Peer Review File

Manuscript Title: Phase Transitions in Random Circuit Sampling

Reviewer Comments & Author Rebuttals

Reviewer Reports on the Initial Version:

Referees' comments:

Referee #1 (Remarks to the Author):

This paper describes non-analyticities in an observable, the Cross-Entropy Benchmark (XEB). XEB has been used as a proxy for the fidelity of quantum simulation that is accessible in experiment.

Recently, XEB played a prominent role in discussions of quantum advantage experiments. This paper makes a significant contribution by showing the existence of a novel phase transition in how XEB is related to the fidelity changes and giving definite experimental signatures of this phase transition.

However, despite this advance, the current form of the paper does not look suitable for Nature. The reasons are the following. They are about the theoretical part of discussion. I will leave it to other referees to comment on experimental part.

* First, the presentation is not broadly accessible. Even the beginning of the paper assumes a lot of concepts familiar to experts.

* Many of the theoretical arguments in actual paper seem to be heuristic arguments. It relies on extensive supplemental materials to make precise claims. This would be OK if the supplemental gave detail to backup ideas in the paper. However, it seems that many arguments in the supplemental use totally different logic to the paper. As a result, reading the actual paper only gives a limited idea of why the claims are true.

* The claims made for the physical meaning of the phase transition are another question. As far as I can from the paper, this phase transition is established only as a property of a specific observable (XEB). I do not find evidence that this is a phase transition of the dynamics in general sense. However, it seem the strong claim is made in many places in the paper: last sentence of abstract, second paragraphe etc.

There are smaller clarity issue, eg the phase transition is not defined clearly enough and is confusing. See below specific comments.

* The paper accessibility would benefit from an early-on statement explaining the logic for introducing the XEB.

* The abstract talks about phase transitions observable with the XEB and transitions to a stable complex phase. Is it justified that the phase transition has a meaning beyond XEB for a complex phase? If this is not demonstrated in the paper then the claims should be accordingly moderated.

* First paragraph typo "quantum quantum correlations"

* "The reason is that RCS circuits are optimized to maximize the speed of quantum quantum correlations [2, 18, 19] while preventing potential simplifications in the correspond- ing classical emulations [17]."
I do not understand the sentence. The spread of quantum correlations has been studied in some random circuits and the spreading speed is not-universal, I do not know it to be optimized or maximized:
(1) Phys. Rev. X 8, 021013, 2018
(2) Phys. Rev. X 8, 021014, 2018
These references I think are also relevant to some later theoratical parts of this work.

* Phase transitions occurs in the thermodynamic limit. Clarity would benefit if it is explained early what limit should be taken to see a sharp transition. In Fig 1 the parameter held fixed is d/log(n), in Fig 2 the parameter that is held fixed is d, on page 2 it says the transition becomes a discontinuity when d → ∞. It is not a problem that the paper discusses different limits but explanation is needed. Related, around equations (1,2) it should be clarified when trends with n or trends with d are being talked about.

* At the top of p2 the reference to "wavefunction" seems instead to mixed state. On second colum of same page, it says system "converges" without noise but it seems the state will keep evolving, not converge. End of same paragraph, I/2^n should be I/2^(n/2).

* equation (2) needs more explanation for how the entangling operation affects XEB.

* footnote 27. It refers to [45], this is a related paper that came out on arxiv just after this one. The footnote says that [45] does "numerical study in the case of all-to-all connectivity", however it [45] also contain analytic results for the phase transition and one dimensional circuits. The authors should consider whether more informative reference to [45] is appropriate given the relations between papers.

* Later theoretical discussion: "We can describe both phase transitions mentioned above with a map to a model in statistical mechanics". This is not explained and supplemental material instead uses a different analysis.

* I think the Ising model and theoretical population dynamics analysis in supplemental both use similar tools to the calculations of out of time order correlators in (1), (2) above, in this case these and any other relevant references should be cited.

* "It is evident that these experiments fall well within the weak noise regime, satisfying the requirement to fully utilize the computational capacity of the noisy quantum processors" Related to the question

above, what is the actual relation between weak noise phase for XEB and this computational capacity requirement?

Referee #2 (Remarks to the Author):

This paper studies phase transitions in random circuit sampling (RCS) subject to noise, which can be diagnosed by measuring the XEB and its discrepancy with fidelity. In addition, a new RCS experiment with improved fidelity is demonstrated, leading to the state-of-the-art quantum supremacy experiment.

Here is some background: since Google's breakthrough experiment in 2019, the linear cross-entropy benchmark (XEB) has become a widely used method for benchmarking noisy quantum circuits and quantum simulators. There has been a lot of theoretical and experimental works arguing that XEB is a good proxy of fidelity in the low-noise regime, but it is not clear precisely when it works and why – an important question in quantum device benchmarking. This paper is the first to show that there is an underlying phase transition associated with XEB and to identify the boundary of the phase transition, thus giving XEB a solid physics justification and a deeper understanding of its applicable regime.

Separately, since 2019 there has been a continuing effort to push the boundary of quantum supremacy experiments. At the time there were two important aspects of the experiment: the first is practical hardness: how much classical computing resources does it take to simulate the experiment; the second is scalable hardness: if we keep a constant noise rate per gate (which is the eventual regime of fault-tolerant quantum computation) and scale the system bigger and bigger, will the experiment be exponentially harder and harder to simulate classically. Both questions are extremely important for understanding the boundary of quantum advantage. Much progress has been made on both questions since then. By today's standards, the 2019 RCS experiment is very easy to simulate classically; subsequent experiments by other groups are harder to simulate, but by current standards they are somewhat on the boundary of being simulable with modest resources. On the other front, theoretical works of Ref. [26] and others showed that there is no asymptotic hardness of RCS in the constant noise regime.

There is thus an urgent need for clarification: how to think about RCS and quantum supremacy experiments considering recent developments. This paper provides such a clarification.

First, there are new RCS experiments with 67 qubits and 32 cycles / 70 qubits and 24 cycles. These experiments are very impressive: while increasing the scale relative to USTC's prior experiment on 60 qubits, the overall circuit fidelity increased by an order of magnitude. This clearly puts the experiment far beyond today's classical simulation capabilities. Second, there are theoretical arguments showing that the phase transition associated with XEB can also be used to identify the boundary of quantum advantage. When total noise per cycle (eps*n where eps is noise per qubit and n is the number of qubits) is below some small constant threshold, it was argued that the XEB of a spoofing algorithm cannot match the XEB of the experiment; in addition, the actual experiment is shown to be in this low-noise regime. This gives an important clarification to the complexity of RCS experiments: we should think of them as being in the low-noise regime of small error per cycle, which holds for current small-scale experiments; moreover, it is plausible that in this regime the experiments are hard to spoof classically.

Therefore, I support the eventual publication of this paper, but I do have some major comments that need to be addressed. My overall opinion in these comments is that some of the claims need to be adjusted, and some of the presentations need to be improved/clarified.

1. An important reference is missing. [arxiv 2111.14907] gave rigorous theoretical bounds for the scaling of XEB, although it requires smaller noise than the phase transition boundary identified in this paper. It is very important that this paper be cited properly.

2. Another important reference is not cited properly. [arxiv 2005.02421] is the first to prove anti-concentration in 1D (including the exp(n e^(-d)) scaling) and should be cited whenever anti-concentration is discussed.

3. There may be too much emphasis on the “weak-link model” in the main text. While this model helps build up the intuition, I think it may be an oversimplification of the actual noisy RCS model: it is not obvious why the physics of the “weak-link model” should be a faithful representation of noisy RCS. For example, why only consider dividing the system into two subsystems instead of many subsystems (e.g. a 2D system is divided into many squares, each square is weakly linked to neighboring squares)?

Moreover, unless I misunderstood, there seems to already be some inconsistencies between the weak-link model and noisy RCS model shown in the paper. For example, the analysis at the bottom right corner of page 4 essentially says that the 2D behavior of the weak-link model is very different from 1D, while the statistical mechanics model from Appendix E says that the phase diagram in 2D is not much different from 1D (see Eq. (E32) and Eq. (E16)). A clarification is much needed on this point.

Finally, there are some claims about the physics of noisy RCS that may be inaccurate. A sentence reads “when the error rate per cycle is large, the wavefunction of the system could be approximately represented by multiple uncorrelated subsystems.” While this is evident from the weak-link model, I think it is far from obvious that actual noisy RCS follows this description. In fact, the result from Ref. [26] says that this is false: even in the high noise regime, dividing the system into subsystems is far from the actual noisy RCS because this approximation loses low-weight Pauli paths that cross the boundary between subsystems, and these Pauli paths are very important for a faithful approximation of noisy RCS. Therefore, I suggest a serious revision of this claim (and similar claims elsewhere).

4. Correspondingly, I think the statistical mechanics model in Appendix E deserves a much better emphasis and presentation. This model to me is a much better representation of noisy RCS, and the analysis in this section is very impressive: It is absolutely remarkable to me that a closed-form expression of the XEB can be found, even if it is under some approximations. Also, this analysis does not seem too complicated. I think it would really help the presentation of the paper if some elements of the

analysis can be in the main text (while shortening the weak-link model). In addition, there are two things that are missing:

First, there is a very interesting remark “The noise induced phase transition is driven by a control parameter (analogous to a magnetic field in an Ising model) that scales with the system size (number of qubits). This is loosely alike to Freederiks transitions in liquid crystals [29].” But it was not mentioned at all in the Appendix. I think it is important to expand on this point to really demonstrate the underlying physics, which is one of the main points of the paper.

Second, there is no plot to verify the formulas obtained in Appendix E. It is very important to plot the formulas against numerical/experimental data of noisy RCS, to see how well the approximations holds, and to address the inconsistency mentioned above.

5. I don’t understand Eq. (F7) – Eq. (F9) in Appendix F. This is a remarkable claim about what “general” spoofing algorithms can do, which says that they cannot achieve XEB better than “noiseless XEB - 1”.

Here’s an example of a simple spoofing algorithm, which is a simplification of Ref. [26]: just calculate all weight-(d+1) Pauli paths, which is the bare minimum that a classical spoofing algorithm can do. With Haar random 2-qubit gates, this gives XEB roughly n*(2/5)^d.

Comparing my formula with Eq. (F9), there are two differences:

First, the decay rate 0.4 is much bigger than the decay rate in Eq. (F9) which is 0.14. Second, there is a factor of n.

While the first difference could potentially be explained due to the gate set, the second difference is fundamental. It seems to me that “noiseless XEB - 1” is quite different from my model “calculate low weight Pauli paths”, which in general contains polynomial factors in n. Since we are talking about actual numbers in experiments, where d cannot be too large, this polynomial factor in n is very important and could contribute a lot.

I don’t think it’s necessary for the authors to fully address my point, but the above discussion shows that the claims made in Eq. (F7) – Eq. (F9) may not apply to *general* spoofing algorithms, and some adjustments to the claims are necessary.

6. The claim that the noisy RCS experiments fall into the low-noise regime is very important, but I don’t fully understand its justification. How is the threshold error per cycle eps*n = 0.47 determined? I cannot find any description of how to obtain this (arguably the most important) number.

More importantly, from the description of Fig. 3g, it seems that this threshold number is determined using the weak-link model. As discussed above, this model may not be very reliable for understanding noisy RCS and therefore using this model to calculate the threshold is problematic.

I understand that numerically simulating 2D RCS is hard, but Clifford RCS can be used as a proxy in this case to find the threshold, like Fig. 16.

7. Question about Table 1. A very important aspect for these estimates is that the classical simulation only needs to achieve a tiny fidelity (matching the noisy experiment) and do not need to fully simulate the ideal quantum circuit. The estimates in Table 1 do seem to take this into account, but I cannot find a description about how these numbers (noisy sampling simulation time) are calculated. An account of this estimation is necessary.

8. Suggestion about presentation: the abstract seems to be missing a key sentence that emphasizes the key point of the paper: the conclusion that there is evidence for hardness in the weak-noise phase, and current experiments are within this phase.

9. Typos: typo in paragraph 1 of main text, typo below Eq (E29).

Referee #3 (Remarks to the Author):

In the paper, the Google Quantum AI team has demonstrated that the Sycamore processor has reclaimed its position in the random circuit sampling problem, which remains the only task to date where a quantum advantage has been convincingly shown. The Sycamore experiment marked the first achievement of this milestone in the 2019 Nature paper. However, its supremacy has been challenged in recent years by advances in classical simulations and by the Zuchongzhi experiments conducted by the USTC team.

Upon reviewing the paper and its supplementary materials, I understand that the authors delineate three distinct phases in the random circuit sampling problem—localized, delocalized with strong noise, and delocalized with weak noise—each separated by phase transitions and each corresponding to different types of classical spoofing algorithms.

In the localized phase, cross-entropy benchmarking (XEB) fails to approximate fidelity, and the system is vulnerable to spoofing through method 1, the sub-space post-selection method, as discussed in Physical Review Letters 128, 030501.

In the delocalized phase with strong noise, the relevant classical spoofing approach is method 2, which involves subsystem approximation combined with post-processing, as described in arXiv:2112.01657.

The final phase, the delocalized phase with weak noise, is immune to both the sub-space and subsystem methods, leaving method 3, approximate tensor contractions with rejection sampling (e.g., Physical Review Letters 129, 090502), as the only viable spoofing technique.

The paper's second goal is to show that the latest Sycamore experiments fall within the delocalized phase with weak noise, where methods 2 and 3 are inapplicable, with method 3 being computationally infeasible. Furthermore, the paper suggests that Sycamore is significantly more challenging to simulate than Zuchongzhi using Method 3.

The concept of three phase transitions in random circuit sampling is intriguing, and the new experiments reinforce Sycamore's lead in the race for quantum supremacy, particularly over Zuchongzhi. While I appreciate the phase transition theory, the Zuchongzhi experiments also operate in the delocalized phase with weak noise, where classical spoofing has not yet succeeded. Given Sycamore's lack of exclusivity, it is challenging to assess whether this work meets Nature's publication criteria. It would be beneficial if the authors could clarify the substantial advantages of the Sycamore processor and its experiments over the Zuchongzhi processor.

Additionally, I have several points I would like the authors to address:

1. Regarding the tensor network contraction problem, the computational cost is estimated assuming infinite memory. Does this imply that only the contraction order matters, without considering dynamic slicing? Clarification is needed.

2. It would be useful to differentiate between classical spoofing method 2, which combines subsystem approximation and post-processing, and method 3, which uses tensor network contraction and rejection sampling. Method 2 yields high XEB but low fidelity, whereas method 3's XEB more closely approximates fidelity.

3. The paper's definition of FLOPs is unconventional and initially caused confusion. For instance, Table 1 presents two different FLOP definitions: "single precision complex FLOP" and "machine FLOPs." I suggest the authors replace "single precision complex FLOPs" with "time complexity" for clarity.

4. The term "spoofing" needs a clearer definition. How does one differentiate spoofing from simulation? If classical method 3 produces both high XEB and high fidelity, should it be considered spoofing or simulation?

5. I recommend providing more detailed descriptions of the classical algorithms in the supplementary material, ideally making the circuit parameters and the code for finding the contraction order and dynamic slicing publicly accessible. This would allow other researchers to verify the correctness of the algorithms and data in Table 1.

I believe these questions and comments should be addressed to facilitate a more informed decision on this paper's suitability for publication.

Author Rebuttals to Initial Comments:

Response to referee

Referee #1 (Remarks to the Author):

This paper describes non-analyticities in an observable, the Cross-Entropy Benchmark (XEB). XEB has been used as a proxy for the fidelity of quantum simulation that is accessible in experiment.

Recently, XEB played a prominent role in discussions of quantum advantage experiments. This paper makes a significant contribution by showing the existence of a novel phase transition in how XEB is related to the fidelity changes and giving definite experimental signatures of this phase transition.
We thank the reviewer for the careful reading and valuable comments which we address below.

However, despite this advance, the current form of the paper does not look suitable for Nature. The reasons are the following. They are about the theoretical part of discussion. I will leave it to other referees to comment on experimental part.

* First, the presentation is not broadly accessible. Even the beginning of the paper assumes a lot of concepts familiar to experts.
We have tried to make the manuscript more accessible by exposing definitions. We have also removed some more technical discussions from the main text.

* Many of the theoretical arguments in actual paper seem to be heuristic arguments. It relies on extensive supplemental materials to make precise claims. This would be OK if the supplemental gave detail to backup ideas in the paper. However, it seems that many arguments in the supplemental use totally different logic to the paper. As a result, reading the actual paper only gives a limited idea of why the claims are true.
We have shortened the discussion of the weak-link model and moved some of the explanation of its generalization to 2D from the SI into the main text. This clarifies the physics arguments used in the paper, and their connection between the weak-link model, its generalization, and previous results in anticoncentration and spoofing algorithms. In particular, see the new paragraphs that start with “In order to explain the nature of the noise induced phase transition for two and more dimensions,...”

* The claims made for the physical meaning of the phase transition are another question. As far as I can from the paper, this phase transition is established only as a property of a specific observable (XEB). I do not find evidence that this is a phase transition of the dynamics in

general sense. However, it seem the strong claim is made in many places in the paper: last sentence of abstract, second paragraphe etc.
We explain more generically this phase transition in the weak-link model and its generalization. This phase transition results from a competition between the convergence to a globally correlated state (the Porter-Thomas or ergodic state) and more noise resilient, but transient, local correlations. The new paragraphs in the main text explains this explicitly.

There are smaller clarity issue, eg the phase transition is not defined clearly enough and is confusing.
See below specific comments.

* The paper accessibility would benefit from an early-on statement explaining the logic for introducing the XEB.
We rewrote this sentence “We find that XEB is a proper observable to resolve the aforementioned regimes experimentally, as it is sensitive to the nature of the dominant correlations.”

* The abstract talks about phase transitions observable with the XEB and transitions to a stable complex phase. Is it justified that the phase transition has a meaning beyond XEB for a complex phase? If this is not demonstrated in the paper then the claims should be accordingly moderated.
See response above.

* First paragraph typo "quantum quantum correlations"
We fixed the typo, thank you.

* "The reason is that RCS circuits are optimized to maximize the speed of quantum quantum correlations [2, 18, 19] while preventing potential simplifications in the corresponding classical emulations [17]."
I do not understand the sentence. The spread of quantum correlations has been studied in some random circuits and the spreading speed is not-universal, I do not know it to be optimized or maximized:
(1) Phys. Rev. X 8, 021013, 2018
(2) Phys. Rev. X 8, 021014, 2018
These references I think are also relevant to some later theoretical parts of this work.
We thank the reviewer for pointing out this imprecision in our language. The speed is indeed not universal. The circuits used in the experiment were optimized to maximize this speed using iSWAPs (so that operator spreading proceeds with the light cone velocity), and optimized single qubit gates. We have adjusted our language in the manuscript accordingly and added an additional citation. The sentence now says: “The reason is that RCS circuits can be optimized to maximize the speed of quantum correlations with iSWAP-like gates…”

We thank the reviewer for reminding us of those references. We added them to the SI when we talk about the slower speed of Haar random gate sets and when we introduce population dynamics (see below).

* Phase transitions occurs in the thermodynamic limit. Clarity would benefit if it is explained early what limit should be taken to see a sharp transition. In Fig 1 the parameter held fixed is d/log(n), in Fig 2 the parameter that is held fixed is d, on page 2 it says the transition becomes a discontinuity when $d \to \infty$. It is not a problem that the paper discusses different limits but explanation is needed. Related, around equations (1,2) it should be clarified when trends with n or trends with d are being talked about.
A new paragraph in the main text makes this explicit: “The noise induced phase transition appears in the thermodynamical limit $n \to \infty$ at fixed ϵn and $d / \log n$ (see SI Sec. E for details).” See also Eq. 4 now in the main text.

* At the top of p2 the reference to "wavefunction" seems instead to mixed state. On second column of same page, it says system "converges" without noise but it seems the state will keep evolving, not converge. End of same paragraph, I/2^n should be I/2^(n/2).
We changed “wavefunction” to “state”
Although the system keeps evolving, it converges to a product state.
We fixed the I/2^n to I/2^{n/2}, thank you.

* equation (2) needs more explanation for how the entangling operation affects XEB.
We have changed the paragraph that introduces equation (2) to explain better how the entangling operation affects XEB

* footnote 27. It refers to [45], this is a related paper that came out on arxiv just after this one. The footnote says that [45] does "numerical study in the case of all-to-all connectivity", however it [45] also contain analytic results for the phase transition and one dimensional circuits. The authors should consider whether more informative reference to [45] is appropriate given the relations between papers.
We have updated the citation to the reference.

* Later theoretical discussion: "We can describe both phase transitions mentioned above with a map to a model in statistical mechanics". This is not explained and supplemental material instead uses a different analysis.
This model is now incorporated in the main text, in the paragraphs around Eq. 2 and Eq. 4

* I think the Ising model and theoretical population dynamics analysis in supplemental both use similar tools to the calculations of out of time order correlators in (1), (2) above, in this case these and any other relevant references should be cited.
We thank the referee for the suggestion, and added those references in the SI when introducing population dynamics.

* "It is evident that these experiments fall well within the weak noise regime, satisfying the requirement to fully utilize the computational capacity of the noisy quantum processors" Related to the question above, what is the actual relation between weak noise phase for XEB and this computational capacity requirement?

We added this sentence to the main text after Eq. (3) “[In the strong noise regime] the first term in the right hand side of Eq. (3) prevails and taking small k the state can be approximately represented by multiple uncorrelated subsystems”. This argument is expanded in the SI.

Referee #2 (Remarks to the Author):

This paper studies phase transitions in random circuit sampling (RCS) subject to noise, which can be diagnosed by measuring the XEB and its discrepancy with fidelity. In addition, a new RCS experiment with improved fidelity is demonstrated, leading to the state-of-the-art quantum supremacy experiment.

Here is some background: since Google’s breakthrough experiment in 2019, the linear cross-entropy benchmark (XEB) has become a widely used method for benchmarking noisy quantum circuits and quantum simulators. There has been a lot of theoretical and experimental works arguing that XEB is a good proxy of fidelity in the low-noise regime, but it is not clear precisely when it works and why – an important question in quantum device benchmarking. This paper is the first to show that there is an underlying phase transition associated with XEB and to identify the boundary of the phase transition, thus giving XEB a solid physics justification and a deeper understanding of its applicable regime.

Separately, since 2019 there has been a continuing effort to push the boundary of quantum supremacy experiments. At the time there were two important aspects of the experiment: the first is practical hardness: how much classical computing resources does it take to simulate the experiment; the second is scalable hardness: if we keep a constant noise rate per gate (which is the eventual regime of fault-tolerant quantum computation) and scale the system bigger and bigger, will the experiment be exponentially harder and harder to simulate classically. Both questions are extremely important for understanding the boundary of quantum advantage. Much progress has been made on both questions since then. By today’s standards, the 2019 RCS experiment is very easy to simulate classically; subsequent experiments by other groups are harder to simulate, but by current standards they are somewhat on the boundary of being simulable with modest resources. On the other front, theoretical works of Ref. [26] and others showed that there is no asymptotic hardness of RCS in the constant noise regime.

There is thus an urgent need for clarification: how to think about RCS and quantum supremacy experiments considering recent developments. This paper provides such a clarification.

First, there are new RCS experiments with 67 qubits and 32 cycles / 70 qubits and 24 cycles. These experiments are very impressive: while increasing the scale relative to USTC's prior experiment on 60 qubits, the overall circuit fidelity increased by an order of magnitude. This clearly puts the experiment far beyond today's classical simulation capabilities. Second, there are theoretical arguments showing that the phase transition associated with XEB can also be used to identify the boundary of quantum advantage. When total noise per cycle (eps*n where eps is noise per qubit and n is the number of qubits) is below some small constant threshold, it was argued that the XEB of a spoofing algorithm cannot match the XEB of the experiment; in addition, the actual experiment is shown to be in this low-noise regime. This gives an important clarification to the complexity of RCS experiments: we should think of them as being in the low-noise regime of small error per cycle, which holds for current small-scale experiments; moreover, it is plausible that in this regime the experiments are hard to spoof classically.

Therefore, I support the eventual publication of this paper, but I do have some major comments that need to be addressed. My overall opinion in these comments is that some of the claims need to be adjusted, and some of the presentations need to be improved/clarified.
We thank the referee for his support and the careful review of our manuscript. We have tried to address these comments and we detail our answer below.

1. An important reference is missing. [arxiv 2111.14907] gave rigorous theoretical bounds for the scaling of XEB, although it requires smaller noise than the phase transition boundary identified in this paper. It is very important that this paper be cited properly.
We thank the reviewer for the suggestion and we added that reference in the paragraph before Eq (2) when we use the corresponding noise model.

2. Another important reference is not cited properly. [arxiv 2005.02421] is the first to prove anti-concentration in 1D (including the exp(n e^(-d)) scaling) and should be cited whenever anti-concentration is discussed.
We thank the referee for the suggestion and cited that paper in the context of anticoncentration.

3. There may be too much emphasis on the "weak-link model" in the main text. While this model helps build up the intuition, I think it may be an oversimplification of the actual noisy RCS model: it is not obvious why the physics of the "weak-link model" should be a faithful representation of noisy RCS. For example, why only consider dividing the system into two subsystems instead of many subsystems (e.g. a 2D system is divided into many squares, each square is weakly linked to neighboring squares)?
We have added the generalization of the weak-link model to many subsystems to the main text. See paragraph starting with "In order to explain the nature of the noise induced phase transition for two and more dimensions, we extend the weak link model from two subsystems…"

Moreover, unless I misunderstood, there seems to already be some inconsistencies between the weak-link model and noisy RCS model shown in the paper. For example, the analysis at the bottom right corner of page 4 essentially says that the 2D behavior of the weak-link model is very different from 1D, while the statistical mechanics model from Appendix E says that the

phase diagram in 2D is not much different from 1D (see Eq. (E32) and Eq. (E16)). A clarification is much needed on this point.
We have added the generalization of the weak-link model to 2D to the main text.

Finally, there are some claims about the physics of noisy RCS that may be inaccurate. A sentence reads “when the error rate per cycle is large, the wavefunction of the system could be approximately represented by multiple uncorrelated subsystems.” While this is evident from the weak-link model, I think it is far from obvious that actual noisy RCS follows this description. In fact, the result from Ref. [26] says that this is false: even in the high noise regime, dividing the system into subsystems is far from the actual noisy RCS because this approximation loses low-weight Pauli paths that cross the boundary between subsystems, and these Pauli paths are very important for a faithful approximation of noisy RCS. Therefore, I suggest a serious revision of this claim (and similar claims elsewhere).
We thank this referee for this comment which has helped us clarify the exposition. Eq. (3) now added to the main text explains the generalization of the weak-link model and its relation to the Pauli paths approximation. We added a more technical explanation in a footnote: “Note that the fist term in the right hand side corresponds to the low-weight Pauli paths approximation in Ref. [28]”

4. Correspondingly, I think the statistical mechanics model in Appendix E deserves a much better emphasis and presentation. This model to me is a much better representation of noisy RCS, and the analysis in this section is very impressive: It is absolutely remarkable to me that a closed-form expression of the XEB can be found, even if it is under some approximations. Also, this analysis does not seem too complicated. I think it would really help the presentation of the paper if some elements of the analysis can be in the main text (while shortening the weak-link model). In addition, there are two things that are missing:
We thank the reviewer for this comment and have followed this advice. We added this model to the main text in the two paragraphs that start as “n order to explain the nature of the noise induced phase transition for two and more dimensions, we extend the weak link model.
We have also shortened some of the discussion of the weak-link model so that the paper is now shorter overall.

First, there is a very interesting remark “The noise induced phase transition is driven by a control parameter (analogous to a magnetic field in an Ising model) that scales with the system size (number of qubits). This is loosely alike to Freederiks transitions in liquid crystals [29].” But it was not mentioned at all in the Appendix. I think it is important to expand on this point to really demonstrate the underlying physics, which is one of the main points of the paper.
We added a new paragraph at the end of SI Sec. E2.

Second, there is no plot to verify the formulas obtained in Appendix E. It is very important to plot the formulas against numerical/experimental data of noisy RCS, to see how well the approximations holds, and to address the inconsistency mentioned above.

We have addressed the generalization of the weak-link model to 2D to the main text, as mentioned above. We have found that while numerics support the theory presented, we can not take the thermodynamic limit numerically in 2D, which is the main case of interest.

5. I don’t understand Eq. (F7) – Eq. (F9) in Appendix F. This is a remarkable claim about what “general” spoofing algorithms can do, which says that they cannot achieve XEB better than “noiseless XEB - 1”.

Here’s an example of a simple spoofing algorithm, which is a simplification of Ref. [26]: just calculate all weight-(d+1) Pauli paths, which is the bare minimum that a classical spoofing algorithm can do. With Haar random 2-qubit gates, this gives XEB roughly n*(2/5)^d.

Comparing my formula with Eq. (F9), there are two differences:

First, the decay rate 0.4 is much bigger than the decay rate in Eq. (F9) which is 0.14. Second, there is a factor of n.

While the first difference could potentially be explained due to the gate set, the second difference is fundamental. It seems to me that “noiseless XEB - 1” is quite different from my model “calculate low weight Pauli paths”, which in general contains polynomial factors in n. Since we are talking about actual numbers in experiments, where d cannot be too large, this polynomial factor in n is very important and could contribute a lot.

I don’t think it’s necessary for the authors to fully address my point, but the above discussion shows that the claims made in Eq. (F7) – Eq. (F9) may not apply to *general* spoofing algorithms, and some adjustments to the claims are necessary.
Both the pre-factor of, n, and the exponential, (2/5)^d, the referee brings up are accounted for in the expression “noiseless XEB - 1”. Explicitly, the expression n*(⅖)^d corresponds to the left hand side of Eq. 3 in the main text (which was previously only in the appendix), with k=1 and epsilon =0. The factor ⅖ for Haar random 2-qubit gates becomes ¼ for iSWAP. In this case epsilon=0 because the simulation does not include noise, and k=1 because this simulation would include only one Pauli per cycle. We now explain this in the main text and a footnote.
The argument for Eq. (F7) is essentially what the referee has in mind. Finite-depth correlations outside the Porter-Thomas state, as those captured by “low weight Pauli paths”, can be exploited by spoofing algorithms. They can be bounded by the XEB difference with the Porter-Thomas state, which is the meaning of F7.
Precisely because “we are talking about actual numbers”, as the referee says, we bound these finite-depth correlations numerically to make sure that we account for all of them, including potential finite-size and boundary effects. This is what we do in Sec. F2 of the SI.

6. The claim that the noisy RCS experiments fall into the low-noise regime is very important, but I don’t fully understand its justification. How is the threshold error per cycle eps*n = 0.47

determined? I cannot find any description of how to obtain this (arguably the most important) number.
The threshold of 0.47 is just a bound to the simulations presented in Fig. 3g, as seen in that figure. More importantly, as we said in the text “the noise induced phase transition for the discrete gate set used in the experiment occurs at higher noise rates, see SI Sec. F2. “

More importantly, from the description of Fig. 3g, it seems that this threshold number is determined using the weak-link model. As discussed above, this model may not be very reliable for understanding noisy RCS and therefore using this model to calculate the threshold is problematic.
The numerics in Fig. 3g do not have a weak link, we now made that explicit in the text “We numerically evaluate critical noise rates for systems of different sizes and circuit structures without weak links,...”

I understand that numerically simulating 2D RCS is hard, but Clifford RCS can be used as a proxy in this case to find the threshold, like Fig. 16.
Indeed the numerics in Fig. 16 is the reason why we say that “the noise induced phase transition for the discrete gate set used in the experiment occurs at higher noise rates, see SI Sec. F2. “. Clifford circuits have the same average XEB, so this proxy is reliable. While we feel that this Figure is important, we haven’t found a way to incorporate it in the main text as it introduces several new concepts.

7. Question about Table 1. A very important aspect for these estimates is that the classical simulation only needs to achieve a tiny fidelity (matching the noisy experiment) and do not need to fully simulate the ideal quantum circuit. The estimates in Table 1 do seem to take this into account, but I cannot find a description about how these numbers (noisy sampling simulation time) are calculated. An account of this estimation is necessary.
The estimates in Table 1 indeed take the corresponding fidelity into account. We added a paragraph to the SI explaining the known technique used: “In the context of simulating RCS, slices can also help reduce the computation time while reducing the fidelity of the output state….”

8. Suggestion about presentation: the abstract seems to be missing a key sentence that emphasizes the key point of the paper: the conclusion that there is evidence for hardness in the weak-noise phase, and current experiments are within this phase.
We thank the reviewer for the suggestion and changed this sentence in the abstract: “Furthermore, by presenting an RCS experiment in the weak noise phase with 67 qubits at 32 cycles, we demonstrate that the computational cost of our experiment is beyond the capabilities of existing classical supercomputers.”

9. Typos: typo in paragraph 1 of main text, typo below Eq (E29).
Thank you for pointing this out.

Referee #3 (Remarks to the Author):

In the paper, the Google Quantum AI team has demonstrated that the Sycamore processor has reclaimed its position in the random circuit sampling problem, which remains the only task to date where a quantum advantage has been convincingly shown. The Sycamore experiment marked the first achievement of this milestone in the 2019 Nature paper. However, its supremacy has been challenged in recent years by advances in classical simulations and by the Zuchongzhi experiments conducted by the USTC team.

Upon reviewing the paper and its supplementary materials, I understand that the authors delineate three distinct phases in the random circuit sampling problem—localized, delocalized with strong noise, and delocalized with weak noise—each separated by phase transitions and each corresponding to different types of classical spoofing algorithms.

In the localized phase, cross-entropy benchmarking (XEB) fails to approximate fidelity, and the system is vulnerable to spoofing through method 1, the sub-space post-selection method, as discussed in Physical Review Letters 128, 030501.

In the delocalized phase with strong noise, the relevant classical spoofing approach is method 2, which involves subsystem approximation combined with post-processing, as described in arXiv:2112.01657.

The final phase, the delocalized phase with weak noise, is immune to both the sub-space and subsystem methods, leaving method 3, approximate tensor contractions with rejection sampling (e.g., Physical Review Letters 129, 090502), as the only viable spoofing technique.

The paper's second goal is to show that the latest Sycamore experiments fall within the delocalized phase with weak noise, where methods 2 and 3 are inapplicable, with method 3 being computationally infeasible. Furthermore, the paper suggests that Sycamore is significantly more challenging to simulate than Zuchongzhi using Method 3.

The concept of three phase transitions in random circuit sampling is intriguing, and the new experiments reinforce Sycamore's lead in the race for quantum supremacy, particularly over Zuchongzhi. While I appreciate the phase transition theory, the Zuchongzhi experiments also operate in the delocalized phase with weak noise, where classical spoofing has not yet succeeded. Given Sycamore's lack of exclusivity, it is challenging to assess whether this work meets Nature's publication criteria. It would be beneficial if the authors could clarify the substantial advantages of the Sycamore processor and its experiments over the Zuchongzhi processor.
We thank the reviewer for the careful study and multiple suggestions which we address below.
Zuchongzhi 2.1 achieves 0.0366% fidelity with 60-qubit and 24-cycle. We achieve 0.1% fidelity with 67 qubits and 32 cycles, which sets this experiment comfortably in the classically

intractable regime. We added this sentence: “This increased depth is possible thanks to the substantially reduced errors compared with previous processors.”

Additionally, I have several points I would like the authors to address:

1. Regarding the tensor network contraction problem, the computational cost is estimated assuming infinite memory. Does this imply that only the contraction order matters, without considering dynamic slicing? Clarification is needed.
That’s correct. We have added the clarification to the main text:
“In Fig. 4c we show the time complexity or FLOP count (the number of real multiplications and additions) as a function of number of qubits and cycles required to compute a single amplitude at the output of a random circuit without memory constraints, i.e., optimizing exclusively the contraction ordering of the underlying tensor network. ”.
We also changed this sentence in the SI: “Tensor network contraction orderings is the only optimization needed for the case without memory constraints. When taking into account finite memory, the choice of sliced indices, sparsification of the output of the circuit, and reuse of intermediate computation across slicing instances, are all taken into account simultaneously in the optimization.”

2. It would be useful to differentiate between classical spoofing method 2, which combines subsystem approximation and post-processing, and method 3, which uses tensor network contraction and rejection sampling. Method 2 yields high XEB but low fidelity, whereas method 3's XEB more closely approximates fidelity.
See below for the definition of spoofing. We study post-processing in SI Sec. F3 and F4.

3. The paper's definition of FLOPs is unconventional and initially caused confusion. For instance, Table 1 presents two different FLOP definitions: "single precision complex FLOP" and "machine FLOPs." I suggest the authors replace "single precision complex FLOPs" with "time complexity" for clarity.
We thank the referee for this suggestion. We now use actual FLOP values everywhere for simplicity (we have multiplied the “complex FLOP” values by 8 appropriately).

4. The term "spoofing" needs a clearer definition. How does one differentiate spoofing from simulation? If classical method 3 produces both high XEB and high fidelity, should it be considered spoofing or simulation?
We reserve the term “spoofing” for “classical algorithms that represent only part of the system at a time”, as we say in the text. We have improved the citations of this sentence, and added two paragraphs that explain why this is possible in the strong noise regime. See for instance the explanation of Eq. (3).

5. I recommend providing more detailed descriptions of the classical algorithms in the supplementary material, ideally making the circuit parameters and the code for finding the

contraction order and dynamic slicing publicly accessible. This would allow other researchers to verify the correctness of the algorithms and data in Table 1.
We are happy to make the explicit contraction orderings and slides available to any interested party.

I believe these questions and comments should be addressed to facilitate a more informed decision on this paper's suitability for publication.

Reviewer Reports on the First Revision:

Referees' comments:

Referee #1 (Remarks to the Author):

The main update relevant to my report is near equation (3).
The authors have tried to keep discussion here streamlined, but as a result it is confusing.
It has been presented as extension of weak link model, but it seems new concepts are needed.

In order to parallel the discussion of the weak link model, equation (3) is described as the expansion of steady state (physical steady state). However, I question whether such an interpretation makes sense. In the SI we see this is not an expansion of the physical steady state, instead the steady state of averaged 2-replica problem.
This needs to be explained, or the reader takes away an understanding which is not correct.

" We explain more generically this phase transition in the weak-link model and its generalization. This phase transition results from a competition between the convergence to a globally correlated state (the Porter-Thomas or ergodic state) and more noise resilient, but transient, local correlations. The new paragraphs in the main text explains this explicitly."

If my above understanding is correct this convergence is about 2-replica problem. Therefore, my question remains about whether XEB represents a true phase transition of dynamics, or only a specific observable.

Referee #2 (Remarks to the Author):

My comments are satisfactorily addressed, and the presentation of the manuscript has been significantly improved due to the inclusion of Eqs (3)(4) and other clarifications.

Also, I didn’t appreciate the intuition that “noiseless XEB - 1” is roughly the same thing as low-weight Pauli paths – I think I get it now.

Now I am happy to recommend this manuscript for publication.

Referee #3 (Remarks to the Author):

In their responses, the authors clarified that the new Sycamore achieves 0.1% fidelity with 67 qubits and 32 cycles, which is clearly more advanced than Zuchongzhi 2.1, which achieves 0.0366% fidelity with 60 qubits and 24 cycles. The authors also addressed my comments on the definition of FLOPs and spoofing, and they have updated the manuscript accordingly. I would suggest accepting the manuscript if the authors could eventually make the circuit parameters and the code for finding the contraction order and dynamic slicing public.

Author Rebuttals to First Revision:

Response to referees' comments:

Referee #1 (Remarks to the Author):

The main update relevant to my report is near equation (3).
The authors have tried to keep discussion here streamlined, but as a result it is confusing.
It has been presented as extension of weak link model, but it seems new concepts are needed.

In order to parallel the discussion of the weak link model, equation (3) is described as the expansion of steady state (physical steady state). However, I question whether such an interpretation makes sense. In the SI we see this is not an expansion of the physical steady state, instead the steady state of averaged 2-replica problem.
This needs to be explained, or the reader takes away an understanding which is not correct.

" We explain more generically this phase transition in the weak-link model and its generalization. This phase transition results from a competition between the convergence to a globally correlated state (the Porter-Thomas or ergodic state) and more noise resilient, but transient, local correlations. The new paragraphs in the main text explains this explicitly."

If my above understanding is correct this convergence is about 2-replica problem. Therefore, my question remains about whether XEB represents a true phase transition of dynamics, or only a specific observable.

> At the noise induced phase transition the properties of the wave function change discontinuously. Similarly to the many-body localization transition it manifests in the statistics of the wave function rather than more familiar spectral characteristics. The simplest but not the only observable reflecting this change in the wave function is the XEB. Indeed, as the referee states, we show a discontinuity in the formalism of population dynamics, which the referee calls the averaged 2-replica problem. Therefore, there is a discontinuity in quantities that depend on the second moment, such as XEB. We expect a similar discontinuity in higher order quantities. Note that we use this formalism both for the weak-link model and the more general. Nevertheless, in all cases we try to keep the discussion in the main text as accessible as possible. Therefore, we have added two footnotes in the main text to clarify that we are using population dynamics to formally justify the equations in the text.
> The footnote before Eq. 2 says "We justify this equation formally averaging circuit instances over 2-replicas, which results in the so-called population dynamics

formalism.”. The footnote before Eq 3. says “This equation is derived using the formalism of population dynamics.”

Referee #2 (Remarks to the Author):

My comments are satisfactorily addressed, and the presentation of the manuscript has been significantly improved due to the inclusion of Eqs (3)(4) and other clarifications.

Also, I didn’t appreciate the intuition that “noiseless XEB - 1” is roughly the same thing as low-weight Pauli paths – I think I get it now.

Now I am happy to recommend this manuscript for publication.

Referee #3 (Remarks to the Author):

In their responses, the authors clarified that the new Sycamore achieves 0.1% fidelity with 67 qubits and 32 cycles, which is clearly more advanced than Zuchongzhi 2.1, which achieves 0.0366% fidelity with 60 qubits and 24 cycles. The authors also addressed my comments on the definition of FLOPs and spoofing, and they have updated the manuscript accordingly. I would suggest accepting the manuscript if the authors could eventually make the circuit parameters and the code for finding the contraction order and dynamic slicing public.

We are open sourcing the corresponding code.