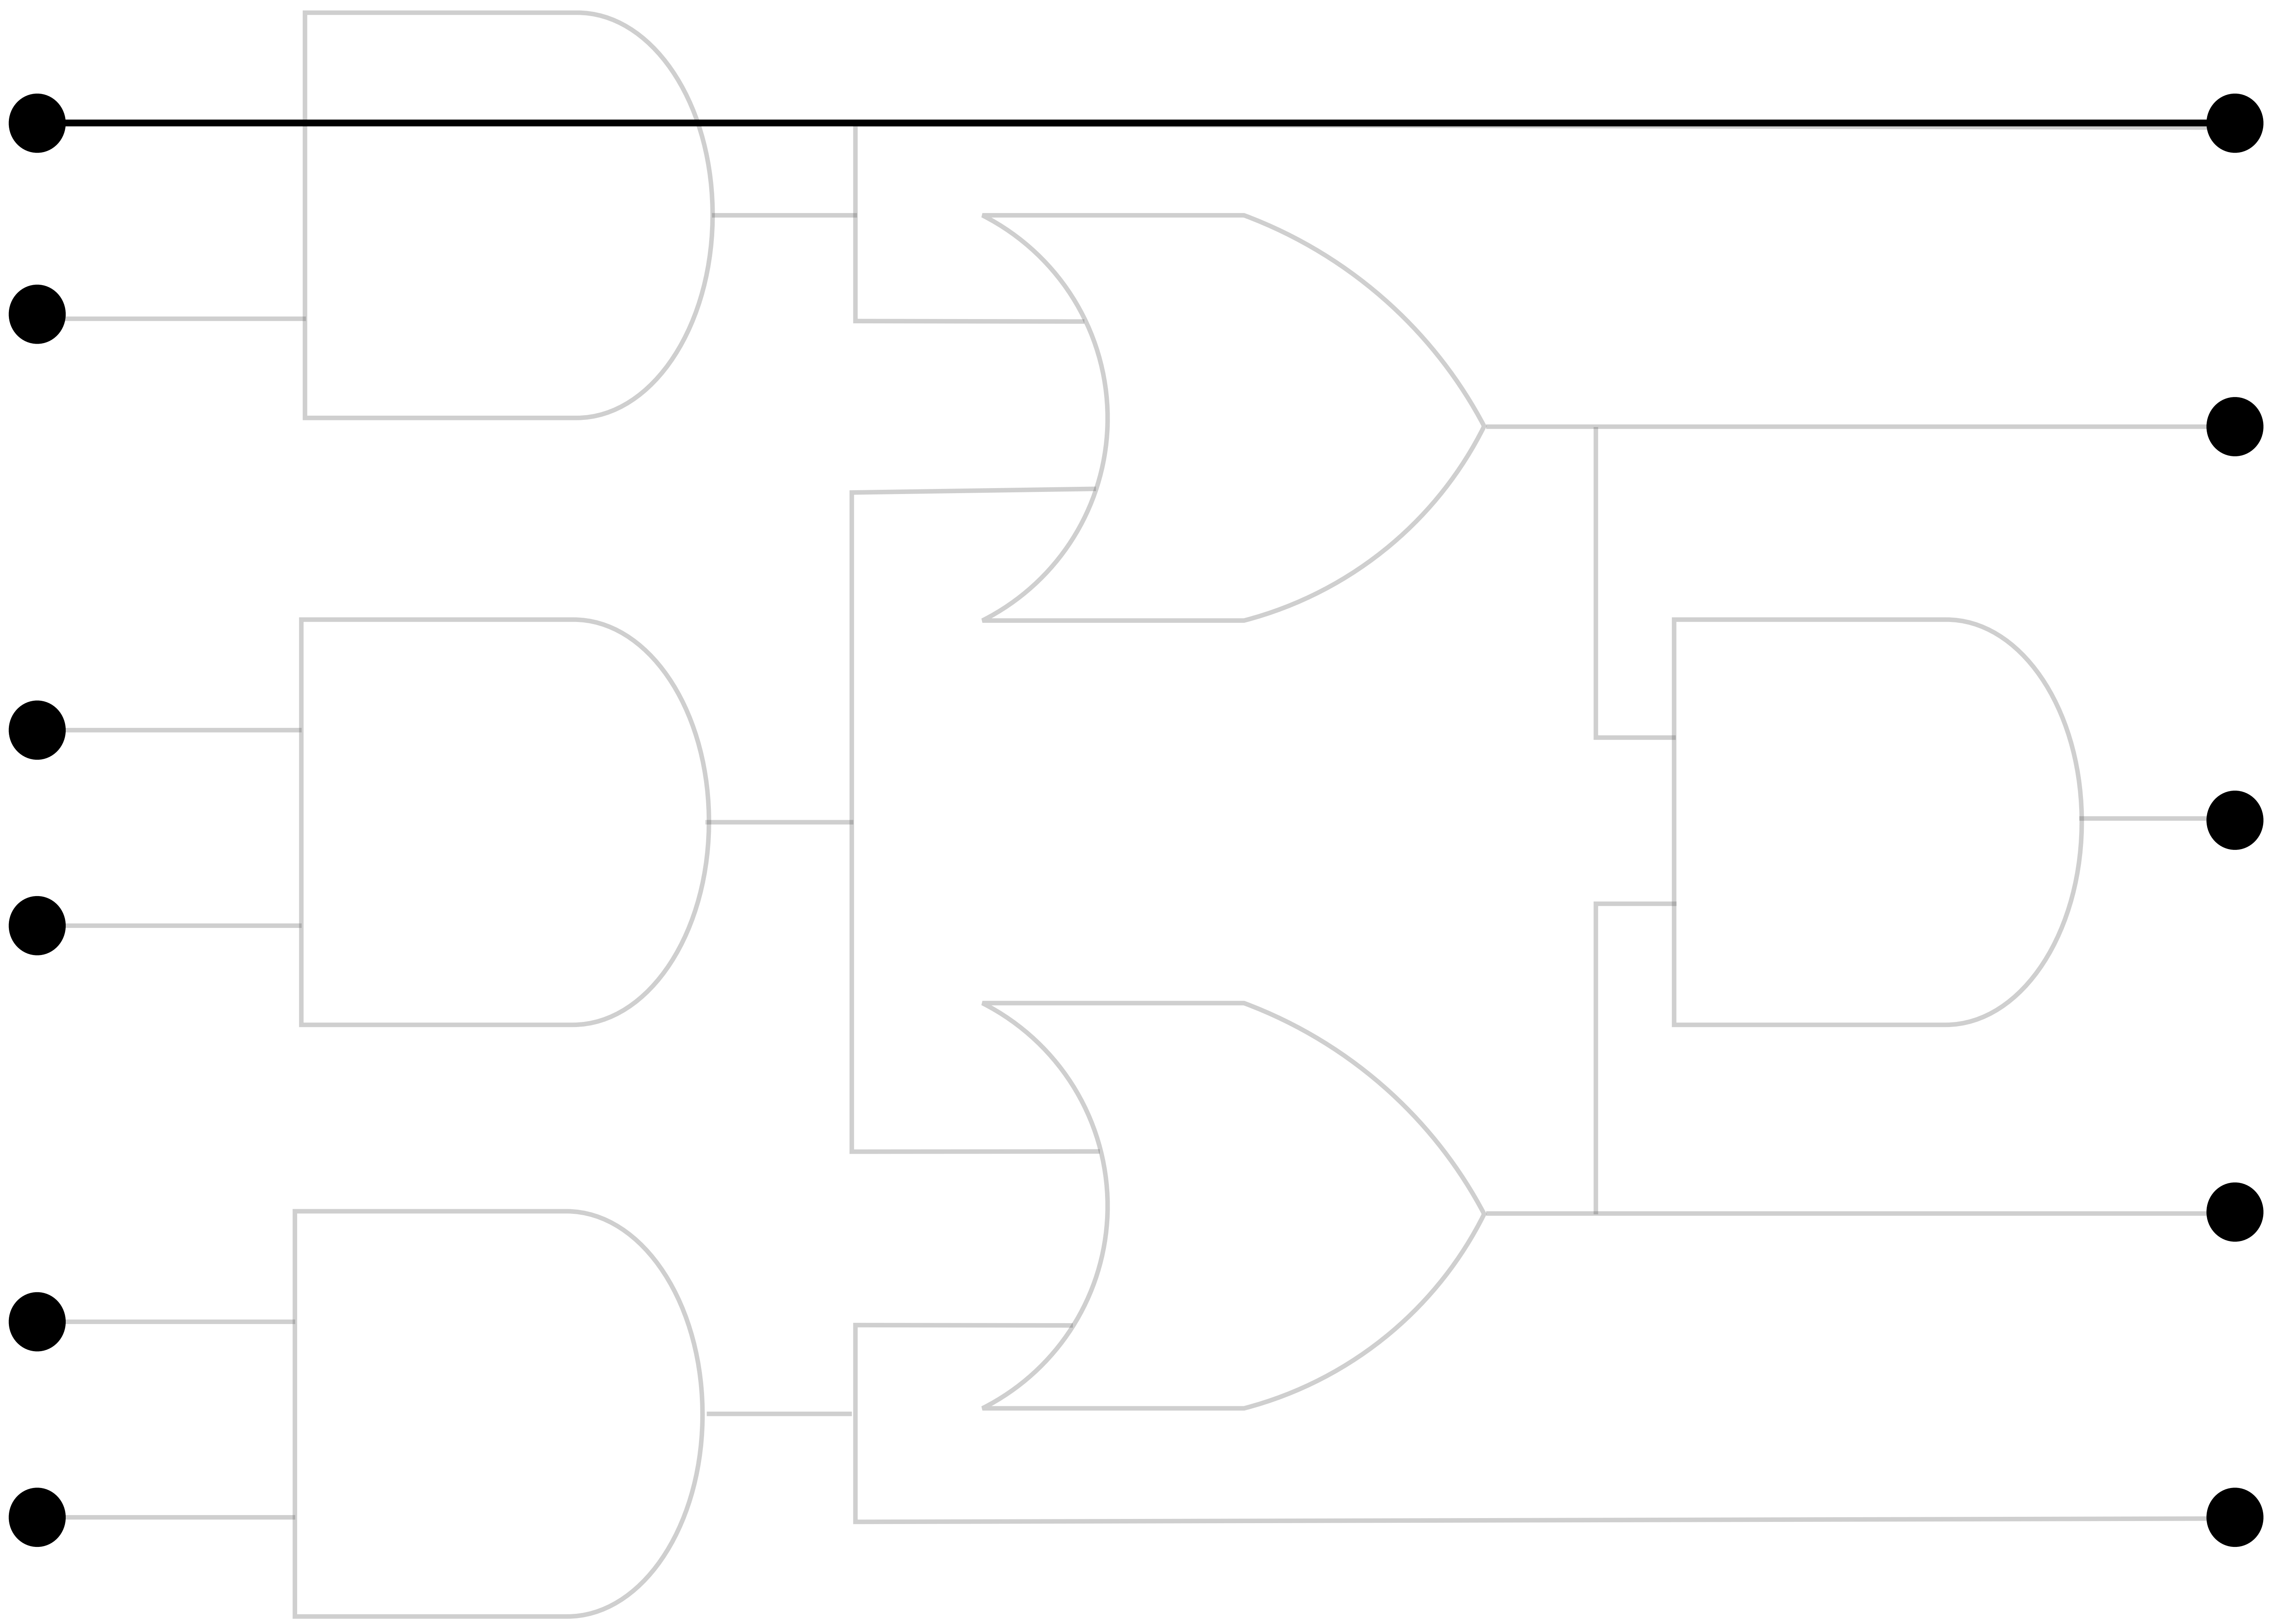

Why can’t all my lightcones be huge?

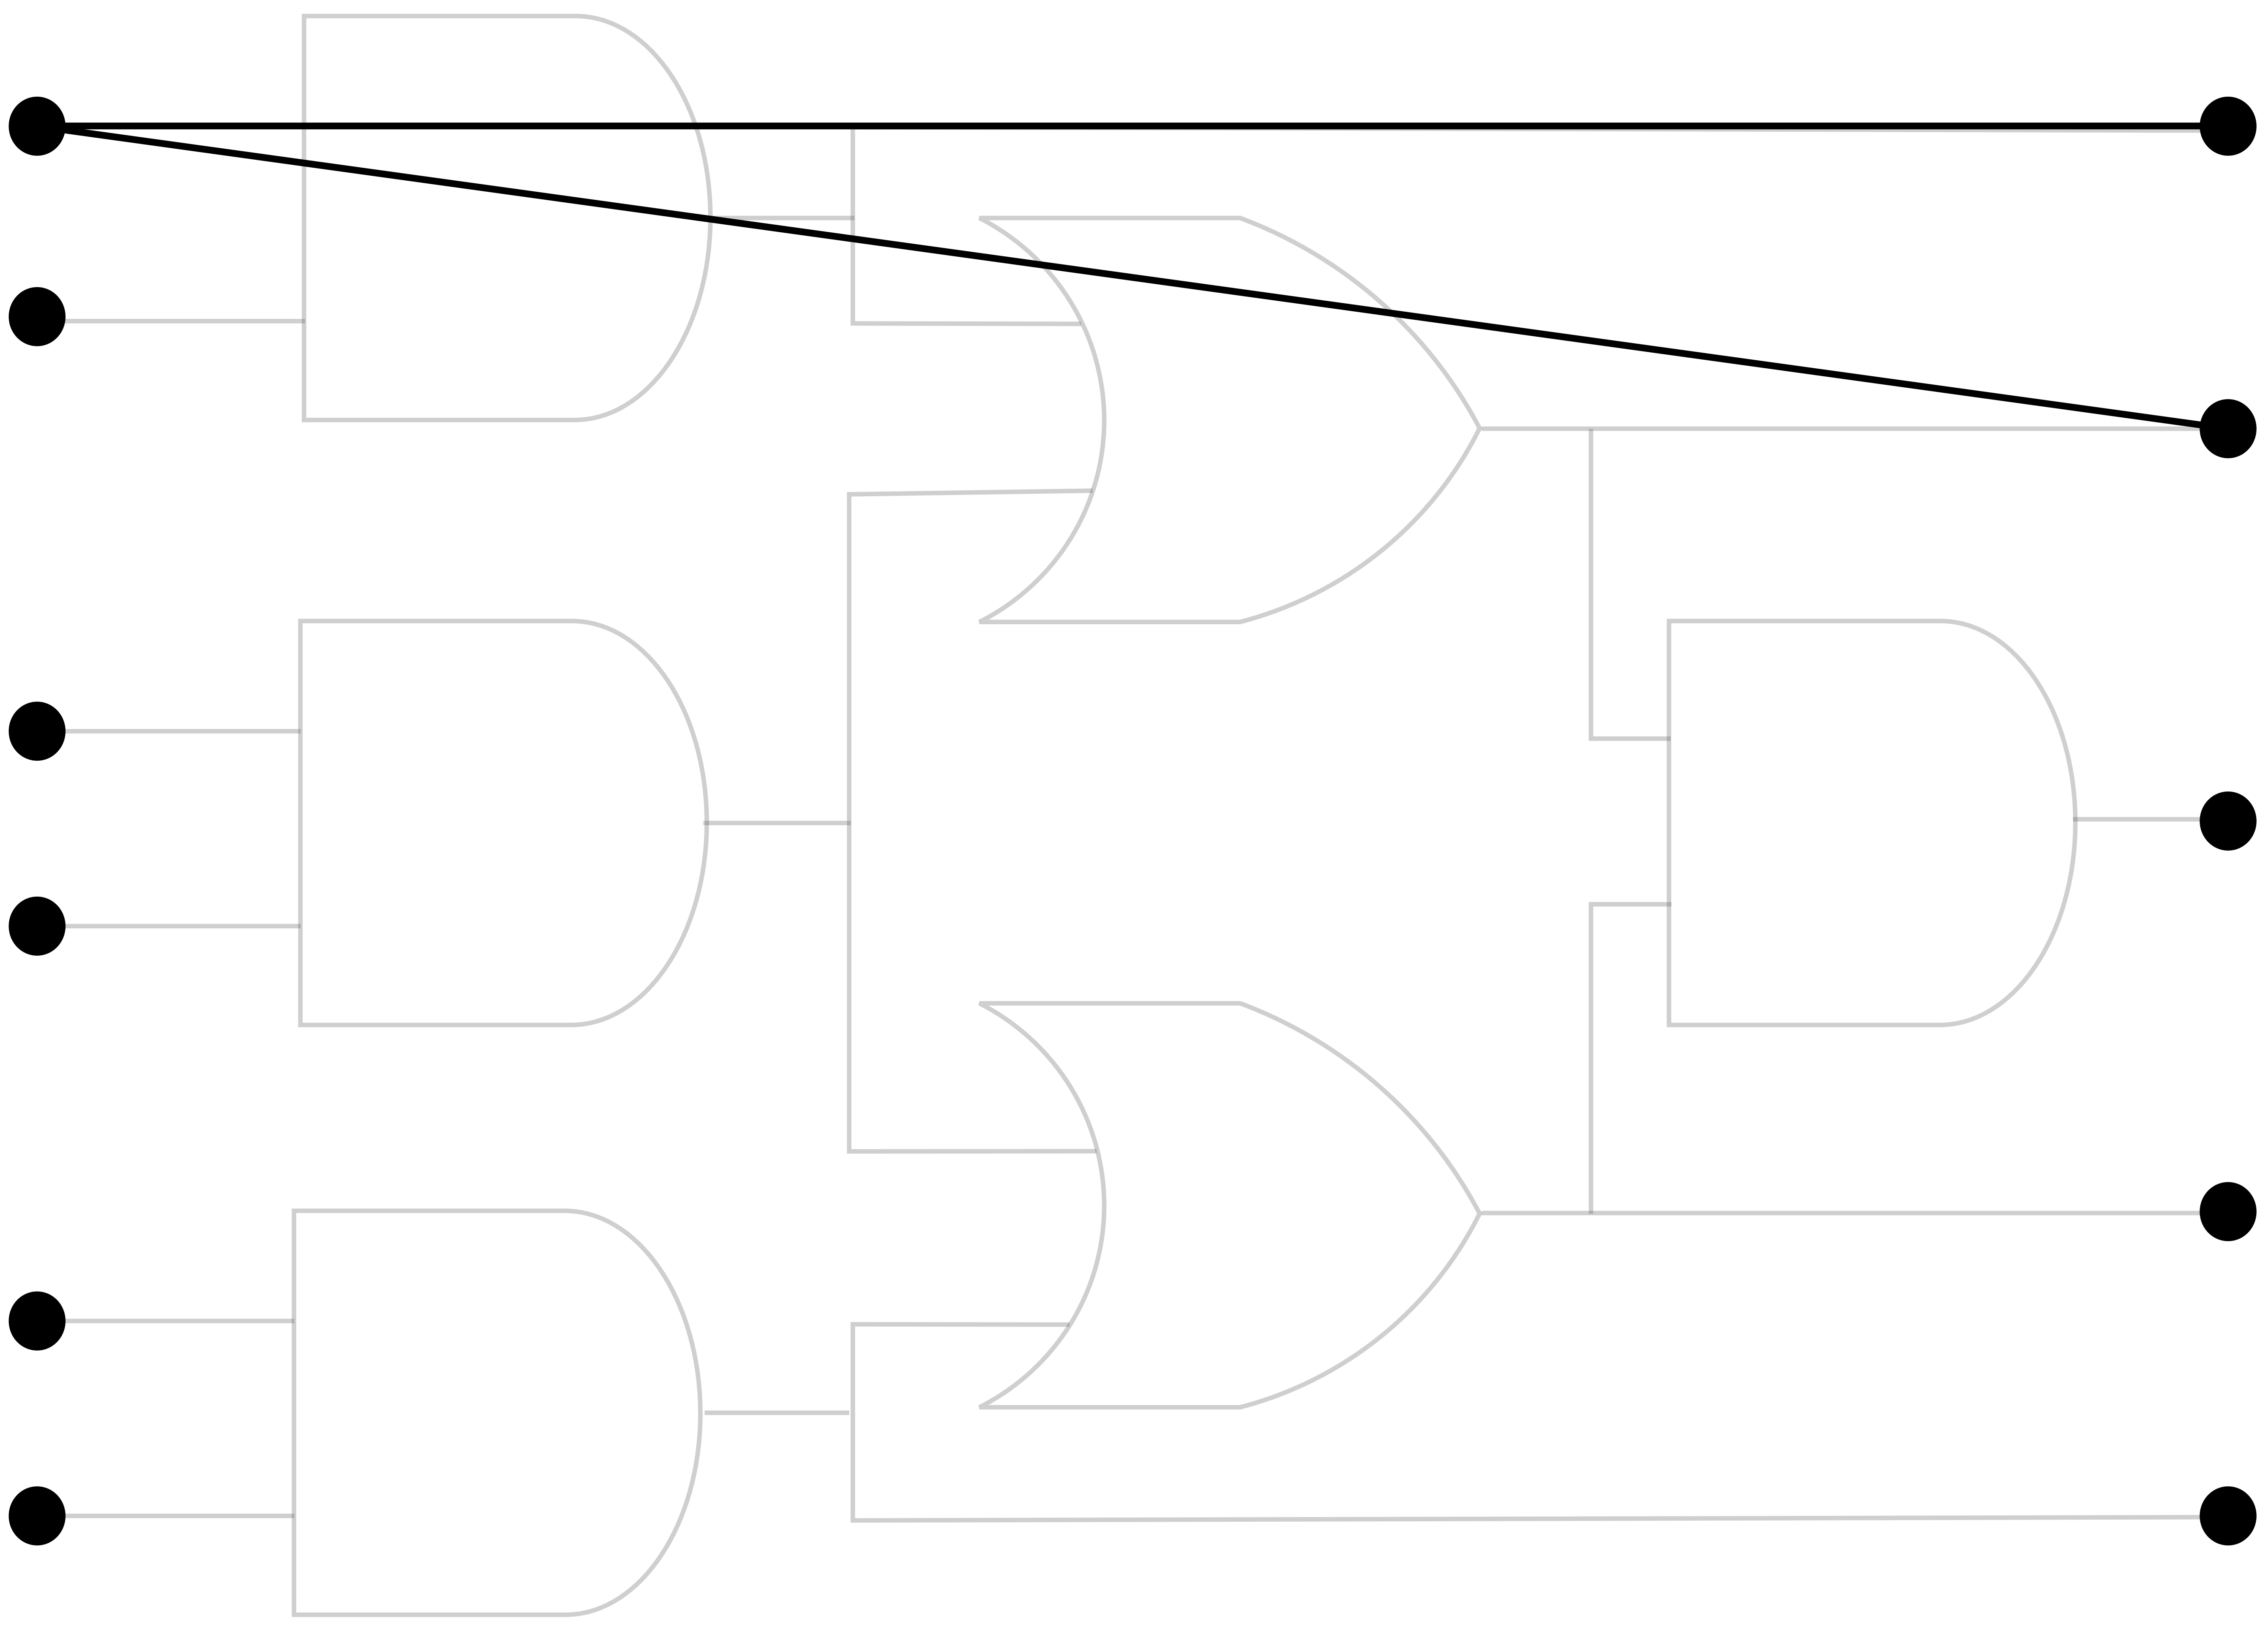

Why can't all my lightcones be huge?

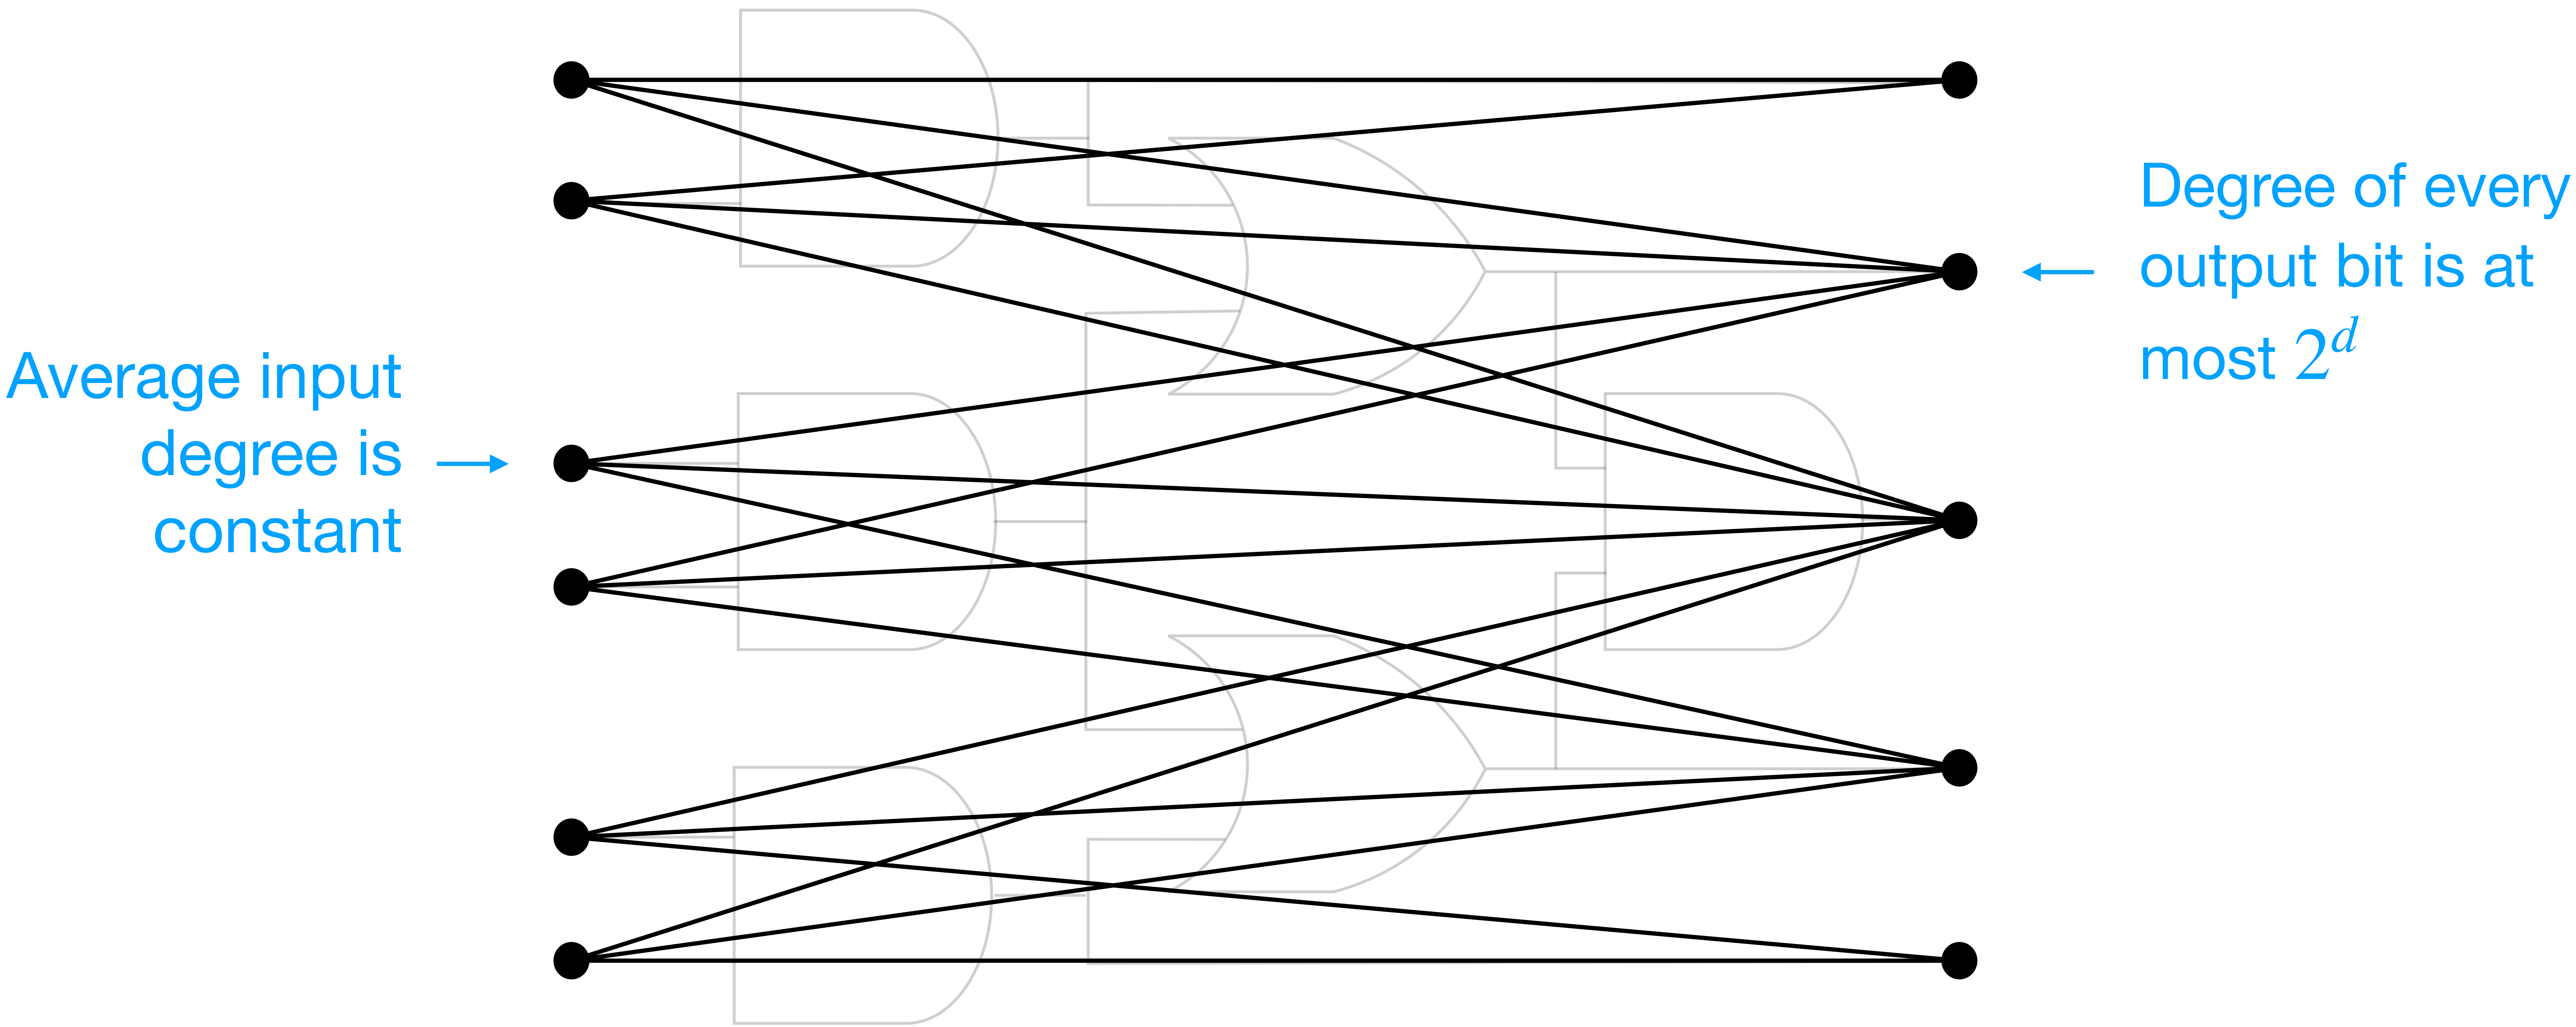

Upshot: Conditioning on high-degree inputs, gives low degree everywhere

Why can’t all my lightcones be huge?

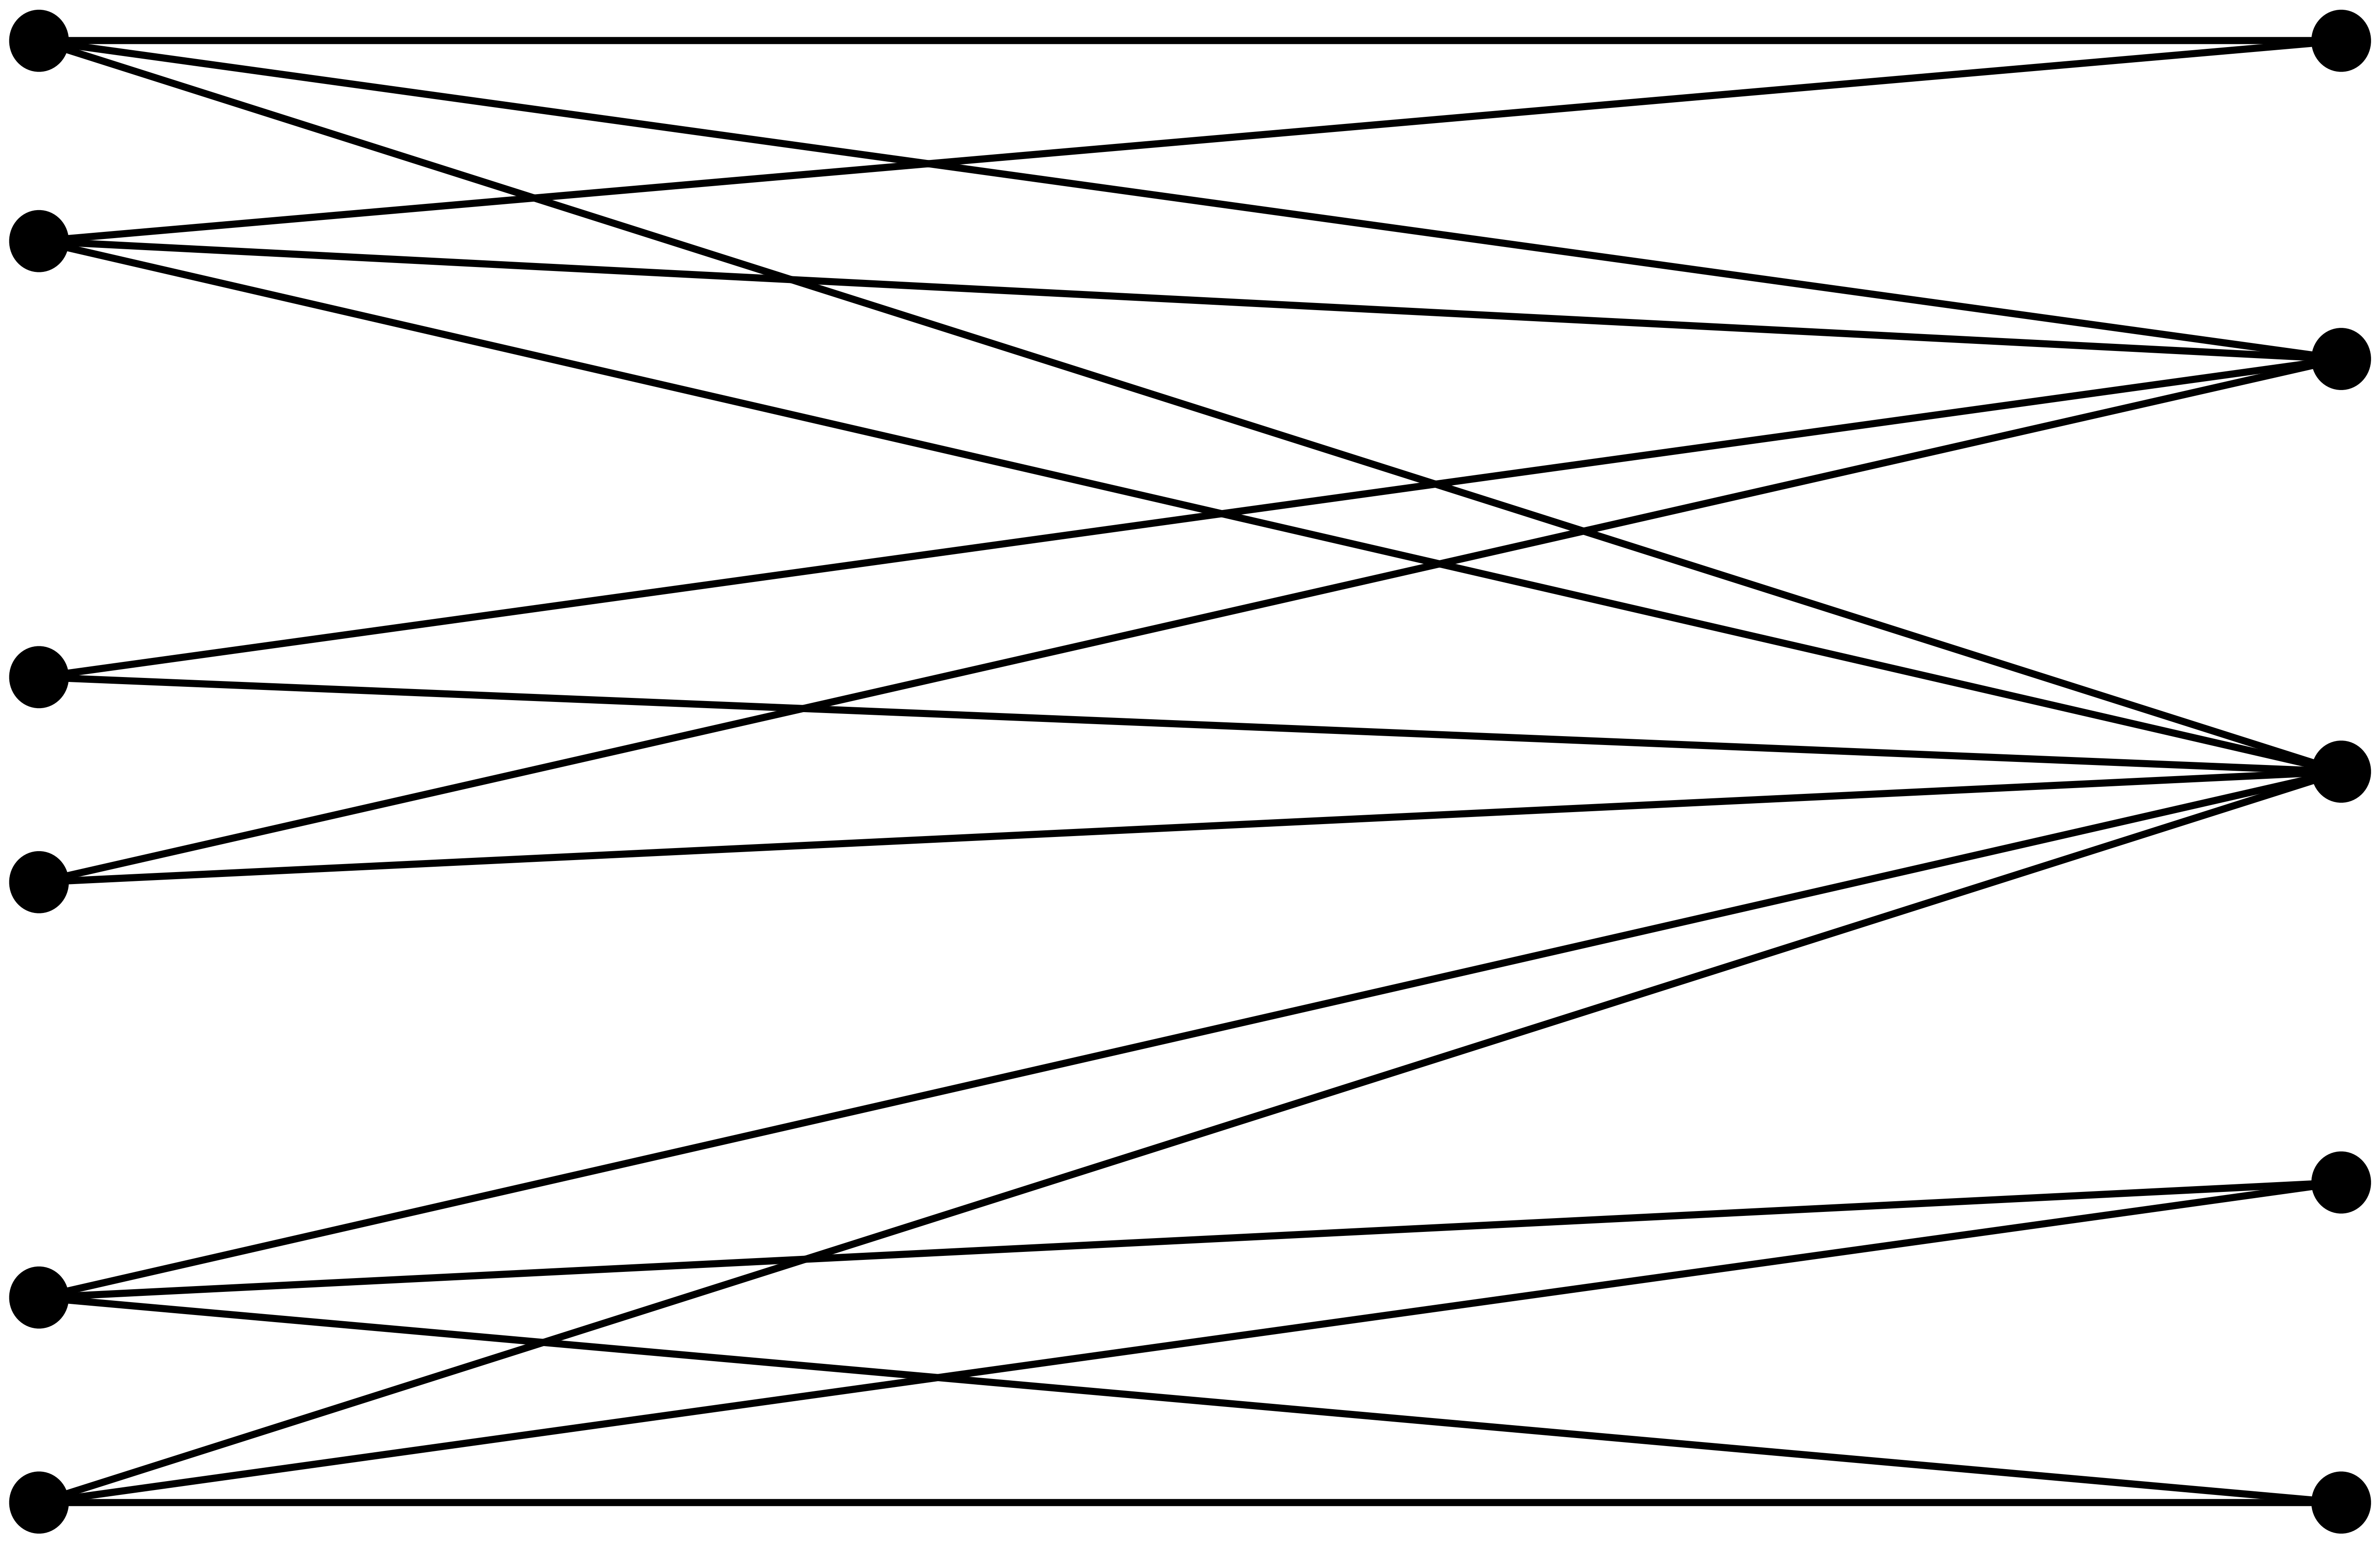

Why can’t all my lightcones be huge?

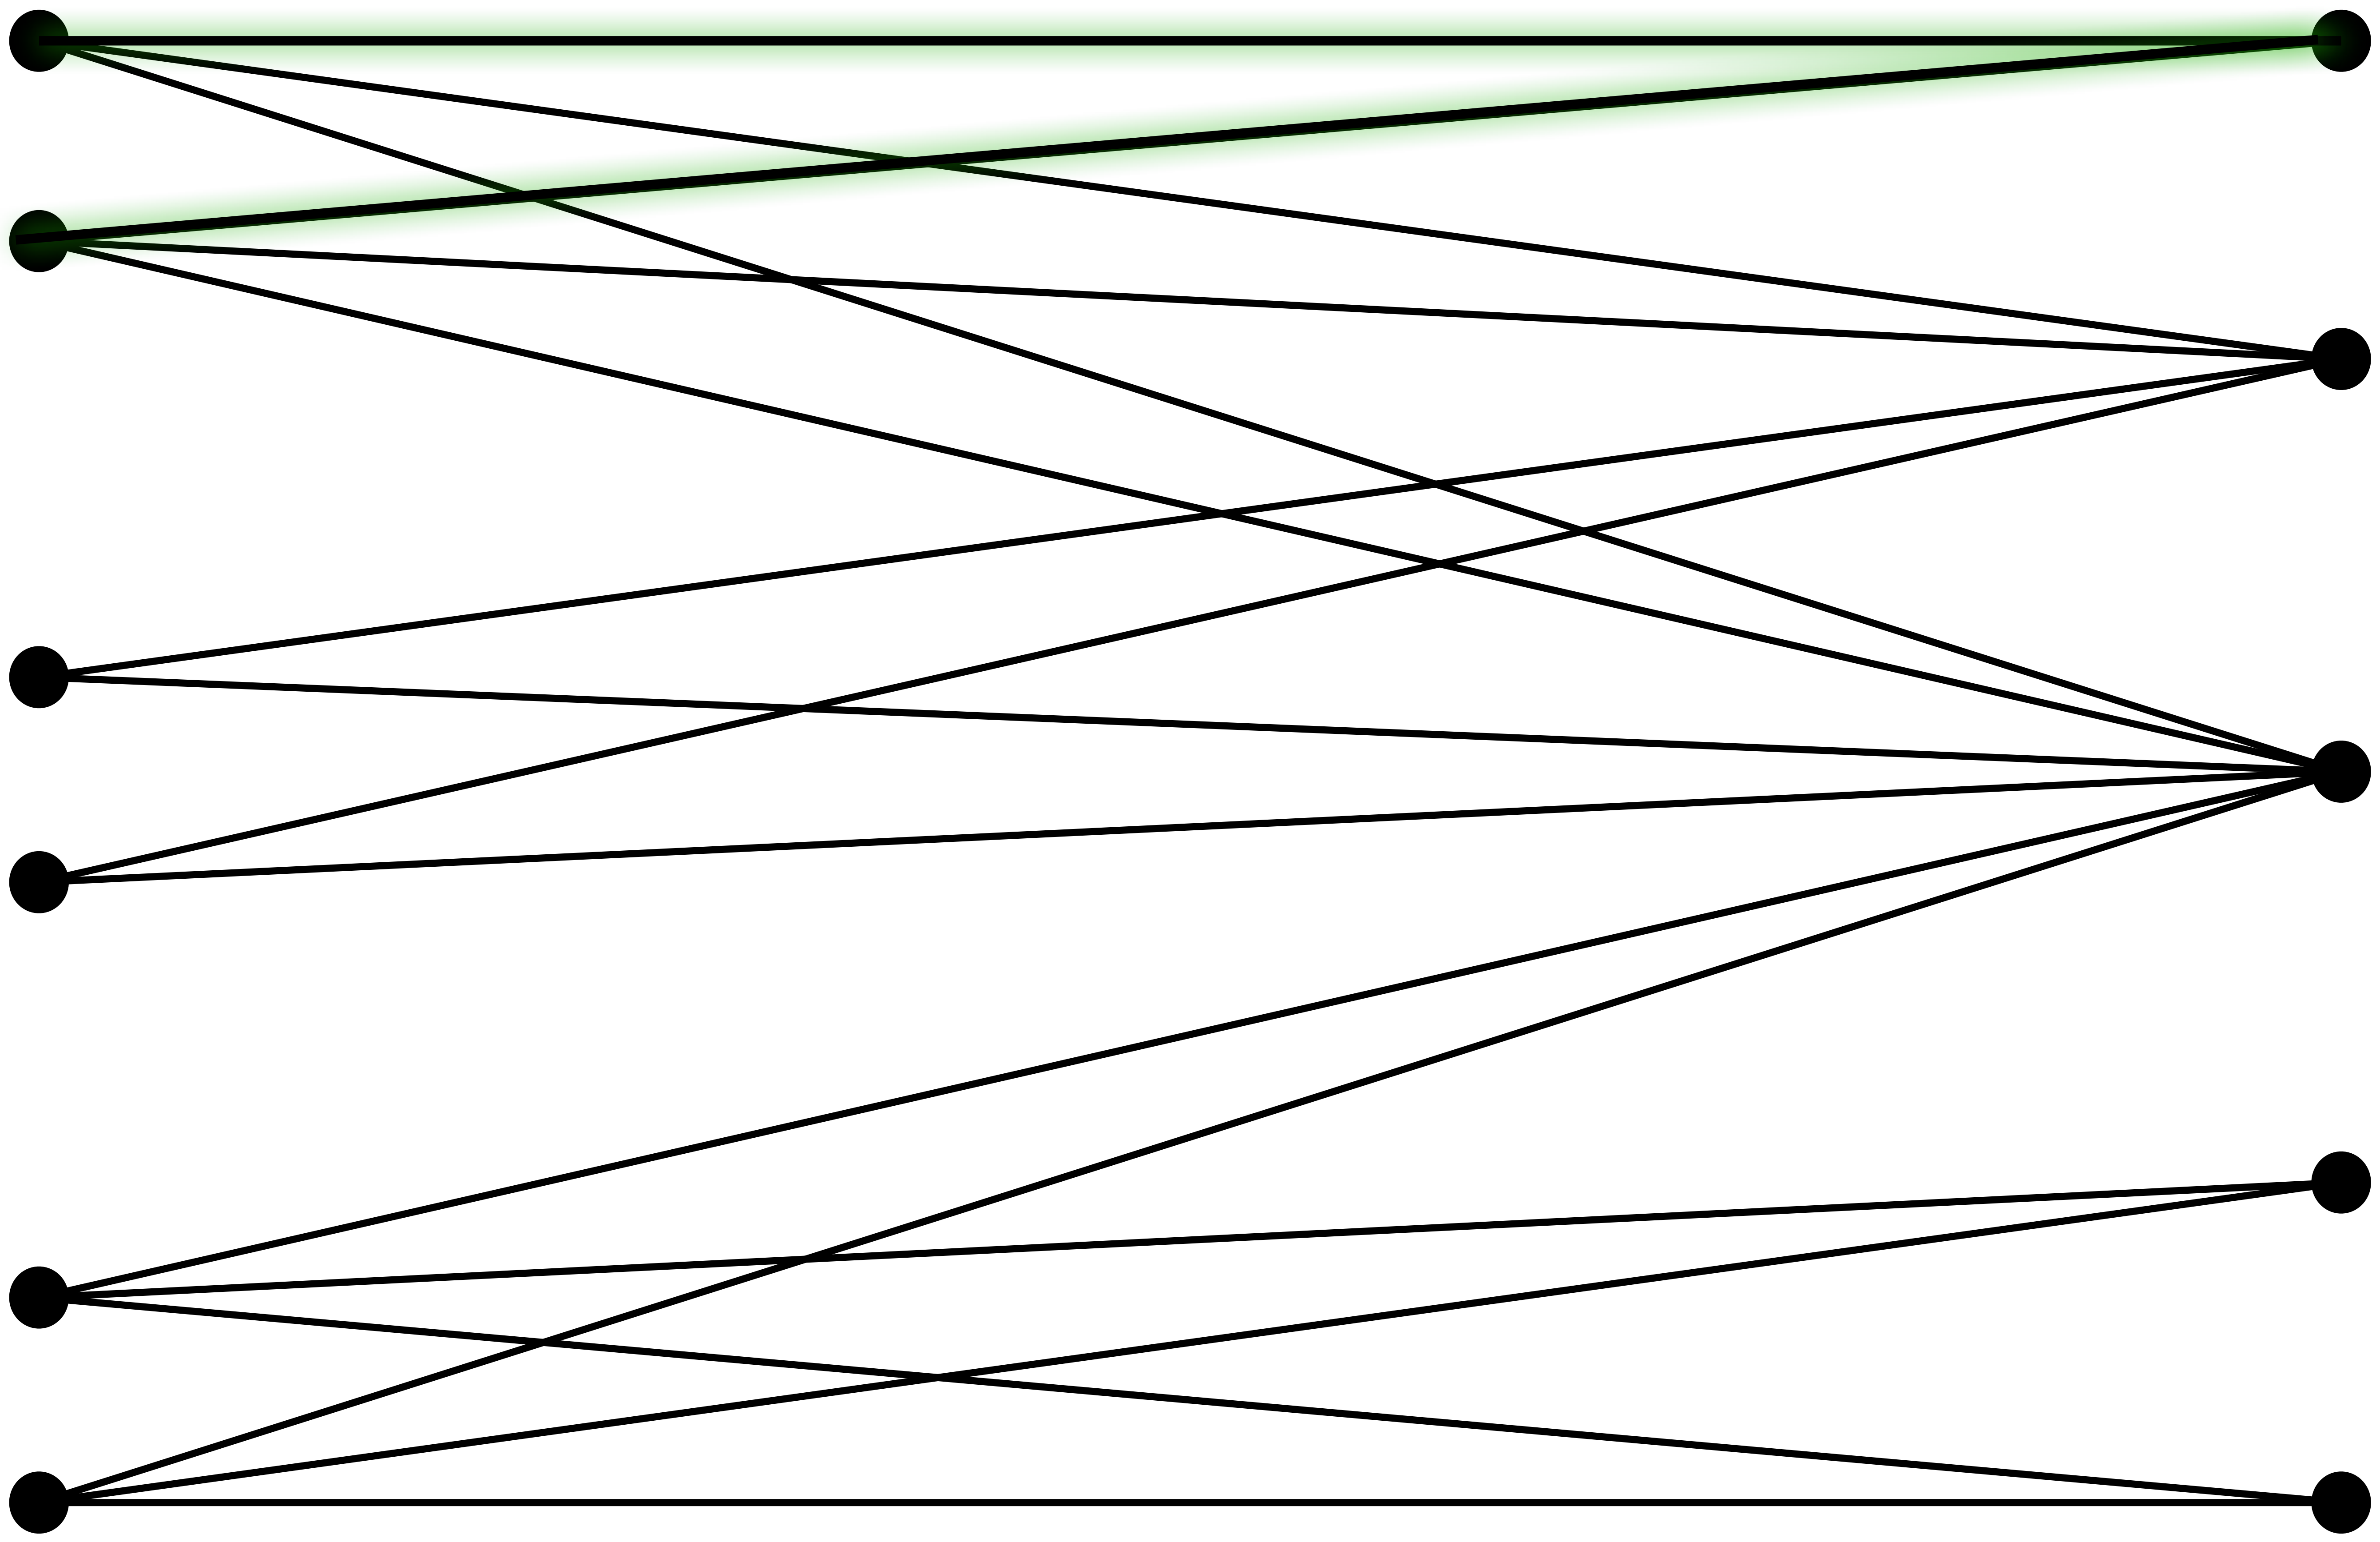

Why can't all my lightcones be huge?

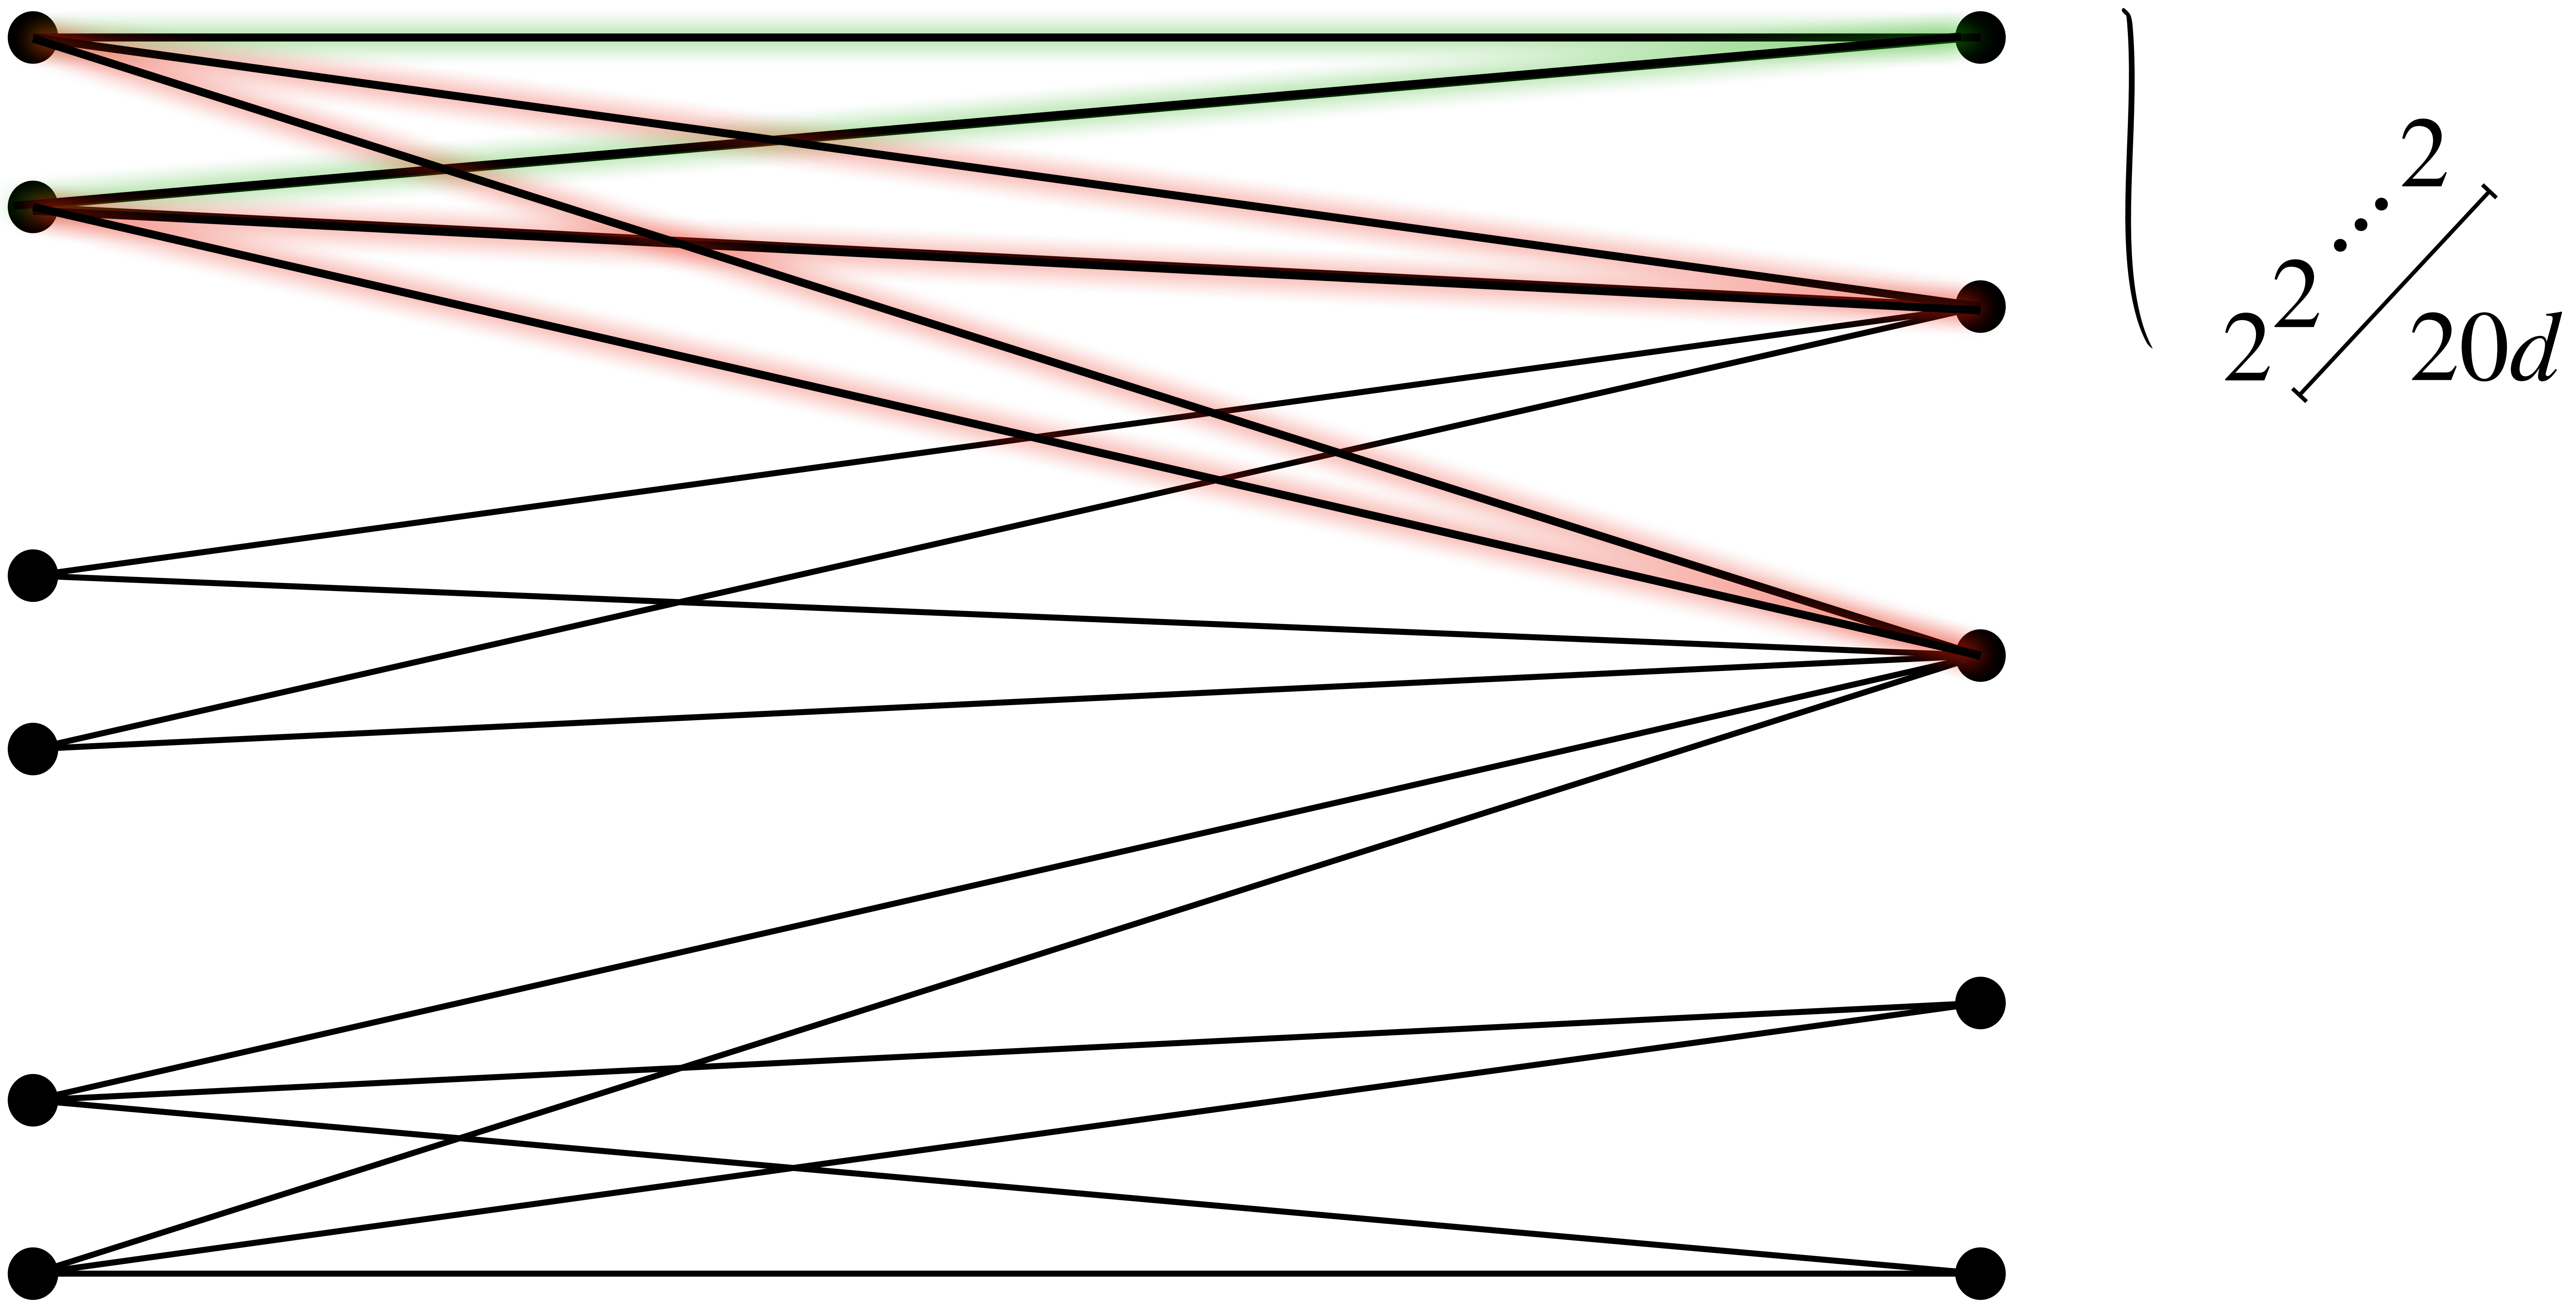

Theorem [KOW 24]**:** Exist conditionings to find many disjoint neighborhoods

Open questions

Question: Can you improve the sampling lower bound to AC^0

- Need a new candidate hard distribution
- Still open for QAC^0 circuits

Question: Can we get stronger separations for other sorts of problems?

- Theorem [G, Schaeffer]: Interactive sampling requires NC^1 circuits