

Black Hole Remnants and Classical vs. Quantum Gravity

Peter Bokulich^{†‡}
University of Notre Dame

Belot, Earman, and Ruetsche (1999) dismiss the black hole remnant proposal as an inadequate response to the Hawking information loss paradox. I argue that their criticisms are misplaced and that, properly understood, remnants do offer a substantial reply to the argument against the possibility of unitary evolution in the spacetimes that contain evaporating black holes. I also claim that an understanding of these proposals -- and, more generally, of attempts to formulate a theory of quantum gravity -- requires a recognition of the significant and controversial nature of assessing the limits of our current theories.

1. Introduction. Any attempt to formulate a fully quantum theory of gravity faces the difficulty of establishing the dynamical behavior of the theory when spacetime itself is subject to quantum uncertainties. Typically, our only justification for offering any picture of truly quantum gravitational processes comes from arguing that in certain limits we can trust the models provided by classical general relativity (GR), low energy quantum field theory (QFT), and their hybrid offspring semi-classical gravity. However, we recognize that these models will break

[†] Department of Philosophy, 336 O'Shaughnessy Hall, University of Notre Dame, Notre Dame, IN 46556; e-mail: Bokulich.1@nd.edu

[‡] Thanks to Laura Ruetsche, Don Howard, Jim Cushing, and the Chicago Area Philosophy of Physics Group for helpful comments and discussions.

down in certain regions, specifically in regions where interactions on the scale of the Planck energy are likely to occur. We then face the challenge of using our classical models to the extent that we can, marking off the areas where these models can no longer be trusted, and making reasonable guesses at what takes place beyond those borders.

We can witness many variations of this delicate balancing act in the debate surrounding the Hawking information loss paradox. Both the argument supporting the paradox, and the responses to it, rely on semi-classical methods and guesses about the behavior of systems in the quantum gravitational regime, but there is considerable controversy over the success of these various arguments. Belot, Earman, and Ruetsche (BER) argue in their 1999 overview of responses to the Hawking information loss paradox that a class of proposed solutions to the paradox, namely those that postulate the existence of black hole remnants, fail to adequately answer the challenge posed by Hawking. They claim that the proposed solutions either fail to rebut the argument driving the paradox or they escape only by denying the existence of the explanandum, that is, by denying the existence of black holes.

Here I argue that this criticism of the remnant proposals is misguided on three important counts. First, BER either fail to consider, or misrepresent, a prominent remnant scenario that clearly escapes their argument. Second, they neglect or misunderstand the limits of the classical description of the spacetime postulated by the remnant theorists, and thereby miss the substance of the remnant proposal. Third, the definition of black holes they appeal to is overly restrictive in a quantum gravitational context and leads them to casually dismiss a substantial response to the paradox posed by Hawking. I conclude the paper by arguing that one cannot adequately

understand current debates about quantum gravity without recognizing the role played by arguments over the correct limitations of our classical and semi-classical theories.

2. The Information Loss Paradox. The basis of the information loss paradox is Hawking's 1974 discovery that the theory of quantum fields in curved spacetime implies that black holes will give off thermal radiation at a temperature inversely proportional to their mass.

Conservation of energy implies that the black hole will lose mass through this process, and if nothing halts the evaporation the black hole will eventually disappear completely. This premise of complete evaporation is essential to Hawking's argument, and is denied by some of the remnant proposals that we will be considering in the next section.

Figure 1 is a Penrose conformal diagram representing the formation and complete evaporation of a black hole.¹ The shaded area represents the black hole, the region that is not in the causal past of future null infinity \mathcal{I}^+ . The causal curves that pass through the black hole cannot be extended to \mathcal{I}^+ but instead are assumed to terminate in a curvature singularity indicated by the jagged line. Although the existence of a black hole implies that a spacetime is singular, we will see that many remnant theorists deny that black holes harbor curvature singularities in their centers.

¹ The Penrose conformal diagrams that are used throughout this paper preserve the causal structure of spacetime, but systematically distort spatial and temporal intervals. The null rays (light paths) of any spacetime are represented as lines of 45° , regardless of the curvature of the spacetime. By clearly indicating the lightcone structure of the spacetime, these diagrams capture possible causal (luminal or sub-luminal) paths. For example, they allow us to represent a black hole as a region from which nothing can escape without traveling faster than light, i.e., on a path more than 45° from vertical.




Figure 1

BER employ algebraic QFT to formulate their argument for information loss because of worries about non-unitarily equivalent representations of quantum fields in curved spacetimes. However, in the interests of brevity and clarity, and because nothing in my discussion will depend on this choice, I will use the more standard Hilbert space vocabulary. Consider the three spacelike hypersurfaces, Σ_{early} , Σ_0 , and Σ_{late} , on which we can define our global state at a "time." The hypersurface, Σ_0 , is composed of a section inside the black hole, which we will call Σ_{bh} , and a section outside, labeled Σ_{ext} . We assume that our quantum field begins in a pure state, ρ_{early} , at an early time, that is, on our hypersurface Σ_{early} . Quantum mechanics tells us that this state will then unitarily evolve into another pure state, ρ_0 , on Σ_0 .

At this point in the argument we appeal to the validity of local QFT, which claims that ρ_0 will be a state in $\mathcal{H}_{\text{bh}} \otimes \mathcal{H}_{\text{ext}}$, where these two Hilbert spaces will be defined by local field operators on Σ_{bh} and Σ_{ext} respectively. We can then define the component states ρ_{bh} and ρ_{ext} by tracing ρ_0 over the degrees of freedom associated with Σ_{ext} and Σ_{bh} . These component states, ρ_{bh}

and ρ_{ext} , will be pure states if and only if ρ_0 is an uncorrelated, factorizable state, that is, if and only if

$$\rho_0 = \rho_{\text{bh}} \otimes \rho_{\text{ext}}.$$

However, we have every reason to think that there will be quantum correlations between the field on Σ_{bh} and the field on Σ_{ext} , that is, we expect that ρ_0 will be an entangled state, and thus ρ_{bh} and ρ_{ext} will be mixed states.

A further requirement of locality, often called microcausality, requires all spacelike related observables to commute. This implies that no act represented by a local operator -- for example, performing a measurement, or making local changes to the Hamiltonian density -- can change the state of the fields outside of the light cone of region corresponding to that local operator. BER appeal to microcausality and the fact Σ_{late} is completely spacelike related to the interior of the black hole to justify their “commutation condition.” This condition claims that any local observables associated with the black hole region, for example, any observable that can be constructed from the local field operators on Σ_{bh} , should commute with all local observables defined on the late time slice Σ_{late} . In the framework of QFT, this implies that our late time state, ρ_{late} , will be independent of the values on Σ_{bh} , and that ρ_{late} will unitarily evolve from ρ_{ext} .

But we now recall that ρ_{ext} is a mixed state, which implies that ρ_{late} is also mixed (since unitary evolution preserves the purity or mixedness of states). It therefore appears that the formation and complete evaporation of a black hole cannot be described quantum mechanically. Quantum mechanics describes all evolution as unitary transformations of quantum states, but the evolution described above cannot be unitary since it begins with a pure state, ρ_{early} , and ends with

a mixed state, ρ_{late} . This is Hawking's "paradox."

3. Black Hole Remnants and Remnant Black Holes. Responses to the information loss paradox typically argue that some aspect of the above picture of evaporating black holes will be excluded by a full theory of quantum gravity, and thus we can expect such a theory to retain unitary quantum mechanical evolution. The remnant proposal, in particular, points to the fact that the semi-classical methods used to derive the Hawking effect are clearly invalid when the black hole reaches the Planck mass. Perhaps Planckian physics will offer some way of saving the information contained in the black hole, thus allowing for the unitary evolution of our global quantum state.

We can divide remnant proposals into two categories: stable remnants and long-lived remnants. If a remnant is stable, then Planck scale quantum gravitational effects shut down the Hawking radiation, and the Planck sized black hole (or the black hole remnant, if this is something distinct from a black hole) remains for all future time. All of the information that fell into the black hole remains hidden inside, and we are left in a situation similar to the one predating Hawking's discovery of black hole radiation: black holes are permanent objects into which information can fall, but out of which no information can escape. As long as the black hole, and its information, are prevented from disappearing completely, the troubling conclusion of Hawking's argument (namely non-unitary evolution) is avoided.

Long-lived remnants, on the other hand, eventually disappear. This scenario is motivated by the recognition that, although our semi-classical models indicate that no information can

escape in the Hawking radiation, these models will break down at the Planck scale. This seems to open the possibility that Planckian processes could return the information to the external universe, and once the coherence of the external universe is secured, the remnant could safely pass out of existence. Although this seems to answer Hawking's argument, a glance at Figure 1 might make us worry about the apparent superluminal information transfer that this proposal seems to imply. I explore responses to this concern in Section 5.

Is a remnant supposed to be a black hole, or is it supposed to be some new entity that a black hole transforms into? Both positions have been taken, although some remnant theorists claim that this distinction may not be appropriate for these Planck sized objects. For the purposes of this discussion we will consider remnant black holes, or "residual black holes," and black hole remnants as two distinct proposals.

4. Belot, Earman, and Ruetsche's Assessment. BER's paper gives a rigorous derivation of pure to mixed state evolution in the spacetime of an evaporating black hole. They then offer a taxonomy of proposed solutions to the Hawking paradox based on which premise of the derivation is denied by a particular solution. My attempt to defend remnant proposals from BER's conclusions will require me to focus on the expected breakdown of local QFT and GR at Planck scales. However, BER explicitly assume the validity of local QFT, and of the classical general relativistic spacetime indicated in Figure 1. We therefore face the question of whether my criticisms actually confront BER's argument, or whether they simply point to a possibly fruitful extension of their project. I wish to set this question to one side here. My goal is not so

much to decide whether BER's assessment of the remnant proposal is fair given the project they set out, as to clarify and defend this scenario as a viable solution to the information loss paradox.

Given their goal of offering a taxonomy of responses to Hawking's Paradox based on how the response rejects a premise of their argument, what is BER's evaluation of the remnant proposal? They claim that the proposal faces a "fundamental difficulty" which they pose in the following dilemma.

Either remnants are remnants--that is, of black holes--in which case they do not provide for a satisfying resolution of the Hawking paradox, or they are not remnants--at least, not of black holes--in which case they can do nothing to address the problem of *black hole* evaporation. (BER 1999, 216)

Thus either remnants should be grouped with inadequate, confused responses to the paradox, or they should be seen as (apparently uninteresting) denials of the original assumption that black holes exist. I will argue below that the solution offered by remnants is both satisfying and completely relevant to the question of evaporating black holes, regardless of which horn of BER's dilemma we decide to face.

First, however, let us consider the case of residual black holes, that is, Planck scale black holes that no longer Hawking radiate. It is not clear whether BER take their criticisms to apply to this remnant scenario as well, or whether they consider it to be outside the scope of their discussion, as it denies the premise of complete evaporation. Regardless, this is a prominent remnant scenario that clearly deserves our attention.

Such a scenario has been advocated by Banks (1995), who offers the following Penrose diagram of the situation.

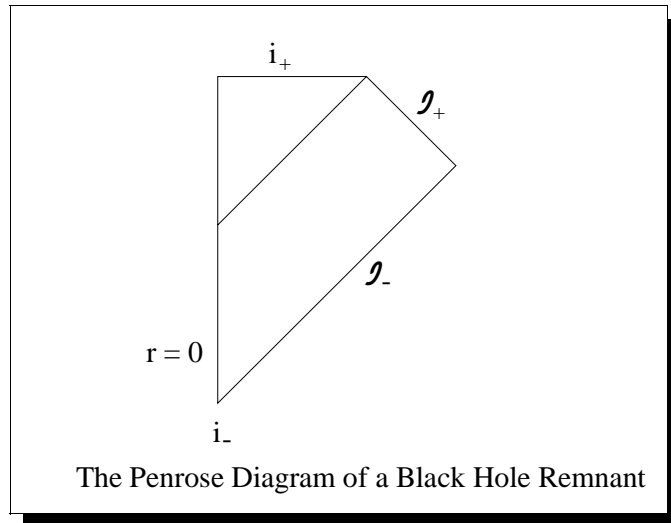


Figure 2

He argues that Planck scale physics may shut down the Hawking radiation and prevent the formation of a singularity in the center of the black hole. The remnant (or residual black hole) will be stable, and will remain entangled with the external world, thus safeguarding the unitarity of the evolution of the universe. The *effective* description of an external observer will not be unitary since the exact state of the interior of the remnant will be inaccessible, but we have still eluded Hawking's argument. I submit that this proposal violates no premise of the information loss argument except the premise of complete evaporation, and while it is beyond the scope of this paper (or BER's) to explore the potentially troubling physical consequences of such a model, conceptually it certainly constitutes an adequate response.

5. Remnants vs. The Commutation Condition. While Banks' remnant proposal explicitly retains the spatiotemporal nature of the interior of a remnant black hole, there are remnant

proposals that assume that Planck scale processes will *replace* the black hole with a new sort of fundamental particle, namely a black hole remnant. This seems to be the proposal that BER are criticizing. I reproduce their diagrams in Figure 3.

A way of picturing remnants which does not involve a residual black hole is given in [Figure 3]. The spacetime in question still has the event horizon structure constitutive of a black hole, so while the remnant (the ??? of [Figure 3]) is not a remnant black hole, it is [a] remnant of a black hole, and so confronts the dilemma's first horn. In this situation one can proclaim as loudly as one wants that information is stored in the remnant. Be that as it may, observables in the algebra associated with post-evaporation slice Σ_{10} ... (stable remnant) or Σ_{11} ... (long-lived remnant) ought to commute with observables associated with the black hole interior. (BER 1999, 216)

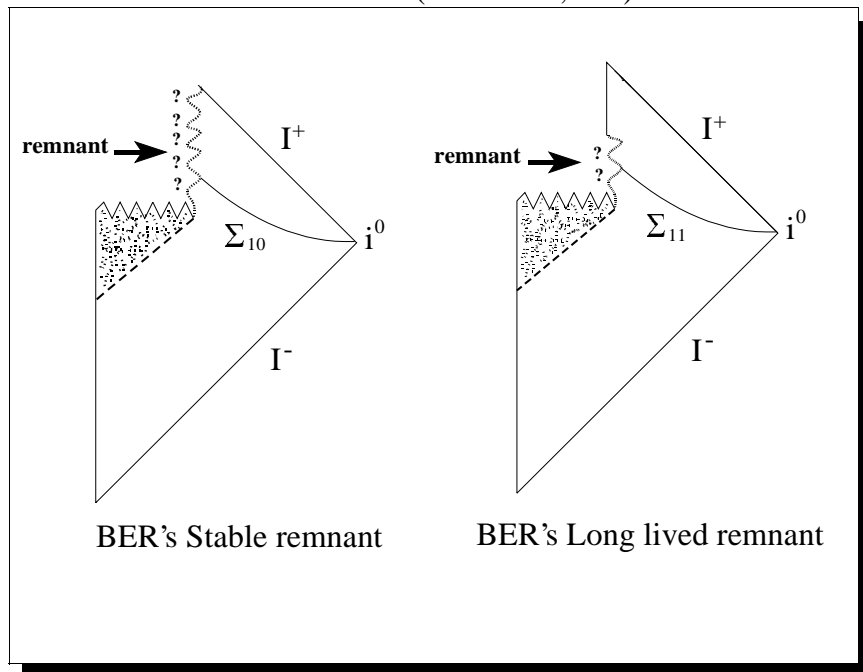


Figure 3

They consider one possibility for fleshing out the remnant story. The remnant (the ??? of their diagrams) could be a placeholder for a set of boundary conditions imposed at the singularity that would preserve unitarity.

But until remnant enthusiasts produce the new physics that incorporates the boundary conditions in a natural way, the present proposal 'solves' the information loss paradox

only by inserting the missing information by hand, and 'remnant' is just a name that does nothing to justify the sleight of hand. (BER 1999, 218)

Leaving aside the question of whether this is a fair appraisal of the attempt to save unitarity through the imposition of boundary conditions, we should recognize that remnant theorists have other resources available. The most common direction for the remnant proposal to take, and I think the most promising, is to deny that there is a true singularity in the center of black holes. (Whether these objects then deserve the title "black hole" is a question I address in the next section.) Indeed, I believe that this is the interpretation that Giddings intended in his (1995) figure that served as the model for BER's diagrams reproduced above.

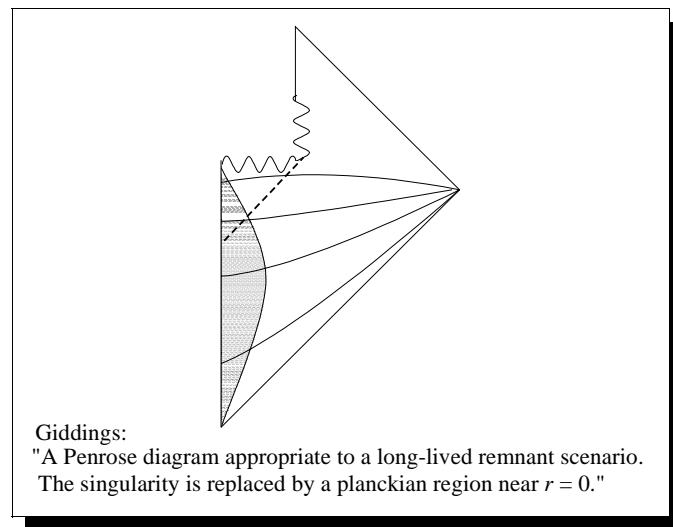


Figure 4

While Giddings' figure (Figure 4) does not explicitly label the remnant, it seems clear that he intends the remnant, and not a curvature singularity, to be the terminus of all causal paths in the black hole. Why should this be a remnant rather than a singularity? Consider the claim that the remnant theorist is making. A large amount of matter undergoes gravitational collapse. This matter steadily loses mass through Hawking radiation, but it remains entangled with the external

world. Eventually the entire Planck mass black hole is replaced by a Planckian remnant. In the absence of evidence for a contrary interpretation, it would seem that we should take the claim that the information of the infalling matter is transferred to the remnant to *imply* that the remnant is in the causal future of this matter.

Giddings' figure represents a long-lived remnant, but we can easily adjust BER's representation of a stable remnant as indicated in Figure 5.

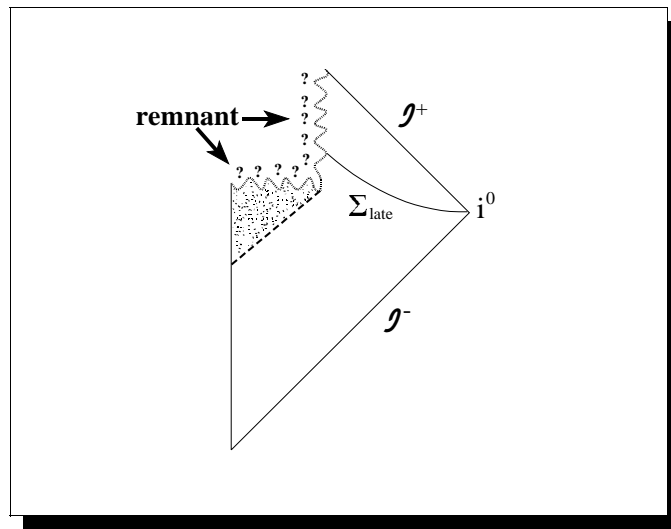


Figure 5

It is worth comparing this figure with Figure 2, the depiction of Banks' residual black hole. If matter is allowed to enter the remnant, and if the remnant subsequently gives off thermal radiation, the two pictures could arguably be considered two representations of the same situation, although there might be some specifically spatiotemporal degrees of freedom in a residual black hole that would not be present in a fundamental remnant.

Now the mere act of postulating a spacelike remnant replacing the singularity of classical GR will not assuage our worries about violations of microcausality. Even with this revised

picture, operators in the black hole are apparently still spacelike related to all operators on Σ_{late} , even if these degrees of freedom are not supposed to be accessible to the outside world. But the remnant theorist has (at least) two ready responses to this worry.

The first option is to deny that the commutation condition holds at the Planck scale. For example, remnant theorists typically claim that remnants are *fundamental particles*, which implies that it is ill defined to ask whether spacelike related operators on the remnant will commute with each other. If remnants are *fundamental*, then there can be no "independent parts" of these objects. This means that there are no *separate* observables (associated with different locations on the remnant) that might or might not commute. This move is not necessarily *ad hoc*, as the majority of remnant proponents are string theorists, or at least sympathetic to the string picture of quantum gravity. In string theory we may not *have* local observables *within* the string, or representing *parts* of the string. If there are no observables that can be associated with the "part of the string here" and the "part of the string there" then we cannot sensibly ask whether such observables commute. Thus one can deny the commutation condition by denying the validity of the framework in which it is phrased. We may wonder why string theorists would accept the fundamentals of the argument for information loss when this argument is based on local QFT. The short answer is that they accept this argument because they have an understanding of the *domain of applicability* of these theories. I will have more to say about this in Section 7.

The remnant theorist's second option is to retain the framework of local QFT, but argue that the microcausality fails at the Planck scale because the causal structure of spacetime is not

completely definite. Fluctuations in the spacetime geometry will imply that the question of whether two points are spacelike related or timelike related is not well defined. This gives us reason to expect the commutation condition to break down at these scales, not necessarily because spacelike related operators will fail to commute, but because the question of whether two operators *are* spacelike related will not admit of a univocal answer.

If the black hole region is not definitely spacelike related to Σ_{late} in Figure 5, we face the question of whether this diagram is a coherent representation of the postulated spacetime. The key to answering this question is to recognize that the diagram is *limited*. The classical regions have been accurately portrayed, but not all of spacetime is amenable to this classical description. Our classical models offered by GR simply do not accurately capture the Planckian region in and around the remnant. We misunderstand the proposal if we argue that because there are two regions of spacetime that *would* be spacelike related if the spacetime were completely classical, therefore operators associated with these regions "ought to commute." An essential part of the black hole remnant scenario (as opposed to the residual black hole scenario) is that we can specify certain regions where our classical approximations break down, and it would be a mistake to insist on the commutation condition in this regime. Remnant theorists are typically quite clear on this point. For example, immediately before his diagram of a long-lived remnant (Figure 4 above), Giddings writes:

Another possibility is that the information is radiated after the black hole reaches $M \sim m_{\text{pl}}$ and the semiclassical approximation fails. Here ordinary causality no longer applies to the interior of the black hole, and it's quite plausible that the information escapes.
(Giddings 1995, 551)

Although an explicit rejection of the commutation condition at Planck scales is prevalent in the

remnant literature, theorists often do not distinguish between the two possible justifications for this rejection. This is partly based on a suspicion that the distinction may not hold up in a full theory of quantum gravity, and partly based on the fact that theorists are more interested in developing viable models than in clarifying the expected failure of this condition.

6. Remnants vs. the Definition of Black Holes. The above defense of remnants might seem to simply leave us facing the second horn of BER's dilemma, for on my reading of the remnant proposal there will be no points of spacetime that are not, strictly speaking, in the causal past of \mathcal{I}^+ . BER criticize Giddings' (1992) massive remnant scenario on the grounds that it denies the existence of a global event horizon. Giddings' proposal is sketched in Figure 6 (reproduced from Giddings (1992), Figure 3). Here the collapsing matter does not form a singularity, but rather a Planckian "core" that superluminally expands past the horizon, thus allowing the information to return to the external universe.

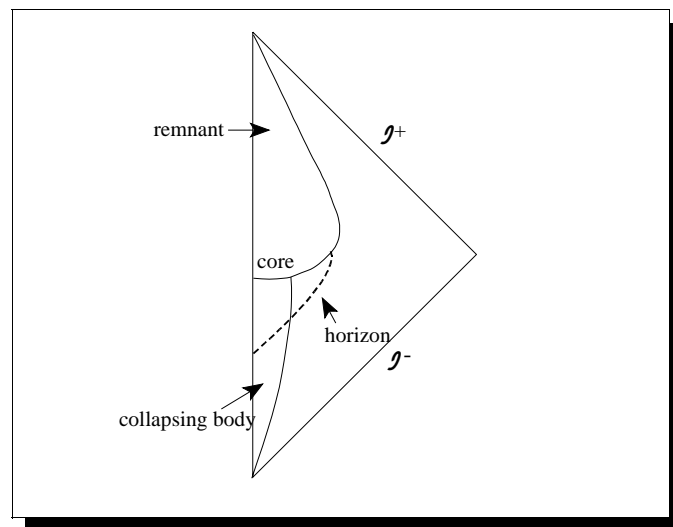


Figure 6

BER question whether it is appropriate to refer to this proposal as a solution of the problem of information loss in black holes, since there is no true black hole in Figure 6.

One disturbing feature of this proposal is that the core expands at a superluminal rate. Worse, [Figure 6] incorporates no singularity of gravitational collapse. ... Because the spacetime lacks a genuine black hole, the surface labeled 'horizon' in [Figure 6] must be an apparent horizon, an object locally delineated, rather than a true event horizon... And because the relevant event horizon structure is missing, the theorems that underwrite Hawking radiation do not apply. In short, the labels of 'black hole' and 'black hole evaporation' strike us as misnomers when applied to [Figure 6]. ... It is less a solution to the information loss paradox than a sweeping denial of the problem. But perhaps that is his point. (BER 1999, 219)

While I cannot fully address all of the objections BER raise here, a number of points should be made on behalf of remnant proposals.

The question of how we should use the term "black hole" is, of course, not the significant issue here. We can, if we like, identify black holes by the formation of apparent horizons. This was the concept that Wheeler first applied the term to in 1967, five years before Hawking introduced his account of global event horizons. Alternatively, we could, with very little violence to the term, identify a black hole with the region from which all causal curves necessarily terminate in a Planckian remnant. We might call the border of this region a "remnant horizon" if we wanted to distinguish it from an apparent horizon. Figure 6, illustrating Giddings' massive remnant proposal, clearly intends the horizon to be an apparent horizon, since the core is supposed to protrude beyond it. On the other hand, Figure 4 above picks out the black hole by a remnant horizon. Note that nothing could escape this region without passing through the Planckian remnant, a prospect not even a fundamental particle is likely to survive. But nothing of consequence hangs on this terminological issue.

The substantial question is whether apparent horizons or remnant horizons will be adequate for the derivation of Hawking radiation. Presumably it is this worry that prompts BER to charge Giddings with offering a "sweeping denial of the problem." If he were denying the force of the argument for Hawking radiation, there might be something to this charge (although I would still worry about the distinction between solving a problem and denying that there is a problem). But, while it is strictly true that Hawking's arguments rest on the presence of a global event horizon, it is also true that both the argument *for* information loss and the *responses* to it rest on approximations and expectations based on theories that *break down* at some point when trying to describe an evaporating black hole. Hawking's calculations are clearly invalid when we reach the Planck scale, so the extrapolation that the black hole completely disappears, *and the assumption that we have global event horizon*, rest exclusively on our expectations of how a full theory of quantum gravity will describe the situation.

Without a true global event horizon the derivation of Hawking radiation does not rigorously go through, but we nonetheless have good reasons to expect that complete gravitational collapse and the formation of an apparent horizon *will* lead to the generation of Hawking radiation. The situation is analogous to our expectation that an observer accelerating for a finite period in flat spacetime will measure a temperature proportional to her acceleration, even though the derivation of the Unruh effect strictly requires us to consider an observer who is constantly accelerating for *all* time. We consider an idealized situation in order to perform our calculations, but then we expect that these calculations tell us something about the real world where observers only accelerate for a finite time and where (perhaps) there are no true global

event horizons to black holes. The question is how far we can *trust* the semi-classical arguments, given that our world surely does not strictly obey semi-classical laws. There is nearly universal agreement that the prediction of Hawking radiation is robust enough for us to expect it to be a feature of any full theory of quantum gravity, even one that denies collapse singularities and global event horizons.

But perhaps BER are wondering what all of the hullabaloo is about if at the end of the day there are no true (global) event horizons. Isn't it obvious that we can save information by allowing all causal curves to extend into the future? The answer is that this solution is not as easy as it seems. Remnant theorists need to do more than simply deny the existence of a global event horizon, they also must argue that we can reasonably *expect* the behavior they describe from a theory of quantum gravity. Further, it appears that remnants would have to have an infinite number of degrees of freedom to hold all of the information that could fall into a black hole, despite the fact that they have a finite mass. This seems to imply that remnants should be pair produced without bound, which is highly problematic. The difficulty facing remnants is not a conceptual one, but rather the *physical* challenge of squaring this proposal with our other cherished theories.

7. Quantum Gravity and Border Disputes. I would like to conclude by suggesting that the forgoing discussion offers some lessons for exploring scenarios in which we expect our current theories to break down, as in the case of evaporating black holes. The first point is that we should take note of the limits that various proposals suggest for our current approximate theories.

It is important, for example, to distinguish between proposals, such as some remnant scenarios, that postulate violations of the commutation condition over Planckian distances, and other proposals that claim that this condition will break down over macroscopic, or even astronomical scales.

Secondly, we should realize that we can typically do more than report the fact that our theories are not completely reliable; we often have the resources to develop some picture of *where* and *how* the theories will break down. In the case of black hole evaporation, for example, we can appeal to the fact that local QFT begins to run into problematic infinities when we consider very short distances, and the fact that interactions at Planckian energies should themselves create black holes, to argue that our current theories are inadequate in such a regime. But by the same token, these arguments seem to indicate that the semi-classical approximation *should* be adequate where these energies are absent.

BER recognize that QFT and GR are only approximately true, but they claim that "we cannot know how good the approximation is, or even what 'approximately' means, until we know how to combine QM and GR in one theory" (BER 1999, 221). While there is a good deal of truth in this statement, it neglects the fact that we do have resources, however imperfect, for finding the limits of our existing theories. Further, we can hope to glimpse some features of a full theory of quantum gravity by working to establish the proper domain of the low energy theories we expect it to reduce to. For this reason we should see the ??? in Figure 5 not merely as a placeholder for some future physics, but also (and perhaps more substantially) as a proposal for the proper *border* of our semi-classical description of the situation. Of course, remnant

theorists do more than demarcate the expected breakdown of our current theories; they also search for models (however idealized) that exhibit the predicted behavior. But the question of when we should expect the effects of Planckian physics to become manifest is an important and controversial question.

To see this it will be helpful to consider, briefly and superficially, one other response to the information loss paradox: black hole complementarity. Advocates of this view, who are again nearly all string theorists, argue that Planckian physics will not be restricted to the center of a black hole, but will become apparent at the event horizon as well. They justify this assumption by pointing out that the modes of the quantum fields that support Hawking radiation become highly energetic when propagated back in time to the event horizon, and by arguing that there are in principle limits to the low energy experiments that one could perform around the event horizon of a black hole.

All parties to the debate agree that our semi-classical theory will break down when we encounter Planckian energies. This particular debate focuses on what *sort* of energies are relevant to this characterization. Remnant theorists appeal to the equivalence principle to argue that Planckian effects will be negligible until one reaches the center of a black hole, black hole complementarians counter by pointing to the high energies involved in descriptions that external observers will offer of the region around the event horizon. While we clearly cannot explore this debate here, it is important to realize that the issue cannot even be framed unless we recognize the significance of arguments that attempt to establish the proper limits of our current theories.

REFERENCES

- Banks, T. (1995), "Lectures on Black Holes and Information Loss", *Nuclear Physics B Proceedings Supplement* 41: 21-65.
- Belot, G., Earman J., and Ruetsche, L. (1999), "The Hawking Information Loss Paradox: The Anatomy of a Controversy", *British Journal for the Philosophy of Science* 50: 189-229.
- Giddings, G. (1992), "Black Holes and Massive Remnants", *Physical Review D* 46: 1347-52.
- _____. (1995), "Quantum Mechanics of Black Holes", in E. Gava, A. Masiero, K. S. Narain, S. Randjbar-Daemi, and Q. Shafi (eds.), *1994 Summer School in High Energy Physics and Cosmology. ICTP Series in Theoretical Physics-Vol. 11*, Singapore: World Scientific, 530-74.