

Does Black Hole Complementarity Answer Hawking's Information Loss Paradox?

Peter Bokulich^{†‡}

A proper understanding of black hole complementarity as a response to the information loss paradox requires recognizing the essential role played by arguments for the applicability and limitations of effective semiclassical theories. I argue that this perspective sheds important light on the arguments advanced by Susskind, Thorlacius, and Uglum—although ultimately I argue that their position is unsatisfactory. I also consider the argument offered by 't Hooft for the breakdown of microcausality around black holes, and conclude that it relies on a mistaken treatment of measurement collapse. There is, however, a legitimate argumentative role for black hole complementarity, exemplified by the position of Kiem, Verlinde, and Verlinde, that calls for a more subtle analysis of the limitations facing our effective theories.

1. Black Hole Complementarity. In his 1976 article, “Breakdown of Predictability in Gravitational Collapse,” Stephen Hawking argues that his prediction that black holes emit thermal radiation, and thereby shrink and eventually disappear, implies that the evolution of black holes cannot be described by standard unitary quantum mechanical evolution. This nonunitary evolution is popularly described as representing a loss of “information”: if a pure state nonunitarily evolves into a mixture, then we can no longer predict with certainty the outcome of any complete set of measurements; thus it appears that some previously existing information has been destroyed. This conclusion has been generally viewed as unacceptable by high energy physicists—who have therefore characterized Hawking's argument as a ‘paradox’ that requires a resolution.

[†]To contact the author, please write to: Department of Philosophy, Boston University, Boston, MA 02215; e-mail: pbokulic@bu.edu.

[‡]The content of a substantial part of this paper is drawn from Bokulich (2003, Chapter 5).

Philosophy of Science, 72 (December 2005) pp. 1336–1349. 0031-8248/2005/7205-0056\$10.00
Copyright 2005 by the Philosophy of Science Association. All rights reserved.




Figure 1.

A somewhat less heuristic account of Hawking's argument¹ considers a spacetime with an evaporating black hole, as represented in Figure 1. Because there will be quantum correlations between the field values localized inside and outside the black hole, the exterior state on Σ_{ext} will be mixed. Microcausality then seems to imply that the late-time state on Σ_{late} will also have to be mixed, even if the global state on Σ_{early} was pure. This process therefore cannot be described by unitary quantum evolution, because it begins with a pure state and ends with a mixed state.

The variety of proposed solutions to the information loss paradox grew sharply in the early 1990s with the development of two-dimensional toy models of evaporating black holes. One of the most popular of these responses suggested that once the black hole gets sufficiently small, quantum gravitational processes might shut down the Hawking radiation, leaving behind a black hole remnant that would remain entangled with external fields, thus preserving the purity of the overall state. Other solutions consider novel spacetime features such as branching universes or null

1. For a more careful analysis of the information loss argument, see Belot et al. (1999) or Bokulich (2003).

singularities. Belot, Earman, and Ruetsche (hereafter, BER) have offered a helpful overview of these various proposals in their 1999 paper “The Hawking Information Loss Paradox: The Anatomy of a Controversy.”²

However, the number of solutions that are actively being pursued has dwindled sharply since the so-called “second superstring revolution” of the mid 1990s. In particular, the success of Strominger and Vafa (1996) in recovering the Bekenstein entropy for certain extremal black holes, along with the ensuing interest in the holographic principle and the AdS-CFT correspondence, has led many of the participants in this debate to abandon their previous positions. One no longer finds arguments for black hole remnants, thunderbolt evaporation, or nonunitary superscattering operators in the literature. This is not because these proposals have been decisively ruled out, but is rather because the currently most popular avenue of quantum gravity research, namely superstring theory, promises to offer an escape from the information loss paradox that is incompatible with most of replies that have been offered. Although the matter is far from settled, it appears that superstring theory may offer a completely unitary description of the formation and evaporation of black holes without creating black hole remnants, baby universes, etc.³

This raises the question of how this proposed model responds to Hawking’s argument for nonunitary evolution. The standard understanding among string theorists is that the correct way of resolving the information loss paradox is essentially captured by black hole complementarity.⁴ This position claims that an outside observer can effectively describe the black hole as a heated membrane situated just above the event horizon. This membrane has a finite number of degrees of freedom, given by the Bekenstein entropy of the black hole, and an area one Planck unit (10^{-66} cm²) larger than the event horizon. According to the outside observer,

2. For a critique of BER’s assessment of remnant scenarios and an analysis of the conceptual strategy underlying these scenarios, see Bokulich (2001).

3. More precisely, it is assumed that a (yet to be formulated) nonperturbative theory underlying string theory will be able to describe accurately the evolution of black holes. This theory has been dubbed “M-Theory,” although it is, as yet, indeterminate exactly what the “M” stands for. Historically it seems most likely that the term comes from “Membrane” because the nonperturbative extensions of string theory include multi-dimensional “branes” as well as one-dimensional strings. However, it may turn out to stand for “Matrix” as one contender for this nonperturbative theory is known as “matrix theory.” Although we do not know precisely what this theory looks like, we can catch glimpses of it through certain nonperturbative techniques. This is how we land in the seemingly paradoxical situation of knowing that a theory “exists” but not knowing what that theory is.

4. This assessment is based on personal communication with string theorists such as Banks, Giddings, Maldecena, and Susskind.

any infalling information will interact violently with this heated membrane, and will eventually be reemitted to the exterior universe, thus keeping the late-time state pure.

A difficulty facing this suggestion is the fact that the event horizon of a black hole is a globally defined property of a spacetime; we would not expect a freely falling observer to notice anything unusual at the horizon—and we certainly would not expect her to be destroyed there. Black hole complementarity postulates that *from her perspective* the infalling observer's passage through the heated membrane and the event horizon is indeed unremarkable. Her description of the situation is claimed to be *complementary* to the external observer's description, rather in the way that the descriptions of a quantum particle in terms of position and momentum are complementary.

The deepest conceptual challenge facing this position is that of reconciling these two apparently (if not *obviously*) contradictory descriptions. The most prominent route of justification appeals to a seemingly naive verificationist philosophy. Black hole complementarians argue that it would be impossible for the infalling observer to confirm that the external observer did in fact see her burn up on a heated membrane. Likewise, the external observer cannot verify that the infalling observer survived this part of her trip. Thus, the argument proceeds, we need not insist that these descriptions be consistent with each other, for each observer can consider the other's description to be unverifiable, and therefore “unphysical” on verificationist or operationalist grounds.

In their discussion of BHC (black hole complementarity), BER rightly criticize this appeal to operationalism as an unconvincing invocation of a widely discredited philosophical position. However, there are several avenues of justification for BHC that go beyond this simple appeal to verificationism and also go beyond BER's treatment of the position. A weakness of BER's analysis is that they fail to recognize that BHC, like other positions in the debate over the information loss paradox, is crucially about the applicability and limitations of our *effective* theories. In particular, the primary question that these proposals are meant to address is that of when and where we should expect the descriptions of an effective quantum field theory on a background spacetime of classical general relativity to be valid. This fact is crucial for an understanding of a prominent thread of argument supporting BHC.

2. Nice Slices and the Applicability of Local QFT. Belot, Earman, and Ruetsche are critical of the complementarian's appeal to operationalism not only because it is philosophically unattractive, but also because they do not see this verificationist reasoning playing any important role in responding to Hawking's paradox. As BER correctly point out, most

advocates of black hole BHC are string theorists, who therefore reject the quantum field theoretic framework of Hawking's argument. From this they infer that "in some sense the most philosophical plank of the Complementarian platform is also the most otiose" (BER 1999, 215). Why, given their commitment to a fundamentally non-local theory of quantum gravity, do the complementarians not simply argue that this non-local nature of strings will somehow secure the purity of the late time state?

This avenue is not immediately open to them because even though local quantum field theory (QFT) may fail in the quantum gravitational regime, it clearly offers an extremely accurate account of processes that involve only typical (sub-Planckian) energies. Further, it appears that we have compelling reasons to believe that the event horizon of a large black hole is safely within the domain of validity of low-energy QFT. The argument for this claim plays a foundational role in justifying both the information loss argument and the remnant proposal, although this argument is often not fully articulated. Typically one points out, as we did above, that there is nothing locally significant about the event horizon, and that the local spacetime curvature there can be as small as we like for sufficiently large black holes. Therefore we do not expect any quantum gravitational effects to come into play in this area. However, for this argument to be convincing, we need to be clearer about exactly what a "quantum gravitational effect" would be, and about how we are to specify when we can, and cannot, legitimately ignore these effects.

A more rigorous argument appeals to the fact, commonly raised in the context of renormalization and effective field theories, that high energy degrees of freedom will generally decouple when all energies involved are low—meaning that a description involving only low energy degrees of freedom should adequately describe such a process. Because high energies map onto short distances, this seems to imply that the nonlocal nature of strings—whose length is on the order of the Planck length—can only manifest itself if a process includes Planck-scale energies. If it is possible to run the information loss argument in a situation that avoids such high energies, then it would seem that the *effective* validity of QFT would secure Hawking's conclusion. Such a description of the situation does seem to be possible, however, for we can construct a series of spacelike time slices that avoid all regions of high curvature for as long as possible, *and* for which both the infalling bodies and the outgoing radiation have low energies in the frame of the time slice. Such foliations of spacetime are referred to as 'nice slices'.

Because all of the local energies on nice slices are low, it seems that all processes should be adequately described by a QFT obeying microcausality. Further, it is possible to construct a family of nice slices that pass through the infalling matter and the outgoing Hawking radiation for most

of the evaporation of a large black hole. Therefore, proceeds the argument, the information cannot escape when the black hole is large, but must wait until it shrinks to Planck size and we can no longer define the state of the system as a state of a nice-slice field. At this point we can expect Planckian physics to come into play—but if these quantum gravitational effects are to rescue unitary evolution at this late stage they will have to slow or halt the Hawking radiation and leave us with a remnant.

Black hole complementarians generally want to argue that this argument is flawed. The mere fact that one can construct a low energy field theory on a nice slice need not imply that the nonlocal effects of the underlying theory of quantum gravity cannot manifest themselves. The heart of several accounts of BHC, such as that advanced by Susskind et al. (1993), is the denial of the claim that the validity of local QFT at low energies gives us sufficient grounds to agree with Hawking that the evolution of black holes will be nonunitary.

3. Operationalism and Planckian Physics. Setting aside their most blatant appeals to operationalism, we can find in the presentations of Susskind, Thorlacius, and Uglum (hereafter STU) a more promising line of argument, a *tu quoque* of sorts, for calling into question the legitimacy of a local field-theoretic description. They argue that the assumed validity of a local QFT description in the context of black holes is an assumption standing in need of justification as much as their own assumption that the suggested effective descriptions of the two observers is compatible. This position can be found most clearly in the stated goal of a 1994 paper by Susskind and Thorlacius:

Our aim is limited to challenging the commonly held view that, as there is no strong curvature or other coordinate invariant manifestation of the event horizon, an information paradox can be posed without detailed knowledge of the underlying short-distance physics. (Susskind and Thorlacius 1994, 966)

If the advocates of remnants or information loss are forced to appeal to the nature of Planck-scale physics in their arguments, then it seems that these arguments will only be persuasive to one who shares their assumptions about the nature of a full theory of quantum gravity. If someone has different expectations of quantum gravity, then the argument for information loss or remnants will not be compelling.

Susskind and Thorlacius are focusing on this issue in part because questions had been raised in response to their earlier paper over whether it really would be impossible, for example, for an external observer both

to see the infalling observer destroyed on the heated membrane,⁵ and also to receive verification that she survived this passage. It is clear that no signal from the infalling observer can reach the external observer, for any such signal would have to be superluminal. However, we might ask whether it would be possible for the infalling observer to leave a signal *inside* the black hole, which the external observer could then receive by entering the black hole himself.

It is in this context that Susskind and Thorlacius argue that any such strategy would take us out of the realm of low-energy effective QFT and into the truly quantum gravitational regime. Their arguments rest on the geometry of the black hole spacetime and the fact that the energy required for sending any message will be inversely proportional to the amount of time one has to send it. In our above scenario, for example, if we calculate the amount of time that the two observers would have inside the black hole before encountering the central singularity, we find that the frequency of the signal sent would have to be far beyond the Planck scale. From this, Susskind and Thorlacius conclude that our present knowledge of physics does not allow us to pose legitimately the information loss paradox.

This claim seems to invoke a rather odd argumentative strategy that is worth further scrutiny. Notice that we are not here ruling out the *possibility* of using Planckian energies to send messages. Further, it does not seem that the objection could be that we would not know how to generate, manipulate, or decipher Planck-scale messages. One would assume that Morse code using Planckian gamma rays is not *conceptually* any more problematic than using radio signals. Susskind and Thorlacius' point must be that without a full theory of quantum gravity we have no legitimate grounds on which to rule out their suggestion that the seemingly incompatible effective descriptions of the infalling and external observers are in fact consistent effective descriptions of some underlying theory of quantum gravity.

However, as we saw above, the argument for the applicability of standard QFT relies on our ability to run the argument in such a way that the energies involved are low. The fact that it would take Planck-scale energies to experimentally *verify* this low-energy description seems to be irrelevant. At the very least we are owed an account of why considerations of the energies required to verify the descriptions should be a decisive factor in evaluating the proposals before us, and Susskind et al. do not provide us with such an account. We therefore now turn to other argu-

5. More rigorously, the question is whether the external observer would be able to perform a sufficient number of measurements on the outgoing Hawking radiation to infer the quantum state of the infalling observer. See BER (1999) or Bokulich (2003).

ments that have been offered for the incompatibility of the descriptions of our infalling and external observers.

4. Of Collapse and Commutation. Gerard 't Hooft was the first real champion of the unitary evolution of evaporating black holes and of the position that came to be known as black hole complementarity ('t Hooft 1985). 't Hooft suggests taking as an *Ansatz* the claim that black hole evolution is unitary and trying to develop on this basis a full theory of black hole interactions. The theory he proposes resembles in crucial respects the membrane picture of black holes offered by STU, and it likewise faces the challenge of reconciling the external unitary description of the black hole evolution with the fact that an observer should be able to pass through the event horizon of a black hole unharmed. 't Hooft's response to this apparent incompatibility is much more direct than STU's appeal to operationalism: He argues that measurements on the outgoing Hawking radiation and measurements on the infalling matter performed inside the apparent horizon are measurements of *incompatible observables* represented by noncommuting Heisenberg picture operators, despite the fact that these observables are spacelike related.

This incompatibility is supposed to be grounded in the fact that the commutators between operators $\hat{O}(x_{\text{late}})$ describing the late-time Hawking radiation, and the longitudinal component of the stress-energy tensor at the black hole horizon, $\hat{T}_{--}(h_s)$, grows exponentially with time. It is worth quoting at length the consequences that are claimed to follow from this fact:

In itself, this uncertainty relation would not have been a disaster if the particles causing the large $\hat{T}_{--}$ had been completely transparent. But they are not, because they must be associated with a gravitational field which, because of the infinite energy shifts involved, has the ability to destroy everything attempting to cross the horizon. . . . Thus, we conclude that one cannot describe Hawking particles while at the same time one describes observables, i.e., expectation values of local operators, beyond the horizon. The corresponding operators have commutators which are far too large. One must choose the basis in which one wishes to work: either describe particles beyond the horizon or the particles in the Hawking radiation, but do not attempt to describe both. Physically this means that one cannot have "super observers", observers that register both Hawking radiation and matter across the horizon. The corresponding operators have explosive commutators. (Stephans et al. 1994, 626)

The claim here is that the *incompatibility* between observables localized inside the black hole and late-time observables localized outside the black




Figure 2.

hole is due to the large commutators between these late-time external observables and the stress-energy operator at the horizon. 't Hooft seems to be suggesting that a late-time measurement will, apparently quite literally, destroy all infalling observers and measuring apparatus. This claim is reiterated in his later overview of his approach (I reproduce 't Hooft's figure in Figure 2):

Any measurement made by B , implies the introduction of states obtained from the Hartle-Hawking vacuum by acting on it with operators that create or remove particles seen by B , which for A would be outrageously energetic. These particles would cause gravitational shifts that seriously affect the ingoing objects, including the fragile detectors used by A . Thus these observations cannot be independent. ('t Hooft 1996, 4684)

Taken at face value, this claim that a late-time measurement “implies the introduction of states” that contain particles that are “outrageously energetic” according to the infalling observer is misleading or false. The picture that he has in mind seems to be the following: We begin in a global state that a radially infalling observer would find indistinguishable from a Minkowski vacuum state. This implies that such an observer will detect no particles whatsoever, even when passing through the black hole horizon. A late-time “Rindler” observer, however, would detect thermally distributed particles coming from the black hole. Let us suppose that this observer detects one such particle at time t with energy ω . If we evolve this particle back in time to the point on the horizon where it meets the path of the infalling observer, we find that it has extremely high energy on the order of $\omega e'$. In leaving the gravitational well of the black hole,

the particle's energy was red-shifted down to the relatively small value T recorded by our late-time observer. It therefore appears that this observer can conclude that this highly energetic particle would have destroyed any observer trying to enter the black hole. Although the mixed Hawking state is compatible with—indeed, it is derivable from—a vacuum state experienced by an infalling observer, it would seem that any *actual* measurement outcome is incompatible with such a vacuum. 't Hooft appears to conclude that therefore the late-time Heisenberg-picture measurement will collapse the time-independent state of the field and destroy anything crossing the black hole horizon.

There are, of course, a number of problems facing accounts of measurement collapse (or von Neumann's "type one" evolution), even aside from using the Heisenberg picture to invoke backwards-in-time causation. One difficulty is that of specifying what sort of interactions are to count as 'measurements' that will induce collapse. A second difficulty is that of reconciling the apparently instantaneous nature of collapse with relativity theory. A third difficulty, which would seem to be particularly pressing in the context of the debate under discussion, is that such collapse would be an example of non-unitary evolution. If one thought that the unitary evolution of quantum theory is violated each time we perform a measurement, it is unclear why we should be shocked that the evolution and evaporation of a black hole should also produce such a violation.

However, we can set these worries aside; for even if we allow the appeal to measurement collapse, we still find that 't Hooft's argument is mistaken. The easiest way to see this is to recast the scenario in the Schrödinger picture, which should be equivalent to the Heisenberg picture for all questions of principle. In the Heisenberg picture the causal relations underpinning time evolution are captured by the nonvanishing commutators between operators representing timelike and lightlike related observables. In the Schrödinger picture of particle mechanics and field theory, all operators representing spatially separated observables will commute, and the time dependence will be given by the evolution of the state. Thus the claim that $\hat{T}_{--}(\mathbf{h}_s)$ and $\hat{O}(x_{\text{late}})$ have a large commutator will go over to the Schrödinger picture claim that a state that is an eigenstate of $\hat{T}_{--}$ at a time t_h , corresponding to $\mathbf{h}_s$, will evolve into a late-time state that has a large dispersion for $\hat{O}$. This implies that if we perform a $\hat{T}_{--}$ measurement at t_h the state will change in such a way that we will be able predict very little about the late-time $\hat{O}$ measurement.

Whatever criterion we adopt to decide when the measurement collapse occurs, it is clear that at time t_h either a $\hat{T}_{--}$ measurement has occurred or it has not. Thus the state is either given by unitary Schrödinger evolution or it has discontinuously collapsed onto an eigenstate of $\hat{T}_{--}$. However, neither of these possibilities poses any problem for the health of an

infalling observer, and there seems to be no reasonable mechanism by which a measurement by a late-time Hawking observer could possibly affect the happenings near the horizon. The non-vanishing commutator between the Heisenberg picture operators indicates that the late-time observer would, in principle, be able to tell whether or not the infalling observer performed a measurement of the stress-energy of the field. It does not indicate that the actions of the late-time observer can change the past and destroy the infalling observer. Nor does it indicate that operators localized inside the black hole will fail to commute with operators localized outside the black hole at late times. We therefore still have no reason to believe that the descriptions of our two observers are incompatible with each other.

5. Limitations of Effective Descriptions. It is worth remarking that both the operationalist arguments of Susskind et al. and the demolition argument presented by 't Hooft go to quite extreme lengths in an attempt to block the legitimacy of theoretical descriptions that all parties to the debate admit will be fundamentally limited. We might expect that a more promising avenue would be to investigate more carefully the presuppositions going into the application of low-energy QFT and ask whether there might be a plausible mechanism in a fully quantum gravitational treatment of black holes that could undermine these presuppositions. Such an approach is pursued by Kiem, Verlinde, and Verlinde (hereafter, KVV) in a 1995 paper in which they suggest that black hole complementarity is but one manifestation of a more general limitation of semiclassical descriptions. The substance of their account of these limitations, which they title *space-time complementarity*, is as follows.

According to quantum field theories in curved (classical) spacetime, the state of a system resides in a Hilbert space defined on a Cauchy slice of the spacetime. However, the stress-energy of the field will also influence the geometry of spacetime. This implies that in general different quantum states should, in principle, be defined on *differing* background geometries. Typically, however, we argue that so long as the energy-momentum of the fields we are studying is sufficiently small, we can safely ignore the resulting alterations of the spacetime. This will require us to rule out any states, or interactions, that are so energetic that their gravitational effects cannot be legitimately ignored. We can specify these restrictions by introducing a cutoff length, $\varepsilon(x)$, which will be the limit of allowable field modes.

KVV suggest that the appropriate measure of when our semiclassical measure is no longer reliable is given by the stress-energy fluctuations associated with the cutoff, which typically grow as $\varepsilon(x)^{-4}$ as the cutoff scale shrinks. Once these fluctuations reach the same scale as the cutoff itself, the classical description of the spacetime will no longer be accurate,

and our semiclassical account will break down. However, it is not clear that merely specifying a short-distance cutoff will offer an adequate semiclassical theory, for any such specification seems to violate Lorentz invariance. Two different observers with a large relative velocity will give different accounts of which wavelengths are very small and therefore highly energetic; any choice of cutoffs would seem to privilege one of these observers over the other. KVV's proposal is the claim that *both* observers may employ cutoff scales appropriate their particular frames, and the resulting truncated theories yield *equivalent* effective quantum descriptions.

KVV argue that black hole complementarity is derivable as a special case of this space-time complementarity. The nice-slice argument discussed above demonstrates that this claim is not trivial, for it would seem possible to pick a slicing of spacetime such that we could describe the measurements of both an infalling observer and a late-time external observer using the *same* energy cutoff—which should then imply that these spacelike related observables commute. The central task that KVV undertake in their 1995 paper is a demonstration that this is not the case.

The crucial assumption in KVV's account is the claim that the cutoffs of different observers yield equivalent semiclassical descriptions of the situation. By this, they mean that if both observers use the smallest possible cutoff $\epsilon(x)$ that does not introduce fluctuations larger than $\epsilon(x)$, then each of the resulting set of operators will be a *complete* set of commuting operators for the *same* physical Hilbert space. From this KVV conclude that such operators will, in situations such as black hole evaporation, fail to commute with each other even though they are associated with spacelike separated regions.

This assumption is clearly incompatible with the high energy theory being a field theory, and on its face is a very surprising claim. Why should we assume that these (spacelike related) observables exist in the same Hilbert space, particularly when we have arrived at this effective semiclassical Hilbert space only by throwing out states that support the observables that are supposedly 'complementary' to the observables safeguarded by our cutoff? Presumably, the best hope of providing a plausible answer to this will come from analysis of string-theoretic models, such as provided in Lowe et al. (1995).

Despite its inability to justify this crucial assumption, KVV's position has the advantage of offering us a more plausible account of the incompatibility between the interior and exterior measurements than we found in the positions of 't Hooft and Susskind et al. KVV do not follow 't Hooft in claiming that performing a late-time measurement will destroy any infalling observer or apparatus, thus rendering impossible any field measurement inside the black hole, nor do they appeal to the impossibility

of communicating the results of the two measurements to any single observer. While they do not develop a clear account of the incompatibility of the measurements in question, KVV seem to be taking the much more reasonable route of arguing that even though both of these measurement procedures could be performed, they cannot both *count* as measurements of the observables of interest. The relevant characterization of the low-energy observables measured by our two observers relies on an effective semiclassical theory requiring an appropriate cutoff length, and there is no such theory that includes the two observables in question. We might be able to know the outcomes of both measurements, but there would be no legitimate effective theory that we could use to make sense of these results in an applicable low-energy semiclassical theory. To reply to Susskind and Thorlacius' worry of verifying information duplication, we could argue that there is no way of *knowing* whether the information has been 'duplicated' in the two measurements, for there is no effective field theory that allows us to extract the 'information' from the measurement results in our hands.

In conclusion, while the more extreme arguments undergirding black hole complementarity do appear to be misguided, the position itself presents a coherent and important response to Hawking's argument for the non-unitary evolution of black holes. Ultimately, this response appeals to a quantum gravitational violation of the locality assumptions of QFT that are encoded in microcausality. The conceptually important work that black hole complementarity adds to this response is that of offering an acceptable account of the applicability and limitations of our low-energy semiclassical theories.

REFERENCES

- Belot, Gordon, John Earman, and Laura Ruetsche (1999), "The Hawking Information Loss Paradox: The Anatomy of a Controversy", *British Journal for the Philosophy of Science* 50: 189–229.
- Bokulich, Peter (2001), "Black Hole Remnants and Classical vs. Quantum Gravity", *Philosophy of Science* 68 (Proceedings): S407–S423.
- (2003), "Horizons of Description: Black Holes and Complementarity", Ph.D. dissertation, Notre Dame University.
- Hawking, Stephen W. (1976), "Breakdown of Predictability in Gravitational Collapse", *Physical Review D* 14: 2460–2473.
- Kiem, Youngjai, Herman Verlinde, and Erik Verlinde (1995), "Black Hole Horizons and Complementarity", *Physical Review D* 52: 7053–7065.
- Lowe, David, Joseph Polchinski, Leonard Susskind, Larus Thorlacius, and John Uglum (1995), "Black Hole Complementarity versus Locality". *Physical Review D* 52: 6997–7010.
- Stephens, Christopher R., Gerard 't Hooft, and Bernard F. Whiting (1994), "Black Hole Evaporation without Information Loss", *Classical and Quantum Gravity* 11: 621–647.
- Strominger, Andrew, and Cumrun Vafa (1996), "Microscopic Origin of Bekenstein-Hawking Entropy", *Physics Letters B* 379: 99–104.

- Susskind, Leonard, and Larus Thorlacius (1994), “Gedanken Experiments Involving Black Holes”, *Physical Review D* 49: 966–974.
- Susskind, Leonard, Larus Thorlacius, and John Uglum (1993), “The Stretched Horizon and Black Hole Complementarity”, *Physical Review D* 48: 3743–3761.
- ’t Hooft, Gerard (1985), “On the Quantum Structure of a Black Hole”, *Nuclear Physics B* 256: 727–745.
- (1996), “The Scattering Matrix Approach for the Quantum Black Hole: An Overview”, *International Journal of Modern Physics A* 11: 4623–4688.