“From EPR to BQP”

Quantum Computing as 21st-Century Bell Inequality Violation

Scott Aaronson (MIT)

Why Quantum Computing Is Like Bell Inequality Violation

Revolutionary insight about **what can be done** using QM—and about what **can't** be done by any classical simulation of some kind

At one level, "just" a logical consequence of 1920s QM—yet wasn't discovered till decades afterward

Sheds light on murky philosophical issues ("spooky action at a distance" / "huge size of configuration space") by **operationalizing** the issues

Challenges an "obvious" classical assumption (Local Hidden Variables / Extended Church-Turing Thesis)

Why Quantum Computing Is Like Bell Inequality Violation

Bell: People think it lets you signal faster than light

QC: People think it lets you solve **NP**-complete problems

But the truth is subtler! (You can "merely" win CHSH 85% of the time / factor integers)

Even in QM, signaling is still impossible, and **NP**-complete problems are still believed to be hard

Tsirelson bound, collision lower bound, etc. constrain QM even more sharply

Classically, the resources needed to win CHSH could also signal, while those needed to factor could also solve **NP**-complete problems. But quantum is different!

Why Quantum Computing Is Like Bell Inequality Violation

Immediately suggests an **experiment**—one that's beyond the technology at the time it's proposed, but not obviously beyond the technology of a few decades later

Some: "Ho-hum, the outcome will just confirm QM"

Others: "This is so crazy, it amounts to a proof that new physical principles have to prevent it"

Even after an experiment is done, it remains to close various loopholes. (For example, related to the use of **postselection**)

Ah, but quantum computing is (supposed to be) useful! Isn't that an important difference?

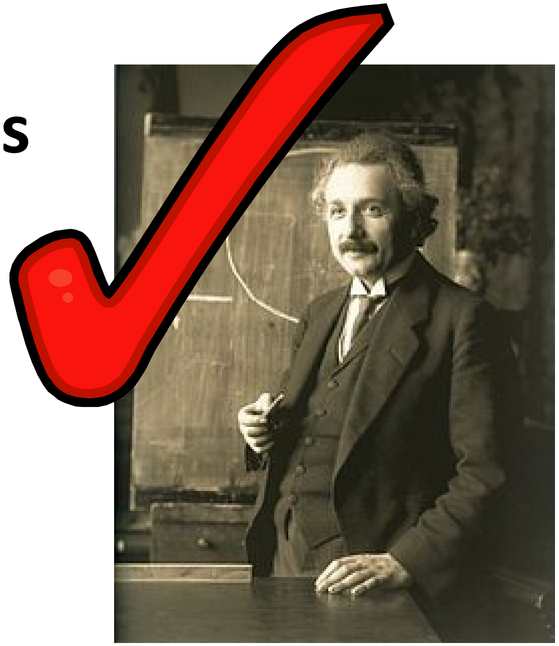

Einstein-certified random numbers

01010110000101111110

Device-independent QKD

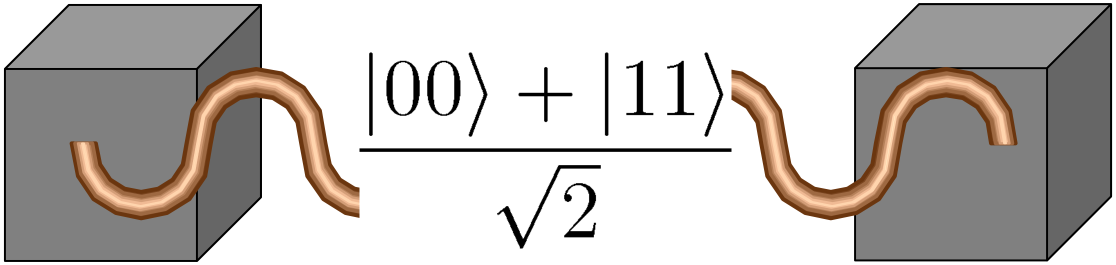

$$\frac{|00\rangle + |11\rangle}{\sqrt{2}}$$

Turns out Bell inequality violation is useful too!

OK, suppose we bought this analogy. So what? What would we do differently?

My Claim: The analogy with Bell's Inequality helps us focus on what's essential for QC experiments (at present), and away from what's nice but inessential

Nice But Inessential:

Universality

Practical applications

Clever quantum algorithms

"Traditional" types of problem

Essential:

Evidence that a classical computer can't do equally well

For me, focus on this issue is the **defining attribute** of quantum computer **science**

BosonSampling (A.-Arkhipov 2011)

A rudimentary type of quantum computing, involving only **non-interacting photons**

Classical counterpart:
Galton's Board

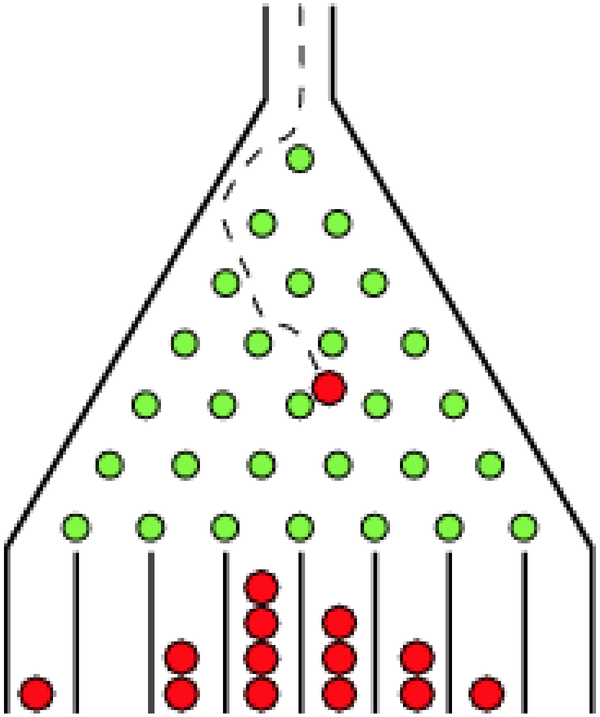

Replacing the balls by photons leads to famously counterintuitive phenomena, like the **Hong-Ou-Mandel dip**

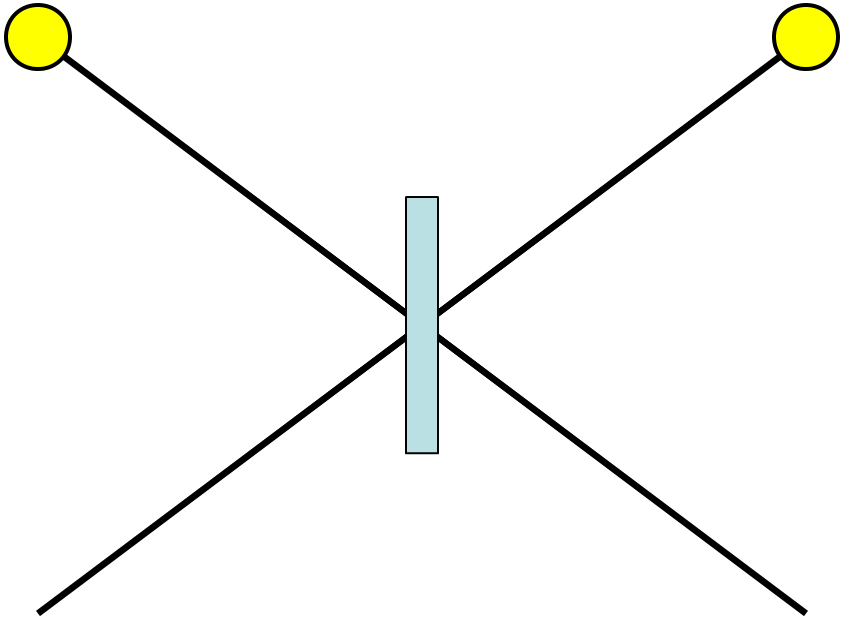

In general, we consider a network of beamsplitters, with n input modes and m≥n output modes (typically $m \sim n^2$)

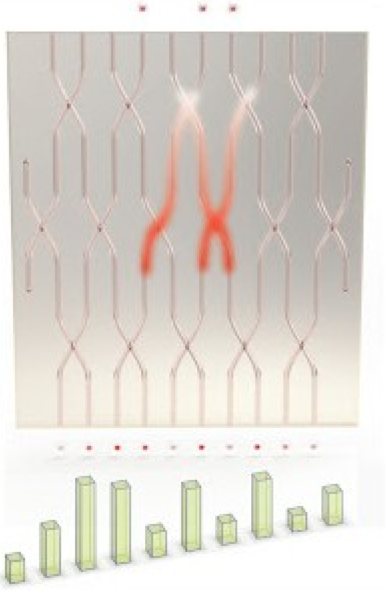

n single-photon Fock states enter

Assume for simplicity they all leave in different modes—there are $\binom{m}{n}$ possibilities

The beamsplitter network defines a column-orthonormal matrix $A \in C^{m\times n}$, such that $$\Pr[\text{outcome } S] = |\mathrm{Per}(A_S)|^2$$

where $$\mathrm{Per}(X) = \sum_{\sigma \in S_n} \prod_{i=1}^{n} x_{i,\sigma(i)}$$

n×n submatrix of A corresponding to S

is the matrix permanent

For simplicity, I'm ignoring outputs with ≥2 photons per mode

Example

For Hong-Ou-Mandel experiment,

$$\Pr\left[\text{output } |1,1\rangle\right] = \left|\operatorname{Per}\begin{pmatrix} \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{2}} \end{pmatrix}\right|^2 = \left|\frac{1}{2} - \frac{1}{2}\right|^2 = 0$$

In general, an n×n complex permanent is a sum of n! terms, almost all of which cancel out

How hard is it to estimate the “tiny residue” left over?

Answer: **#P**-complete. As hard as any combinatorial counting problem, and even harder than **NP**-complete!

So, Can We Use Quantum Optics to Solve a #P-Complete Problem?

That sounds way too good to be true…

Explanation: If X is sub-unitary, then $|\text{Per}(X)|^2$ will usually be exponentially small. So to get a reasonable estimate of $|\text{Per}(X)|^2$ for a given X, we'll generally need to repeat the optical experiment exponentially many times

Better idea: Given A∈ $C^{m\times n}$ as input, let **BosonSampling** be the problem of merely *sampling* from the same permanental probability distribution D_A that the beamsplitter network samples from

Theorem (A.-Arkhipov 2011): Suppose BosonSampling is solvable in classical polynomial time. Then **$P^{\#P}$=BPP^{NP}**

Harder Theorem: Suppose we can sample D_A even *approximately* in classical polynomial time. Then in **BPP^{NP}**, it's possible to estimate Per(X), with high probability over a Gaussian random matrix $X \sim N(0,1)_C^{n\times n}$

Upshot: Compared to (say) Shor's factoring algorithm, we get *different/stronger* evidence that a *weaker* system can do something classically hard

Experiments

Last year, groups in Brisbane, Oxford, Rome, and Vienna reported the first 3-photon BosonSampling experiments, confirming that the amplitudes were given by 3x3 permanents

# of experiments > # of photons!

Was there "cheating" (reliance on postselection)? Sure! Just like in many other current quantum computing experiments…

Obvious Challenges for Scaling Up:

-Reliable single-photon sources (optical multiplexing?)

-Minimizing losses

-Getting high probability of n-photon coincidence

Goal (in our view): Scale to 10-30 photons

Don't want to scale much beyond that—both because

(1)you probably can't without fault-tolerance, **and**

(2)a classical computer probably couldn't even verify the results!

Theoretical Challenge: Show that, even with (say) Gaussian inputs or modest photon losses, you're still solving a classically-intractable sampling problem

Recent Criticisms of Gogolin et al.

(arXiv:1306.3995)

Suppose you ignore which actual photodetectors light up, and count only the **number of times** each output configuration occurs. In that case, the BosonSampling distribution D_A is exponentially-close to the uniform distribution U

Response: Dude, *why on earth* would you ignore which detectors light up??

The output of Shor’s factoring algorithm is also gobbledygook if you ignore the order of the output bits…

Recent Criticisms of Gogolin et al.
(arXiv:1306.3995)

OK, so maybe D_A isn't close to uniform. Still, the very same arguments [A.-Arkhipov] gave for why polynomial-time classical algorithms can't sample D_A, suggest that they can't even **distinguish** D_A from U!

Response: Dude, that's exactly why we said to focus on 10-30 photons—a range where a classical computer **can** verify a BosonSampling device's output, but the BosonSampling device might be "faster"!
(And 10-30 photons is probably the best you can do anyway, without quantum fault-tolerance)

Even More Decisive Responses

(paper in preparation)

Theorem (A. 2013): Let A∈ $C^{m\times n}$ be a Haar-random BosonSampling matrix, where $m>>n^2$. Then with overwhelming probability over A, the BosonSampling distribution D_A has variation distance at least 0.313 from the uniform distribution U

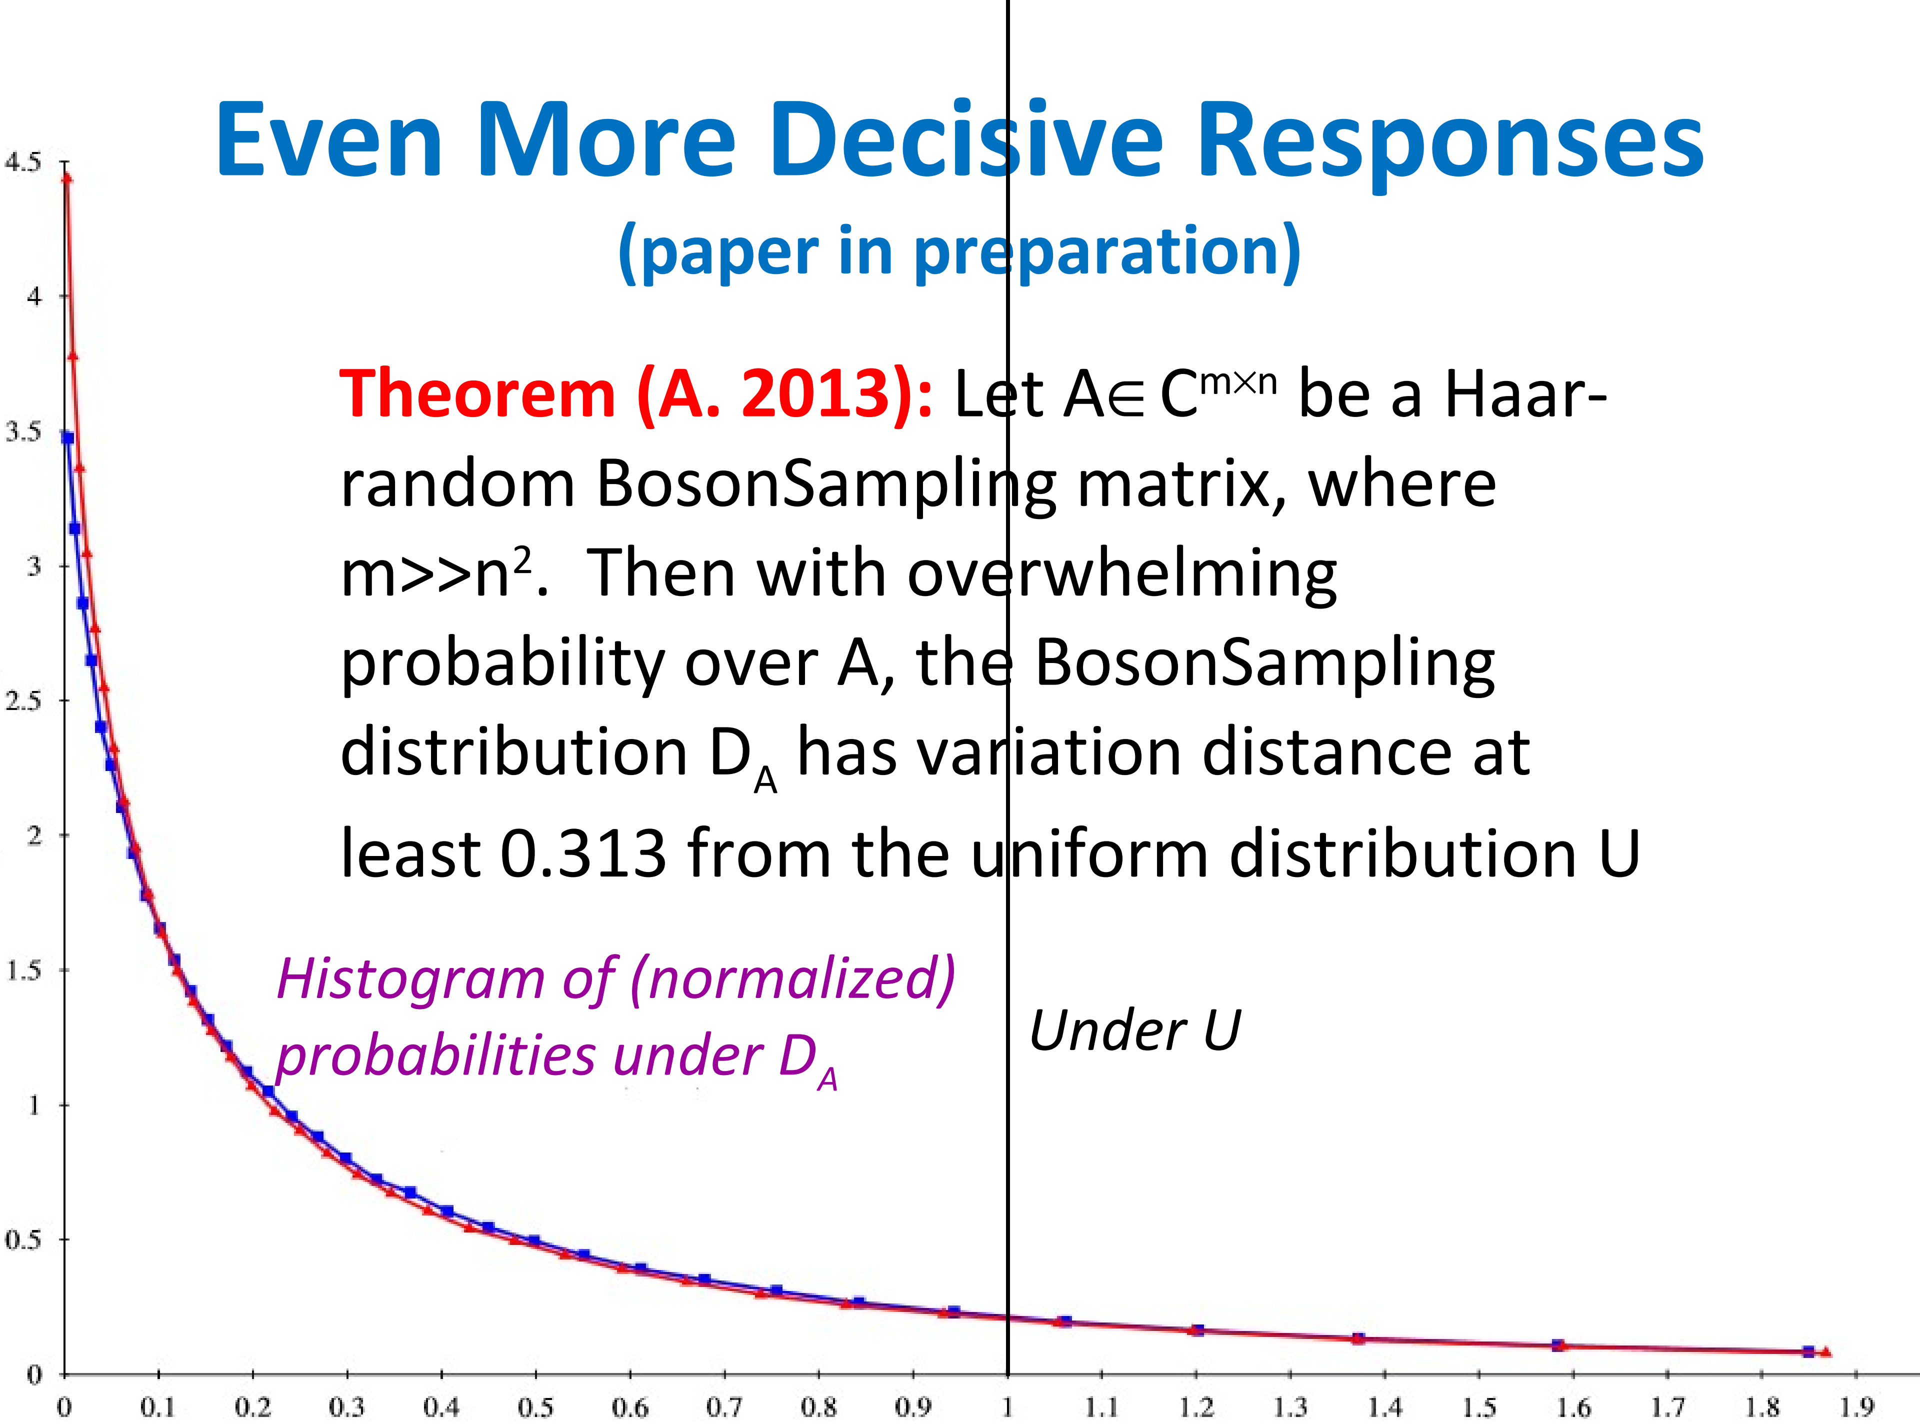

Histogram of (normalized) probabilities under D_A

Under U

Theorem (A. 2013): Let $A \in C^{m \times n}$ be Haar-random, where $m >> n^2$. Then there **is** a classical polynomial-time algorithm C(A) that distinguishes D_A from U (with high probability over A and constant bias, and using only O(1) samples)

Strategy: Let A_S be the n×n submatrix of A corresponding to output S. Let P be the product of squared 2-norms of A_S's rows. If P>E[P], then guess S was drawn from D_A; otherwise guess S was drawn from U

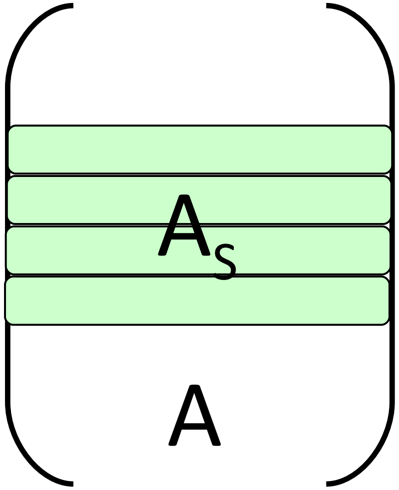

$$P = \left\|v_1\right\|^2 \cdots \left\|v_n\right\|^2 \geq \left(\frac{n}{m}\right)^n ?$$

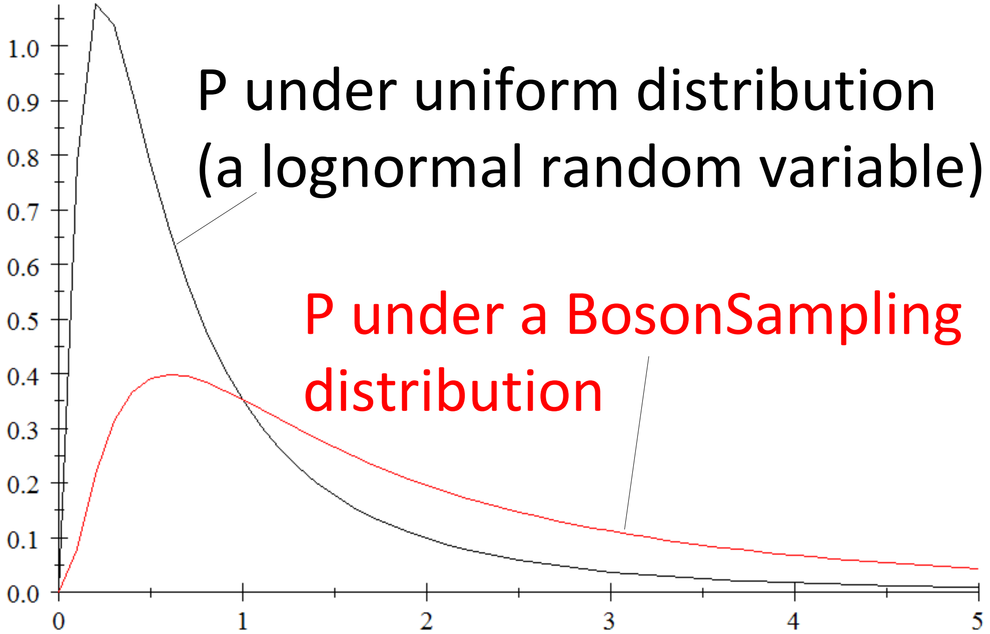

P under uniform distribution (a lognormal random variable)

P under a BosonSampling distribution

Summary

I advocate that our community approach QC experiments as we approached the Bell experiments: as an exciting scientific quest to rule out “polynomial-time hidden-variable theories”

(with any practical applications a “bonus” for later)

This perspective is **constraining**: It puts the question of classical hardness front and center

But mostly it’s **liberating**: It means we can aim, not only for universal QC, but for *any quantum system whatsoever* that does *anything* that we can argue is asymptotically hard to simulate classically

BosonSampling is just one example of what this perspective can lead us to think about. I expect many more!