

Tsirelson bounds for quantum correlations with indefinite causal order

Corresponding Author: Professor Giulio Chiribella

This file contains all reviewer reports in order by version, followed by all author rebuttals in order by version.

Version 0:

Reviewer comments:

Reviewer #1

(Remarks to the Author)

In this paper the authors provide a method to compute the upper bound of causal inequalities arising from quantum causal models.

The quantum causal modeling framework, differently to its classical counterpart, allows a much more general class of causal structures, that cannot be reduced to probabilistic mixtures of classical ones.

For this reason this kind of structures are said to represent a set operations having indefinite causal order (ICO).

Conversely causal inequalities are a set of constraints implied by imposing definite causal order.

The interest in this kind of constraints derive from the fact that, similarly to what happens with the device-independent certification of entanglement in Bell-like scenario, a violation of such inequality, would reveal the presence of ICO only by looking at the output statistics.

Possible violations of causal inequalities have been identified, but there it is still not clear whether they can represent physical processes.

The question answered by the authors is how to find the maximal value, or more generally an upper bound, for a given causal inequality for quantum causal structures.

While any explicit realization always represent a valid lower bound, the question of finding an upper bound to the set for a general scenario is an interesting one that has remained open until now.

To obtain this upper bound the authors first solve the problem for a specific class of inequalities, single-trigger inequalities, that they show can be maximized by a single SDP.

Then by using this result, they show that one can obtain an upper bound for general inequalities, by decomposing them in terms of single-trigger ones.

As a consequence this upper bound is also computable by single (albeit larger) SDP, which represent a relaxation of the original problem.

Using this method they obtain several results for known causal inequalities based on variations of "guess your neighbour input" games.

In particular I found very neat how the analysis of from the dual SDP can obtain analytical bounds and reconstruct the geometry of the definite and indefinite causal order correlations sets in the OCB inequality case, and the LGYNI inequality.

Moreover they also explore the question of the gap between the ICO and the logical maximum, gap that one can now numerically obtain using this method.

In this case they show that, generalizing the possible maps to include ones with indefinite time direction (a further generalization of possible causal orders), the maximum logical bound can always be attained even with classical systems. This result give some intuition on the power of different classes of generalized causal structures, drawing an interesting parallelism with the usual Local/Quantum/GPT correlations in Bell-like scenarios.

I found the presentation is very clear and both the paper and the SI, and I particularly appreciated the care that the authors put in giving intuitive explanations of the more technical parts in the method section.

This method constitute an important missing piece in quantum causal modeling.

In particular it allows for "observational" causal inference and the ability to falsify causal assumptions from observational data only, a fundamental task in causal inference, that was missing in the quantum case (which relied usually on intervention schemes).

This represents a first step to brings this framework closer to its classical counterpart.

In conclusion, I think the technique that the paper describes, represent a relevant addition to the field of quantum causal modeling, and can hopefully bring some more interesting development in this direction.

I would therefore recommend it for publication in Nature communications.

Additional comments:

- In the second to last paragraph of page 2, the authors say that the optimization over ICO correlation set is a non-convex problem.

But shortly before they defined this set as the one generated by the S matrices that "contains all possible processes with definite causal order, as well as other processes that are incompatible with any pre-defined causal order".

I would say that this is a convex set, despite likely not being directly solvable by expressing it in some efficient optimization method as one can do for other convex cones, that is, it is not a linear cone, a PSD cone, a second-order cone or similar.

Indeed, exactly because of the analogy with the Bell case, since we are not specifying the dimension of the system, we should end up having a convex set of correlations also for the ICO set.

If for some reason that I may have missed this is not the case, I think it should be explicitly pointed out why.

- In page 4 the authors present a suggestive representation of the different correlations sets, that resembles the usual one in a Bell scenario.

While this similarity (together with the corresponding representation) were already derived before by lower bounds to the set (for instance in ref. [30]), I am not aware if there is a simple explanation on why this correspondence exists.

Is there a simple mapping or an intuitive explanation for the similarity of this set to the Bell case? If so it would be nice to mention it.

Minor corrections an typos:

- At the end of the first column in page 3, it should be "...setting is equal _to_ the trigger..."

- In the second column of page 3 it should be "do not allow ~to~ any".

- At page 4, at the beginning of the section "Maximum quantum violation of the OCB inequality", the word "bit is repeated.

- In the supplementary note 3, at page 5 of the SI, $P^{\text{ave}}_{\text{succ}}$, is used but was never introduced.

I assume it represents the average success probability for the LGYNI game.

- At the beginning of supplementary note 8, at page 9 of the SI, the definition of η points generically to Appendix instead of Supplementary note 1.

Reviewer #2

(Remarks to the Author)

In their paper, Liu and Chiribella consider the problem of maximizing linear functionals over the set of multipartite quantum correlations with indefinite causal order. Namely, the set of probabilities that a number of separate parties can observe when we require causality and the laws of quantum mechanics to hold within each lab. For some class of such linear functionals or games, which they call single trigger, they find that they can compute the maximum value through a semidefinite program (SDP). The key to this technical achievement is the realization that all single trigger games essentially have the same optimal local strategy. This allows them to frame the optimization over quantum correlations with indefinite causal order as an optimization over process matrices, which only satisfy positivity and linear constraints.

Observing that any general game can be decomposed as linear combinations of single trigger games, the authors propose an SDP to bound the maximum value of arbitrary causal games. The new method allows them to exactly compute the maximum quantum probability of success of several causal games, such as the original Oreshko-Costa-Brukner (OCB) game.

In parallel, the authors show that the maximization of game scores becomes trivial in another type of multipartite causal setting: scenarios with an indefinite time direction. In these scenarios, the operations conducted by each party must be compatible with the two directions of time flow. Such a restriction on the set of feasible operations has the effect of enlarging the dual set, i.e., the corresponding set of classical processes. As the authors show, the dual becomes so large in fact, that the maximization of any linear functional of the observed probabilities always achieves the algebraic maximum.

This paper therefore consists of two very different (and one might add, independent) results. I am not thrilled about the second one, because I cannot picture what kind of strange world would preclude an experimentalist from, e.g., erasing a bit. With regards to the authors' SDP relaxation for causal inequalities, I am quite sure that it will be well received by the quantum foundations community. Indeed, although causal inequalities had been known for more than ten years, so far there

were no rigorous quantum bounds. In this sense, the authors are right when they claim that their analytic results are the analogs of the first Tsirelson bounds for this new causal setup.

In another sense, they are wrong. It makes sense to derive (the standard) Tsirelson bounds because our current model of quantum nonlocality exactly characterizes the set of measurement statistics available to several distant observers living in a quantum world. Indeed, suppose that we are given the mathematical descriptions of the parties' joint quantum state and their local measurements. Assuming certain basic experimental primitives (say, local qubit unitaries, Pauli measurements and the CNOT gate), it is then straightforward to devise a quantum experiment that prepares and distributes said state and allows each party to conduct said measurements.

In contrast, given the mathematical description of an acausal quantum process, in general we would not know how to (or whether one can) implement it in the lab, even in theory. The merit of the OCB paper was to point out that such hypothetical acausal processes are compatible with the laws of quantum mechanics; of course, that does not mean that such processes actually exist or can be created. In fact, 12 years after the OCB paper it is still unclear if, given access to any number of perfect quantum computers and error-free quantum channels, one could experimentally violate a single causal inequality. Thus, this paper is trying to bound the maximum expression of a physical effect that is not known to exist in theory.

To sum up, the first result of the paper is very good, but can only be understood and appreciated by the segment of the quantum foundations community that works on acausal processes. This is not such a large bunch, so I do not think that this paper fits Nature Communications' acceptance criteria. Perhaps a more specialized journal, like Quantum, would be more appropriate.

Reviewer #3

(Remarks to the Author)

Review of „Tsirelson bounds for quantum correlations with indefinite causal order“ by Zixuan Liu and Giulio Chiribella

In this paper the authors answer the important question of the maximal possible violations of causal inequalities within quantum theory with indefinite causal order. They manage to do so by cleverly reformulating this originally non-convex optimization problem as an SDP using so called single trigger correlations and what they call a canonical strategy. This work nicely connects previous partial results and findings (Refs. [30] [32] [37] and [39]) and provides a general procedure for calculating non trivial upper bounds for any causal inequality.

While this main result is very interesting and highly relevant for the field of indefinite causal order, I was somewhat puzzled and less convinced by the part of the paper on indefinite time direction (ICOTD). To begin with, I do not understand the (physical) motivation for considering such scenarios. Moreover, this framework seems to amount to combining processes with their time-reversed version. If that's the case, I would be highly surprised if one could not violate causal inequalities to the algebraic maximum. Causal inequalities can be thought of as expressing the limitations on how the parties involved can signal to one another. A scenario where the parties can all perfectly signal to each other should therefore give the algebraic maximum for the respective correlation function. Whenever a process allows for signaling from A to B, the time reversed scenario will allow signaling from B to A and, hence, the combination of the two will establish two-way signaling. I might be missing a crucial point here, but this is what I understood from the paper.

The work appears technically sound and the details provided in the supplementary information are sufficient for reproducing the results. Personally I would have preferred some of the technicalities to be explained further in the main text, but I understand that this does probably does not fit the format of Nature communications.

All in all, I'm in favor of publishing the article provided the authors either make a stronger case for their results on ICOTD or decrease their prominence in the text.

Version 1:

Reviewer comments:

Reviewer #3

(Remarks to the Author)

In this revised version the authors have satisfactorily addressed the points I raised in my first review and in my opinion improved their already very good paper considerably. The new section on reaching the algebraic maximum is much better connected to rest of the paper and the analogy with Bell inequalities and boxworld serves as a nice motivation. I strongly recommend publication in Nature Communications.

On side note, in my original review I was indeed referring to perfect signaling between the parties, since perfect one way signaling can be achieved by (even causal) processes the combining a process with its time-reversed version leads to perfect two way signaling. However, I appreciate the authors' comments about the fact that not the whole communication protocol is reversed for ICODT.

Open Access This Peer Review File is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

In cases where reviewers are anonymous, credit should be given to 'Anonymous Referee' and the source.

The images or other third party material in this Peer Review File are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

To view a copy of this license, visit <https://creativecommons.org/licenses/by/4.0/>

Response to Reviewer 1

In this paper the authors provide a method to compute the upper bound of causal inequalities arising from quantum causal models.

The quantum causal modeling framework, differently to its classical counterpart, allows a much more general class of causal structures, that cannot be reduced to probabilistic mixtures of classical ones.

For this reason this kind of structures are said to represent a set operations having indefinite causal order (ICO).

Conversely causal inequalities are a set of constraints implied by imposing definite causal order.

The interest in this kind of constraints derive from the fact that, similarly to what happens with the device-independent certification of entanglement in Bell-like scenario, a violation of such inequality, would reveal the presence of ICO only by looking at the output statistics.

Possible violations of causal inequalities have been identified, but there it is still not clear whether they can represent physical processes.

The question answered by the authors is how to find the maximal value, or more generally an upper bound, for a given causal inequality for quantum causal structures. While any explicit realization always represent a valid lower bound, the question of finding an upper bound to the set for a general scenario is an interesting one that has remained open until now.

Response: We thank the Referee for their accurate reading of our manuscript and for effectively summarizing the main question addressed in our work. We are glad that the Referee finds this question interesting, and are grateful for the positive assessment stated in the remaining part of their report.

To obtain this upper bound the authors first solve the problem for a specific class of inequalities, single-trigger inequalities, that they show can be maximized by a single SDP.

Then by using this result, they show that one can obtain an upper bound for general inequalities, by decomposing them in terms of single-trigger ones.

As a consequence this upper bound is also computable by single (albeit larger) SDP, which represent a relaxation of the original problem.

Using this method they obtain several results for known causal inequalities based on variations of "guess your neighbour input" games.

In particular I found very neat how the analysis of from the dual SDP can obtain analytical bounds and reconstruct the geometry of the definite and indefinite causal order correlations sets in the OCB inequality case, and the LGYNI inequality.

Moreover they also explore the question of the gap between the ICO and the logical maximum, gap that one can now numerically obtain using this method.

In this case they show that, generalizing the possible maps to include ones with indefinite time direction (a further generalization of possible causal orders), the maximum logical bound can always be attained even with classical systems.

This result give some intuition on the power of different classes of generalized causal structures, drawing an interesting parallelism with the usual Local/Quantum/GPT correlations in Bell-like scenarios.

Response: Thank you for the accurate summary of our technical contributions, and for appreciating the results on the OCB and LGNYI games, as well as the results on alternative theories that achieve the algebraic maximum. In the revised version, we provided some clarifications on this last point and included more details in the Methods section.

I found the presentation is very clear and both the paper and the SI, and I particularly appreciated the care that the authors put in giving intuitive explanations of the more technical parts in the method section.

This method constitute an important missing piece in quantum causal modeling. In particular it allows for "observational" causal inference and the ability to falsify causal assumptions from observational data only, a fundamental task in causal inference, that was missing in the quantum case (which relied usually on intervention schemes). This represents a first step to brings this framework closer to its classical counterpart.

In conclusion, I think the technique that the paper describes, represent a relevant addition to the field of quantum causal modeling, and can hopefully bring some more interesting development in this direction.

I would therefore recommend it for publication in Nature communications.

Response: We thank the Referee for their positive assessment of our work, and for the recommendation toward publication in Nature Communications. We are delighted that the Referee found the presentation insightful and appreciated the fundamental relevance of our findings.

Additional comments:

- In the second to last paragraph of page 2, the authors say that the optimization over ICO correlation set is a non-convex problem.

But shortly before they defined this set as the one generated by the S matrices that "contains all possible processes with definite causal order, as well as other processes that are incompatible with any pre-defined causal order".

I would say that this is a convex set, despite likely not being directly solvable by

expressing it in some efficient optimization method as one can do for other convex cones, that is, it is not a linear cone, a PSD cone, a second-order cone or similar.

Indeed, exactly because of the analogy with the Bell case, since we are not specifying the dimension of the system, we should end up having a convex set of correlations also for the ICO set.

If for some reason that I may have missed this is not the case, I think it should be explicitly pointed out why.

Response: Apologies for the confusion, you are correct. The actual problem, as you mentioned, is that the set of ICO correlations is not one of the convex sets that can be treated efficiently by standard convex optimization techniques (basically, due to the presence of an optimization over quantum systems of arbitrarily large dimension.) In the revised version, we amended the statement and removed the expression “non-convex optimization problem.”

- In page 4 the authors present a suggestive representation of the different correlations sets, that resembles the usual one in a Bell scenario.

While this similarity (together with the corresponding representation) were already derived before by lower bounds to the set (for instance in ref. [30]), I am not aware if there is a simple explanation on why this correspondence exists.

Is there a simple mapping or an intuitive explanation for the similarity of this set to the Bell case? If so it would be nice to mention it.

Response: These are indeed intriguing questions. Unfortunately we are not aware of a direct connection between the CHSH game and the OCB game, and in this stage we don't have further insights to offer on this point. Finding a deeper connection between these two games remains as an interesting direction for future explorations.

Minor corrections an typos:

- At the end of the first column in page 3, it should be "...setting is equal _to_ the trigger..."

- In the second column of page 3 it should be "do not allow ~to~ any".

- At page 4, at the beginning of the section "Maximum quantum violation of the OCB inequality", the word "bit" is repeated.

Response: Thank you for spotting these typos, they have been corrected in the revised version.

- In the supplementary note 3, at page 5 of the SI, $P^{\text{ave}}_{\text{succ}}$, is used but was never introduced.

I assume it represents the average success probability for the LGYNI game.

Response: Correct. In the revised version we specified that $P^{\text{ave}}_{\text{succ}}$, is the average probability of success in the LGYNI game.

- At the beginning of supplementary note 8, at page 9 of the SI, the definition of η points generically to Appendix instead of Supplementary note 1.

Response: Thanks for spotting this issue. In the revised version we updated the reference to Supplementary Note 1.

Response to Reviewer 2

In their paper, Liu and Chiribella consider the problem of maximizing linear functionals over the set of multipartite quantum correlations with indefinite causal order. Namely, the set of probabilities that a number of separate parties can observe when we require causality and the laws of quantum mechanics to hold within each lab. For some class of such linear functionals or games, which they call single trigger, they find that they can compute the maximum value through a semidefinite program (SDP). The key to this technical achievement is the realization that all single trigger games essentially have the same optimal local strategy. This allows them to frame the optimization over quantum correlations with indefinite causal order as an optimization over process matrices, which only satisfy positivity and linear constraints.

Observing that any general game can be decomposed as linear combinations of single trigger games, the authors propose an SDP to bound the maximum value of arbitrary causal games. The new method allows them to exactly compute the maximum quantum probability of success of several causal games, such as the original Oreshko-Costa-Brukner (OCB) game.

In parallel, the authors show that the maximization of game scores becomes trivial in another type of multipartite causal setting: scenarios with an indefinite time direction. In these scenarios, the operations conducted by each party must be compatible with the two directions of time flow. Such a restriction on the set of feasible operations has the effect of enlarging the dual set, i.e., the corresponding set of classical processes. As the authors show, the dual becomes so large in fact, that the maximization of any linear functional of the observed probabilities always achieves the algebraic maximum.

Response: We thank the Referee for their careful reading of the manuscript and for the accurate summary of our main results and their technical aspects.

This paper therefore consists of two very different (and one might add, independent) results. I am not thrilled about the second one, because I cannot picture what kind of strange world would preclude an experimentalist from, e.g., erasing a bit.

Response: We agree that the paper contains two different results, and appreciate your concerns about the physical interpretation of the second one. We agree that a “bistochastic world” would be a strange world. On the other hand, here we are not arguing for the empirical validity of bistochastic quantum theory, but rather use it as a toy theory to provide an example in which certain causal inequalities can be violated to their algebraic maximum. In the revised version, we clarified this point, removing the comments on time direction, and focussing on the achievability of the algebraic maximum.

With regards to the authors' SDP relaxation for causal inequalities, I am quite sure that it will be well received by the quantum foundations community. Indeed, although causal inequalities had been known for more than ten years, so far there were no rigorous quantum bounds. In this sense, the authors are right when they claim that their analytic results are the analogs of the first Tsirelson bounds for this new causal setup.

Response: Thank you very much for your kind appreciation of the relevance of our result to the quantum foundations community.

In another sense, they are wrong. It makes sense to derive (the standard) Tsirelson bounds because our current model of quantum nonlocality exactly characterizes the set of measurement statistics available to several distant observers living in a quantum world. Indeed, suppose that we are given the mathematical descriptions of the parties' joint quantum state and their local measurements. Assuming certain basic experimental primitives (say, local qubit unitaries, Pauli measurements and the CNOT gate), it is then straightforward to devise a quantum experiment that prepares and distributes said state and allows each party to conduct said measurements. In contrast, given the mathematical description of an acausal quantum process, in general we would not know how to (or whether one can) implement it in the lab, even in theory.

Response: We agree that the physical realizability of arbitrary acausal processes is an open problem. Nevertheless, finding the maximum violation of causal inequalities over all acausal process is important because it provides the starting point for addressing the question of why quantum theory is logically compatible with some amount of acausality, but not with the maximum amount. The in-principle, logical compatibility with a given amount of acausality is a property of quantum theory, independently of whether acausal processes can be implemented in practice.

The merit of the OCB paper was to point out that such hypothetical acausal processes are compatible with the laws of quantum mechanics; of course, that does not mean that such processes actually exist or can be created. In fact, 12 years after the OCB paper it is still unclear if, given access to any number of perfect quantum computers and error-free quantum channels, one could experimentally violate a single causal inequality.

Response: This comment seems to be in contradiction with the recent work of Wechs, Oreshkov, and Branciard [*Nature Communications* **14**, 1471 (2023)], who showed examples of causal inequalities that can be violated with time-delocalized quantum systems. Since the processes achievable with time-delocalized systems are a subset of the set of all quantum ICO processes, our upper bounds apply to them as well. This is useful because finding upper bounds on the violation of causal inequalities is generally a difficult problem, also for the subclass of processes that can be realized with time-delocalized systems. In the revised version of the paper we included a comment on this point.

Thus, this paper is trying to bound the maximum expression of a physical effect that is not known to exist in theory. To sum up, the first result of the paper is very good, but can only be understood and appreciated by the segment of the quantum foundations community that works on acausal processes. This is not such a large bunch, so I do not think that this paper fits Nature Communications' acceptance criteria. Perhaps a more specialized journal, like Quantum, would be more appropriate.

Response: We appreciate your consideration of this matter, but we respectfully disagree with the statement that our results can only be appreciated by the experts of acausal processes. Identifying the maximal violations of causality that are logically compatible with quantum mechanics is the starting point for raising the important question of why quantum theory is in principle compatible with some amount of causal indefiniteness but not with the maximum possible amount. This problem can potentially lead to the discovery of new physical or information-theoretic principles, and, eventually, to a new axiomatization of quantum theory that does not require a pre-defined causal structure. These topics are relevant to the foundations of quantum mechanics, and can appeal a larger audience than just the researchers working on acausal processes.

In fact, it appears that our work has already attracted interest even beyond the quantum foundations community, and in the past months has been the object of both invited and contributed talks at quantum information conference, such as Asian Quantum Information Science (AQIS2024), Central European Quantum Information Workshop (CEQIP2024), and the Seefeld Workshop on Quantum Information Theory 2024. Overall, the interactions we had in these and other venues indicate that the results of this paper can pique the interest of a fairly broad audience of physicists with an interest in causality and quantum mechanics.

On a final note, we feel that research on causality in quantum mechanics can offer valuable insights to the broader physics community, and that it is therefore important to share our finding our results not only with the experts, but also with other fellow physicists who may be interested in the emerging field of quantum causality.

Response to Reviewer 3

In this paper the authors answer the important question of the maximal possible violations of causal inequalities within quantum theory with indefinite causal order. They manage to do so by cleverly reformulating this originally non-convex optimization problem as an SDP using so called single trigger correlations and what they call a canonical strategy. This work nicely connects previous partial results and findings (Refs. [30] [32] [37] and [39]) and provides a general procedure for calculating non trivial upper bounds for any causal inequality.

Response: We thank the Referee their positive assessment of the importance of the question addressed in our work, and we are pleased that the Referee regards our main result as “*very interesting and highly relevant for the field of indefinite causal order.*” We also appreciate the Referee’s thoughtful comments on the additional results on indefinite time direction, which have stimulated a change in the presentation in the revised version.

While this main result is very interesting and highly relevant for the field of indefinite causal order, I was somewhat puzzled and less convinced by the part of the paper on indefinite time direction (ICOTD). To begin with, I do not understand the (physical) motivation for considering such scenarios.

Response: We appreciate that the notion of quantum operations with indefinite time direction is relatively new and may cause some puzzlement. In the revised version we followed your advice to reduce the emphasis on indefinite time direction, and we modified the presentation, clarifying the relation of the related results with the main thread of the paper. More discussion on this points is provided in the replies to your next comments.

Moreover, this framework seems to amount to combining processes with their time-reversed version. If that’s the case, I would be highly surprised if one could not violate causal inequalities to the algebraic maximum. Causal inequalities can be thought of as expressing the limitations on how the parties involved can signal to one another. A scenario where the parties can all perfectly signal to each other should therefore give the algebraic maximum for the respective correlation function. Whenever a process allows for signaling form A to B, the time reversed scenario will allow signaling form B to A and, hence, the combination of the two will establish two-way signaling. I might be missing a crucial point here, but this is what I understood from the paper.

Here there seems to be a bit of confusion. Processes allowing for signalling both from A to B and from B to A are already allowed by the standard framework for indefinite causal order, irrespectively of time reversal. In fact, only processes that signal simultaneously from A to B and from B to A can violate causal inequalities.

The existence of *some* two-way signalling, however, does not imply that this signalling is *perfect*. Now, what we show is that perfect two-way signalling becomes possible in an alternative theory where Alice's and Bob's experiments are described by bistochastic instruments. As you correctly mentioned, processes in this alternative theory can be interpreted as combining Alice's and Bob's local experiments with their time reversed versions. However, this interpretation alone does not *a priori* imply the possibility of perfect two-way signalling: it is not the whole communication protocol between Alice and Bob that is time-reversed, but only the local operations taking place in their local laboratories.

It is also worth noting that the possibility of perfect two-way signalling in bistochastic quantum theory is a rather subtle matter. In the revised version, we clarified that, in fact, such perfect signalling can be achieved when A and B have up to three measurement settings, while for larger number of settings it is an open problem. The complexity of this problem further indicates that perfect two-way signalling is not a trivial consequence of the framework.

This said, we appreciate that mentioning indefinite time direction could cause some misunderstandings. In the revised version, we removed the mention to this notion. The main emphasis is now on the fact that we provide alternative theories where the set of logically admissible processes permits the maximum algebraic violation of certain causal inequalities.

The work appears technically sound and the details provided in the supplementary information are sufficient for reproducing the results. Personally I would have preferred some of the technicalities to be explained further in the main text, but I understand that this does probably does not fit the format of Nature communications.

Response: We understand your thinking, and have been spending quite some time considering how to distribute the technical materials between the main text and supplementary notes. For the Tsirelson-like bounds, the current distribution seemed to be the optimal compromise, and we did not make substantial changes in this revision. For the algebraic maximum result, instead, we realized that it would be useful to give more details, which have now been added in the revised version.

All in all, I'm in favor of publishing the article provided the authors either make a stronger case for their results on ICOTD or decrease their prominence in the text.

Response: Thank you for your appreciation of our work and for the in-principle positive recommendation towards publication in Nature Communications, subject to amendment of the presentation of the ICOTD results. In the revised version, we followed the second option you indicated, reducing the prominence of the results and completely removing the reference to indefinite time direction and the term ICOTD in the main text. Instead, we presented the related results as examples of alternative

theories that violate certain causal inequalities to their maximum algebraic value, in a similar way as the alternative theory of PR boxes leads to a violation of the CHSH inequality to the algebraic maximum.