Quantum paradox of choice: More freedom makes summoning a quantum state harder

Emily Adlam[1] and Adrian Kent[1,2]

[1]*Centre for Quantum Information and Foundations, DAMTP, Centre for Mathematical Sciences, University of Cambridge, Wilberforce Road, Cambridge CB3 0WA, United Kingdom*

[2]*Perimeter Institute for Theoretical Physics, 31 Caroline Street North, Waterloo, Ontario, Canada N2L 2Y5*

The properties of quantum information in space-time can be investigated by studying operational tasks, such as "summoning," in which an unknown quantum state is supplied at one point and a call is made at another for it to be returned at a third. Hayden and May [arXiv:1210.0913] recently proved necessary and sufficient conditions for guaranteeing successful return of a summoned state for finite sets of call and return points when there is a guarantee of at most one summons. We prove necessary and sufficient conditions when there may be several possible summonses and complying with any one constitutes success, and we demonstrate the existence of an apparent paradox: The extra freedom makes it strictly harder to complete the summoning task. This result has practical applications for distributed quantum computing and cryptography and implications for our understanding of relativistic quantum information and its localization in space-time.

A holistic magician (HM) repeatedly performs the following trick. He first asks you to give him an object that you are sure he cannot copy. After working behind a curtain, he presents you with N boxes and asks you to choose one. Opening your chosen box, he reveals the original object inside.

You initially imagine that he has arranged some concealed mechanism that somehow passes the object sequentially through the boxes, allowing him to stop the mechanism and keep the object in one box if you select it. However, you are then puzzled to notice that he is unable to make the trick work if you select more than one box, even though you allow him to choose which of your selections to open. This argues against your mechanical explanation and indeed seems to make any mechanistic explanation problematic. How can giving the magician more freedom make him unable to complete the task?

Strange as it sounds, when quantum mechanics is combined with special relativity it is possible to come up with tasks that have this paradoxical property. There is a well-known "paradox of choice" in economics: More choice can make consumers less happy [1]. By analogy, we describe the phenomenon illustrated by our tasks as a "quantum paradox of choice." Our quantum paradox, however, is much sharper: More freedom in choosing how to complete the task can make it *impossible*. This result has significant consequences for the design of distributed quantum computing routines and quantum cryptographic protocols in contexts where relativistic signaling constraints are important, for example on global financial networks. It also has interesting theoretical implications, for although there is a long-standing tradition of using apparent paradoxes to refine our understanding of quantum theory [2–7], the effect we describe here is a genuinely *relativistic* quantum paradox.

I. SUMMONING TASKS

The study of information-theoretic games and other operationally defined tasks has been an important tool of quantum information theory ever since the birth of the field [8–12], but relativistic quantum tasks are still comparatively unexplored.

Our paradox involves a task known as summoning, in which an agent is given an unknown quantum state and required to produce it at a point in space-time in response to a call made at some earlier point [13]. We generalize this task to allow calls to be made at any number of call points, requiring the agent to return the state at any return point corresponding to one of the calls. We establish necessary and sufficient conditions on the geometric configuration of call and return points in space-time to guarantee that the task can be completed. Further, we show that these are strictly stronger conditions than those established by Hayden and May (HM) for the case where it is guaranteed that at most one call will be made.

Summoning was first introduced [13] as a simple illustration of a task that distinguishes relativistic quantum theory from both relativistic classical theories and nonrelativistic quantum theory. In its original form, the task involves two agencies, A and B, each comprising collaborating networks of agents distributed throughout space-time. B secretly prepares a random quantum state in some agreed Hilbert space—for example, a random qubit—and gives it to A at some point P. At some later point Q in the causal future of P, B will ask for the state back. The point Q is not known in advance by A. Making use of the no-signaling principle from special relativity (see, e.g., [14]) and the no-cloning theorem from quantum mechanics [15,16], it is possible to derive a "no-summoning theorem" [13] which states that no matter how densely Alice's agents are distributed, in general there will be no strategy which guarantees a successful response to Bob's request.[1] This remains true for variations of the task in which time delays in returning and some loss of fidelity in the returned state are allowed [13].

An example of summoning with separated call and return points was given in [18], where it was shown that there exist summoning tasks that cannot be completed by simply propagating the unknown state along a fixed path but can be completed by the use of quantum teleportation. Hayden and May [19] introduced a generalized version of this summoning task defined by a space-time point P_s and a set of N ordered

---

[1]A no-summoning theorem for unknown mixed states follows similarly from the no-broadcasting theorem [17].

pairs of space-time points $\{c_i, r_i\}$. Alice is given a quantum system in some unknown state ψ at a start point s. If a call is made at point c_i, then Alice must return a quantum system in state ψ at the corresponding response point r_i. We follow HM in stipulating that each r_i must be in the causal future of c_i.[2]

For definiteness we work in Minkowski space, although our results and most of our comments apply to more general space-times. Write $x > y$ if the space-time point x is in the causal future of y and $x \geqslant y$ if either $x > y$ or $x = y$. HM define the *causal diamond* D_i to be the set $\{p : r_i \geqslant p \geqslant c_i\}$. They then use iterative applications of quantum teleportation and secret sharing to prove the following beautiful result.

Theorem 1 [19]. Under the assumptions described, summoning is possible if and only if every reveal point $r_i \geqslant s$ and every pair of causal diamonds D_i and D_j are causally related, meaning that there exists $x_i \in D_i$ and $x_j \in D_j$ with $x_i \geqslant x_j$, or vice versa.

These are considerably weaker conditions than naive intuition might suggest. In particular, due to the delocalizability of quantum information, there need not necessarily exist a causal path that starts from s and runs sequentially through the causal diamonds. Hence, it is sometimes possible to complete a summoning task even if it is not possible to complete a comparable relativistic task for a classical object that cannot be duplicated rather than a quantum state. This underlines that summoning is both intrinsically quantum and intrinsically relativistic, since its necessary and sufficient conditions depend crucially on both quantum and relativistic laws.

Summoning is a natural task for distributed quantum computing over networks where relativistic signaling constraints are significant. It is natural for these applications, and also intrinsically theoretically interesting, to introduce another version of summoning, in which calls may arrive at any number of the call points. Since the r_i may be spacelike separated, and since in any case we assume that the unknown state is handed over to another agency when returned, the no-cloning theorem means that Alice cannot return the state several times. We hence define the task such that, if several calls are made at points c_i ($i \in I$), Alice need only return the state at any one of the corresponding return points r_i ($i \in I$).

We now give a characterization of the necessary and sufficient conditions for the existence of a protocol which guarantees a successful response to such a summoning task and show that these are strictly stronger conditions than those established by HM for the case where it is guaranteed that at most one call will be made.

II. NECESSARY AND SUFFICIENT CONDITIONS

We adopt the approximation that quantum states may be effectively localized to a point; see Ref. [19] for further discussion of this approximation and its limitations and follow HM in stipulating that each r_i must be in the causal future of c_i.[3] We then have the following result, with notation defined as before.

Theorem 2. Consider a summoning task defined by a space-time point P_s and a set of N ordered pairs of space-time points $\{c_i, r_i\}$, in which Alice is given a quantum system in some unknown state ψ, and if a call is subsequently made at some set of points K, Alice must return a quantum system in state ψ at any response point $r_i : i \in K$. There exists a protocol which guarantees success for such a task if and only if

(1) every response point $r_i \geqslant s$;

(2) for any subset K of $\{1, \dots, N\}$, there is at least one $k \in K$ such that $r_k \geqslant c_i$ for all $i \in K$.

Proof.

Necessity. The necessity of the first condition is implied by the no-signaling principle.

The necessity of the second condition is established using a proof by contradiction. In order to obtain a contradiction, we make the following assumption.

Assumption 1. There exists a successful protocol for a subtask defined from the original task by selecting some set $\mathcal{M}$ of $M \leqslant N$ call-response pairs, such that for every call-response pair $(c_i, r_i) \in \mathcal{M}$, there exists at least one call-response pair $(c_j, r_j) \in \mathcal{M}$ such that $r_i \not> c_j$.

If c_j is the call point such that $r_i \not> c_j$, it follows from the no-signaling principle that the response made at r_i when calls are made at some set of call points Q with $c_i, c_j \in Q$ is the same as the response made at r_i when calls are made at the set $\{Q \setminus c_j\}$.

Moreover, since each return point is in the causal future of the corresponding call point, Alice may decline to return anything at r_i if no call is made at c_i. We may therefore assume that Alice's strategy returns no state anywhere if no calls are made.

Thus, Assumption 1 implies that for any r_i, if $\mathcal{Q}_i$ is the total number of subsets $Q \subseteq \{c_1, c_2, \dots, c_M\}$ such that ψ is returned at r_i if calls are made at all points in Q, then $\mathcal{Q}_i$ must be even. Therefore, $\sum_{i=1}^{M} \mathcal{Q}_i$ must be even.

However, in order to complete the task successfully, Alice must respond at exactly one point whenever calls are made at any subset $Q \subseteq \{c_1, c_2, \dots, c_M\}$, so we must have $\sum_{i=1}^{M} \mathcal{Q}_i = \sum_{j=1}^{M} C_j^M = 2^M - 1$, which is always odd.

Thus, we have derived a contradiction from Assumption 1, so this assumption can never be true.

Therefore there exists a successful protocol for the subtask defined by a set $\mathcal{M}$ of $M \leqslant N$ call-response pairs only if there exists at least one call-response pair $(c_i, r_i) \in \mathcal{M}$ such that r_i is in the future light cone of the call point for every pair in $\mathcal{M}$.

Since this reasoning may be applied to any subtask, there exists a successful protocol for a summoning task without guarantee with N call-response pairs only if for any subset S of call-response pairs at least one response point r_k belonging to a pair in S lies in the future light cone of all the other call points belonging to pairs in S.

Sufficiency. We now show that the conditions are sufficient by exhibiting a protocol that always succeeds if the call-response pairs satisfy the conditions of Theorem 2. Define $S_N = \{1, \dots, N\}$. From the second condition of Theorem 2, there is at least one $i \in S_N$ such that $r_i \geqslant c_j$ for all $j \in S_N$. Choose one such, i_N, and define $S_{N-1} = S_N \setminus \{i\}$. Similarly, choose $i_{N-1} \in S_{N-1}$ such that $r_{i_{N-1}} \geqslant c_j$ for all $j \in S_{N-1}$, and so on. We thus obtain an ordered sequence of causal diamonds

---

[2]For further discussion of this point, see Appendix A.

[3]For further discussion of this point, see Appendix A.

$D_{i_1}, \ldots, D_{i_N}$ such that the return point of any diamond in the sequence lies in the causal future of the call points of all previous diamonds. We relabel the $D_{i_j}, c_{i_j}, r_{i_j}$, writing j for i_j.

Alice may now proceed as follows. Let d be the agreed dimension of the Hilbert space of the unknown state ψ. Before the protocol begins, she distributes maximally entangled pairs of states in $\mathcal{C}^d \otimes \mathcal{C}^d$ between agents at the spatial locations of $(s, c_1), (c_1, c_2), \ldots, (c_{N-1}, c_N)$. When the state is given to her at s, her agent there teleports it to c_1, broadcasting the classical teleportation data. If a call arrives at c_1, the agent at c_1 transmits the quantum part of the teleported state (i.e., her half of the entangled pair) to r_1, where another agent uses the classical teleportation data to reconstruct the state and return it to Bob.

If a call does not arrive at c_1, Alice's agent there teleports the state to c_2, broadcasting the classical teleportation data. If a call arrives at c_2, the quantum part of the teleported state is sent to r_2, where the classical teleportation data from s and c_1 are used to reconstruct and return the state. Otherwise, the process continues, until either a call arrives at some c_i (and the state is reconstructed and returned at r_i) or there is no call. In the latter case, Alice may reconstruct the state at r_N if she wishes, but does not return it to Bob. ∎

III. COMMENTS

(1) Any set of causal diamonds satisfying the conditions of Theorem 2 also satisfies those of Theorem 1. However, it is easy to construct examples of sets of causal diamonds that satisfy those of Theorem 1 but not those of Theorem 2. For example, Fig. 1 describes one such set. Allowing the possibility of more than one call thus makes the summoning task strictly harder.

(2) Nonetheless, the conditions of Theorem 2 still do not imply that there is a causal path running from the start point through each causal diamond. An example is given in Fig. 2.

(3) The conditions of Theorem 2 imply there is an ordered sequence of causal diamonds $D_{i_1}, \ldots, D_{i_N}$ such that the return point of any diamond in the sequence lies in the causal future
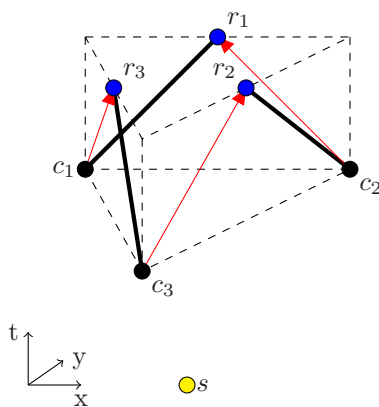

FIG. 2. An example with $c_1, c_2, c_3 < r_1, r_2, r_3$. The diamonds satisfy the conditions of Theorem 2, and so summoning with any number of calls is possible, even though there is no causal path running through all three diamonds. Green lines denote the timelike connection between a call point and its own associated response point, while red lines denote the lightlike connection between a call point and a response point from a different call-response pair.
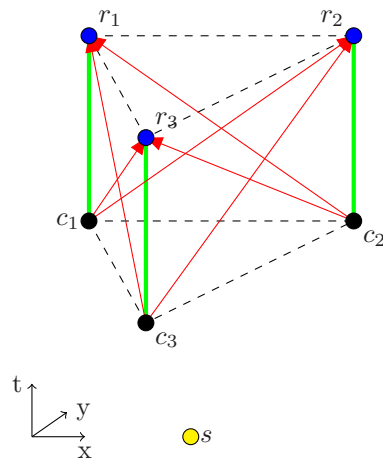
of the call points of all previous diamonds. This ordering is not necessarily unique. For example, a nested pair of diamonds, with $c_i \leqslant c_j < r_j \leqslant r_i$ with appropriate relations to the other diamonds may be taken in either order. More generally, one can construct examples including sets of n nonoverlapping diamonds (c_i, r_i) (for $1 \leqslant i \leqslant n$) for which $c_i < r_j$ for all i, j; Fig. 2 gives an example of this type for $n = 3$.

IV. RESOLUTION

The resolution of the apparent paradox rests on a previously unappreciated feature of summoning tasks. *Prima facie* it seems that the guarantee of at most one call plays no special role in a summoning task other than to ensure that Alice is never required to produce two copies of an unknown state, in violation of the no-cloning theorem. It thus initially seems paradoxical that summoning becomes strictly harder if we allow the possibility of more than one call, even though only one valid response is required.

However, the paradox is dispelled when we recognize that the guarantee of at most one call provides a *resource* that Alice can use to coordinate the actions of her spatially distributed agents. If Alice knows that no more than one call will occur, learning that a call has been made at one point tells her that there are no calls at any other point, and this allows her to coordinate the behavior of her agents via the global call distribution. On the other hand, if multiple calls can occur, learning that a call has been made at one point tells Alice nothing about the distribution of calls at other points. Thus, in a summoning task with unrestricted calls, Alice may have more freedom, but she also has less information, and this last turns out to be crucial.

FIG. 1. A $2 + 1$-dimensional example, taken from Ref. [19], where summoning is possible if one call is guaranteed, but not if more than one call may arrive. Black lines denote the lightlike connection between a call point and its own associated response point, while red lines denote the lightlike connection between a call point and a response point from a different call-response pair.

V. APPLICATIONS

We have given necessary and sufficient conditions for the existence of a protocol that guarantees a successful response to a task that may call for a state at several distinct nodes but requires the state to be produced at only one associated response node. This is a natural condition in many practical contexts, such as a distributed parallel quantum computation in which the output of a subprotocol is to be routed to any one of several parallel computations that call for the output when they reach a certain state. Indeed, in the teleportation model of distributed quantum computation [20], each round of adaptive computation is essentially a summoning task: The measurement result from the previous round determines the measurement to be made in the present round, and thus plays the role of the "call," while the locations of the gates for the various possible measurements play the role of the "response points."

The no-summoning theorem has also already led to new applications in relativistic quantum cryptography [21–25]. The stronger results reported here and in Ref. [19] suggest further ways of exploiting summoning as a general way of controlling the flow of quantum information. We thus expect these results to find application in future cryptographic protocols, as well as in quantum network algorithms.

VI. DISCUSSION

Like the original no-summoning theorem, our results rely crucially on both relativity and quantum theory. In particular, the distinction between summoning tasks with and without guarantees of a single call, and indeed that there are interesting constraints on summoning tasks at all, relies both on the impossibility of superluminal signaling and on the no-cloning theorem. Our results also exploit the delocalizability of quantum information via quantum teleportation, which allows summoning in configurations where sending the state along any given causal path fails.

Our results also clarify the role that the causal diamonds play in summoning tasks. Hayden and May originally described their result (Theorem 1) as follows: "We fully characterize which regions of spacetime can all hold the same quantum information. Because quantum information can be delocalized through quantum error correction and teleportation, it need not follow well-defined trajectories. Instead, replication of the information in any configuration of spacetime regions not leading to violations of causality or the no-cloning principle is allowed. . . . This provides a simple and complete description of where and when a qubit can be located in spacetime, revealing a remarkable variety of possibilities" [19]. However, as we explain in more detail in Appendix B, the idea that the same quantum information is replicated in each of the causal diamonds in a Hayden and May summoning task appears inconsistent with our Theorem 2: If the information were so replicated, it should be possible to guarantee successful completion of our generalized summoning task whenever the conditions of Theorem 1 hold, but as Theorem 2 shows, this is not the case.

We suggest that the role played by the causal diamonds in both Theorem 1 and Theorem 2 is better understood as

a succinct description of the information that is available to Alice in different regions. The future light cones of the call points are regions where Alice is in possession of nontrivial information about which call has been made; the past light cones of the return points are regions in which Alice is able to influence the outcome at those points. The causal diamonds are the intersections of these light cones. The difference between Theorems 1 and 2 can then be explained by the fact that the information that Alice learns at the call points is different for our generalized summoning task than for Hayden and May's original task. Thus, the difference between these two theorems, which seems "paradoxical" when the causal diamonds are described as regions of localization of quantum information, has a simple and straightforward explanation when the causal diamonds are described in terms of available information.

Our apparent paradox and its resolution therefore serve as an important reminder that statements about the spatiotemporal location of quantum states can be misleading. In Appendix C we discuss the ubiquity of such statements in physics and consider their broader implications for the field of quantum foundations. We recognize that such language can be a helpful way of proving and synthesizing operational results in cases where it accurately reflects their implications, but our results reinforce the fact that quantum states and quantum information are *not* persisting physical entities in the ordinary everyday sense, and describing them as such can give rise to fallacies [26].

APPENDIX A: SUMMONING WITHOUT RESTRICTIONS ON THE GEOMETRY

In the main text, we follow HM in assuming that each return point r_i lies in the causal future of its call point c_i. Actually, this assumption is slightly restrictive, since it can be possible to complete the task even when one of the r_i lies outside the causal future of its call point c_i. However, the theorem has a simple generalization that covers this case, as we show below.

A related issue is that HM stipulate that at most one call will be made but do not explicitly define the task requirements if no call is made. One natural definition is to require that Alice should return a quantum state at r_i *only if* she is certain at r_i that a call was made at c_i. Hence, she should not return any state anywhere if no call is made. With this refinement, it follows that each return point must indeed be in the future of the corresponding call point, as HM assume.

However, another natural definition of the task is to allow Alice to return a quantum state anywhere or nowhere, as she chooses, if no call is made. In one version of this task, any returned state must be the state originally supplied, even if no call is made. Another possibility is to allow Alice to return any state when no call is made. Even the first, stronger, version, is

possible for configurations where Theorem 1 does not apply, as we now show.

Recall again that the conditions of Theorem 1 assume that each r_i lies in the causal future of its call point c_i, so that each causal diamond D_i is well-defined and nonempty. The theorem then shows that summoning is possible only if for every pair i,j, the causal diamonds D_i and D_j are causally related.

Consider now the case where we have only two possible call points, c_1 and c_2, with corresponding return points r_1 and r_2. If summoning is possible, r_1 and r_2 must lie in the future of the start point s. However, summoning may be possible even if one of the causal diamonds D_1, D_2 is empty, i.e., even if one of the r_i does not lie in the future of c_i.

This is because in the case of only two call points, learning that no call is made at c_1 implies that either there is a call at c_2 or no call at all. Hence, if $r_2 > c_1$, we can construct a successful protocol even if it is not the case that $r_2 > c_2$.

(1) Before the protocol starts, Alice shares an entangled pair between agents who will be at s and c_1.

(2) When the state is handed over at s, Alice teleports it to c_1, broadcasting the classical data in all directions.

(3) If a call is received at c_1, Alice's agent at c_1 sends the quantum teleportation data (i.e., her part of the entangled pair) to r_1, where it is combined with the classical teleportation data to reconstruct the state, which is handed over to Bob. Alice's agent at c_1 broadcasts the fact that she has received a call. Alice's agent at r_2 then does not hand over anything.

(4) If no call is received at c_1, Alice's agent there sends the quantum teleportation data to r_2, where it is combined with the classical teleportation data and handed over to Bob.

Thus, if a call is made at c_1 the state is at r_1 at the end of the protocol; if a call is made at c_2 (or no call is made at all) the state is at r_2 at the end of the protocol.

Hayden and May's proof [19] shows that a protocol for any number of call points can be built recursively out of a protocol for two call points. Their argument extends to the more general configurations we consider. Theorem 1 can thus be extended as follows.

Theorem 3. A summoning task defined by a set of N ordered pairs $\{c_i, r_i\}$ is possible if and only if the following conditions hold.

(i) Every response point is in the future light cone of the start point.

(ii) For every pair $\{c_i, r_i\}, \{c_j, r_j\}$, both of the response points lie in the future light cone of a single call point.

Note that it follows that for every pair i, j at least one of the causal diamonds D_i, D_j must exist, so there can be at most one call-response point pair for which the causal diamond does not exist.

APPENDIX B: HAYDEN AND MAY ON THE SPATIOTEMPORAL LOCALIZATION OF QUANTUM STATES

There is a tradition in physics, thus far largely unanalyzed, of describing quantum states and quantum information using the language of persisting physical objects [26]. For example, Horodecki *et al.*'s widely cited review [27] states "This is the essence of teleportation: a quantum state is transferred from

one place to another: not copied to other place, but moved to that place." Hayden and May follow in this tradition in interpreting their result in the quotation discussed in the main text: "We fully characterize which regions of spacetime can all hold the same quantum information. Because quantum information can be delocalized through quantum error correction and teleportation, it need not follow well-defined trajectories. Instead, replication of the information in any configuration of spacetime regions not leading to violations of causality or the no-cloning principle is allowed. ... This provides a simple and complete description of where and when a qubit can be located in spacetime, revealing a remarkable variety of possibilities" [19].

We question whether HM's spatiotemporal account of the quantum state during summoning tasks can be supported. Indeed, even before considering the implications of Theorems 1 and 2, we find it a little tricky to understand what exactly Hayden and May's quotation is intended to mean.

One technical complication is that the causal diamonds in a summoning task may overlap or even be identical. In this case, even if two overlapping diamonds in some sense each hold a copy of the state, it could be a single copy held in their intersection. HM do not explicitly consider the possibility of overlaps. Since we wish to consider the case most favorable to their interpretation, we assume nonoverlapping diamonds here.

A more fundamental concern is that, if the quantum information were replicated in the usual sense within each causal diamond in a configuration satisfying HM's conditions, ordinary reasoning would suggest that, given a request at any number of call points, it should be possible to produce a copy of the state at each of the corresponding return points. This would violate the no-cloning theorem, of course, and evidently is not the intended meaning, hence HM's phrase *"replicated but in a restricted fashion."* We are clearly meant somehow to allow for the fact that replication in time is not the same as replication in space, but how, exactly?

HM's comment that "reversibility ... requires that [quantum] information be copied in time" suggests the following analogy. The quantum information in an unknown state with known unitary evolution can be captured and handed over at any given time, or more generally on any given spacelike hypersurface in a foliation. Once this happens, though, it cannot be repeated: If the state is requested again at later times (or on later hypersurfaces), Alice cannot provide a second copy. An analogous statement about summoning would be that, when the conditions of Theorem 1 hold, if several requests are made, the quantum information can be handed over in one causal diamond, but then this cannot be repeated in the other diamonds. However, as our result shows, this is not generally the case. Requesting the state at several HM call points *prevents* a successful response, *unless* the configuration of call and return points also satisfies the stronger conditions of our Theorem 2.

So what *can* we consistently say about the flow of quantum information in summoning tasks, in the light of our results? One safe option, of course, is to remain silent, retreating instead to a purely operationalist perspective. Proving results about possible and impossible quantum tasks in space-time [18] does not necessarily require any discussion of the localization or

location of quantum information in space-time, as Theorems 1 and 2 illustrate. Operationalism avoids the risk of fallacies that can arise from regarding (quantum) information as a physical substance [26].

That said, whatever view one takes on the reality of quantum states, speaking informally about the localization of quantum information can certainly be a very helpful way of proving and synthesizing operational results. Our suggestion is thus not that such language should necessarily be eliminated altogether, but rather that it needs to be analyzed critically and used more precisely. In particular, if the behavior of quantum information in a system depends on external events, such as the calls in a summoning task, we would argue that any discussion of its localization should reflect this explicitly. For example, a more precise informal statement of Hayden and May's result would be that *if* Alice follows their prescribed strategy, and *if* a call arrives at c_i and no call arrives at any other call point, *then* the unknown state's quantum information becomes localized within D_i and the state is reconstructed at r_i. This framing precludes speaking of "replication of the information in [the given] configuration of spacetime regions," since the possibilities of localizing the state within D_i and D_j are alternatives that depend on exclusive possible external events. We suggest too that any statements about the localization of quantum information in space-time—including those we give above for our version of the task—may be best seen as hypotheses that seem to usefully synthesize and summarize operational results, but might yet need to be refined or even rejected.

Finally, while we do not believe that our results imply the validity of any picture of spatiotemporarily localized quantum states, and we do not particularly recommend any such picture, we should note that one can describe the protocol used in our Theorem 2 quite simply in this language. Note first that the conditions of Theorem 2, unlike those of Theorem 1, always allow an *ordering* of the diamonds, albeit not necessarily a unique one. Moreover, they allow a strategy that effectively (although not literally) responds to the calls in the chosen ordered sequence. Hence, *if* Alice follows our strategy, then at each call point c_i she knows that, *if* the state is not being returned as the result of an earlier call, and if she chooses not to teleport her state to c_{i+1}, then—at least in the informal language physicists tend to use when discussing teleportation—the unknown state's quantum information *will* be contained within D_i. That is, there is a quantum state at c_i, together with classical information available at or before r_i, which could be brought together and recombined to reconstruct the state at r_i. In this informal sense, in our protocol, the quantum state is sequentially localized in D_1, D_2, and so on, until it reaches a diamond D_i with a call at c_i, which causes it to be returned at r_i.

APPENDIX C: FURTHER COMMENTS ON THE QUANTUM STATE AND SPATIOTEMPORAL REALITY

Our results directly apply only to specific summoning tasks; they do not give a definitive answer to the general question about when and whether it is correct (or at least consistent) to think of quantum states as though they were real physical objects that can move in space and persist over time (albeit

perhaps discontinuously, as the quotation from Ref. [27] above suggests is the case during teleportation). Still, they suggest to us that this and related questions deserve closer analysis and that examining relativistic quantum tasks [18] might yield further insights.

Our results also allow us to probe intuitions about the nature of quantum states as spatiotemporal entities, a question which has received comparatively little attention in recent debates over the reality of the quantum state [28–35]. To take one interesting example, the celebrated PBR theorem [29,30] purportedly addresses only the question of whether the quantum state is an element of reality in the instantaneous sense set out by Harrigan and Spekkens [28]. It seems to us that some caution is merited here, since the proof of the theorem depends implicitly on assumptions not only about states at a given time, but about the persistence of those states over time. Specifically, the argument hinges on the assumption of preparation independence, the idea that "systems that are prepared independently have independent physical states," and it is assumed without comment that the systems must therefore continue to have independent physical states when they are brought together and measured, in which case it is possible to derive a contradiction between the predictions of quantum mechanics and the predictions of any model which allows some overlap between the ontic states associated with different quantum states. Several authors [30,36,37] have pointed out that the argument fails if we relax the assumption of preparation independence. However, these criticisms of preparation independence have been directed almost exclusively at the claim that the states are independent at the time of preparation. For example, Ref. [36] proposes that even when systems are prepared in a product state they might have global properties not reducible to properties of their subsystems, while Ref. [38] suggests the preparations might fail to be independent due to their common past. Less attention has been given to the fact that PBR must also claim that these independent states persist over time so that the systems still have the same independent states at the time of the measurement. It seems to us that what the PBR theorem really establishes is that *if* quantum systems have localized physical states which persist over time, and the other explicit assumptions of the theorem hold, then those states must carry all the same information as the quantum state. The theorem thus carries no weight against views which deny this spatiotemporal account of physical states in the first place. For example, the theorem *per se* offers no reason to reject any view which suggests that measurement results may depend directly on preparations without being mediated by any persisting physical state, quantum or otherwise.

It is true that there exist later ψ-ontology theorems which come closer to addressing the kind of spatioemporal concerns we raise here. In particular, Hardy's theorem [31] and the Colbeck and Renner theorem [32] both concern themselves with dynamics rather than merely kinematics, and so deal more explicitly with the passage of time. However, these theorems likewise work within the established ontological models framework which presupposes that past events may influence future ones only via the mediation of some persisting physical state, and thus, like PBR, can be circumvented by denying this sort of spatiotemporal account of quantum states. This

suggests to us that a more detailed analysis of possible and impossible relativistic quantum tasks, and of their implications for hypotheses about the temporal persistence of quantum and ontic states, might shed some interesting new light on our understanding of the quantum state and its relationship to reality.

[1] B. Schwartz, *The Paradox of Choice: Why More Is Less* (Ecco/HarperCollins, New York, 2004).

[2] Y. Aharonov and L. Vaidman, *The Two-State Vector Formalism of Quantum Mechanics*, edited by J. G. Muga, R. Sala Mayato, and L. Egusquiza, Time in Quantum Mechanics (Springer-Verlag, Berlin, Heidelberg, 2002).

[3] E. Schrödinger, Die gegenwärtige situation in der quantenmechanik, Naturwissenschaften **23**, 807 (1935).

[4] Y. Aharonov, S. Popescu, D. Rohrlich, and P. Skrzypczyk, Quantum Cheshire cats, New J. Phys. **15**, 113015 (2013).

[5] M. Itano, Perspectives on the quantum zeno paradox, J. Phys.: Conf. Ser. **196**, 012018 (2009).

[6] D. Dieks, Another look at the quantum mechanical entropy of mixing, Am. J. Phys. **56**, 430 (1988).

[7] A. Rae and T. Forgan, On the implications of the quantum-pigeonhole effect, arXiv:1412.1333.

[8] M. L. Almeida, J.-D. Bancal, N. Brunner, A. Acín, N. Gisin, and S. Pironio, Guess your Neighbor's Input: A Multipartite Nonlocal Game with No Quantum Advantage, Phys. Rev. Lett. **104**, 230404 (2010).

[9] T. Vidick, The complexity of entangled games, Ph.D. thesis, University of California, Berkeley, 2011.

[10] L. Mančinska, D. E. Roberson, and A. Varvitsiotis, Deciding the existence of perfect entangled strategies for nonlocal games, arXiv:1506.07429.

[11] M. Tomamichel, S. Fehr, J. Kaniewski, and S. Wehner, One-sided device-independent QKD and position-based cryptography from monogamy games, in *Advances in Cryptology* EU-ROCRYPT 2013*, Lecture Notes in Computer Science, edited by Thomas Johansson and Phong Q. Nguyen Vol. 7881 (Springer, Berlin, Heidelberg, 2013), pp. 609–625.

[12] P. Gawron, J. A. Miszczak, and J. Sladkowski, Noise effects in quantum magic squares game, Int. J. Quantum. Inform. **6**, 667 (2008).

[13] A. Kent, A no-summoning theorem in relativistic quantum theory, Quantum Inf. Proc. **12**, 1023 (2013).

[14] L. Sartori, *Understanding Relativity: A Simplified Approach to Einstein's Theories* (University of California Press, Berkeley, 1996).

[15] D. Dieks, Communication by EPR devices, Phys. Lett. A **92**, 271 (1982).

[16] W. K. Wootters and W. H. Zurek, A single quantum cannot be cloned, Nature (London) **299**, 802 (1982).

[17] H. Barnum, J. Barrett, M. Leifer and A. Wilce, A Generalized No-Broadcasting Theorem, Phys. Rev. Lett. **99**, 240501 (2007).

[18] A. Kent, Quantum tasks in Minkowski space, Classical Quantum Gravity **29**, 224013 (2012).

[19] P. Hayden and A. May, Summoning information in spacetime, or where and when can a qubit be? J. Phys. A **49**, 175304 (2016).

[20] D. Gottesman and I. L. Chuang, Demonstrating the viability of universal quantum computation using teleportation and single-qubit operations, Nature (London) **402**, 390 (1999).

[21] A. Kent, Unconditionally secure bit commitment with flying qudits, New J. Phys. **13**, 113015 (2011).

[22] A. Kent, Unconditionally Secure Bit Commitment by Transmitting Measurement Outcomes, Phys. Rev. Lett. **109**, 130501 (2012).

[23] T. Lunghi, J. Kaniewski, F. Bussières, R. Houlmann, M. Tomamichel, A. Kent, N. Gisin, S. Wehner, and H. Zbinden, Experimental bit commitment based on Quantum Communication and Special Relativity, Phys. Rev. Lett. **111**, 180504 (2013).

[24] E. Adlam and A. Kent, Deterministic relativistic quantum bit commitment, Int. J. Quantum Inf. **13**, 1550029 (2015).

[25] E. Adlam and A. Kent, Device-independent relativistic quantum bit commitment, Phys. Rev. A **92**, 022315 (2015).

[26] C. G. Timpson and H. R. Brown, Proper and improper separability, Int. J. Quantum Inf. **3**, 679 (2005).

[27] R. Horodecki, P. Horodecki, M. Horodecki, and K. Horodecki, Quantum entanglement, Rev. Mod. Phys. **81**, 865 (2009).

[28] N. Harrigan and R. W. Spekkens, Einstein, incompleteness, and the epistemic view of quantum states, Found. Phys. **40**, 125 (2010).

[29] M. F. Pusey, J. Barrett, and T. Rudolph, On the reality of the quantum state, Nat. Phys. **8**, 476 (2012).

[30] M. Leifer, Is the quantum state real? An extended review of ψ-ontology theorems, Quanta **3**, 67 (2014).

[31] L. Hardy, Are quantum states real?, Int. J. Mod. Phys. B **27**, 1345012 (2013).

[32] R. Colbeck and R. Renner, A system's wave function is uniquely determined by its underlying physical state, arXiv:1312.7353.

[33] O. J. E. Maroney, How statistical are quantum states? arXiv:1207.6906.

[34] J. Barrett, E. G. Cavalcanti, R. Lal, and O. J. E. Maroney, No ψ-Epistemic Model can Fully Explain the Indistinguishability of Quantum States, Phys. Rev. Lett. **112**, 250403 (2014).

[35] A. Montina, Exponential complexity and ontological theories of quantum mechanics, Phys. Rev. A **77**, 022104 (2008).

[36] J. Emerson, D. Serbin, C. Sutherland, and V. Veitch, The whole is greater than the sum of the parts: On the possibility of purely statistical interpretations of quantum theory, arXiv:1312.1345.

[37] S. Aaronson, A. Bouland, L. Chua, and G. Lowther, ψ-epistemic theories: The role of symmetry, Phys. Rev. A **88**, 032111 (2013).

[38] M. J. W. Hall, Local Deterministic Model of Singlet State Correlations based on Relaxing Measurement Independence, Phys. Rev. Lett. **105**, 250404 (2010).