NOÛS 34:4 (2000) 622–633

Before-Effect and Zeno Causality[1]

JOHN HAWTHORNE
Syracuse University

In his 1964 monograph *Infinity: An Essay in Metaphysics*, Jose Benardete presents the following intriguing puzzle:

> Let the peal of a gong be heard in the last half of a minute, a second peal in the preceding 1/4 minute, a third peal in the 1/8 minute before that, etc. *ad infinitum*. … Of particular interest is the following puzzling case. Let us assume that each peal is so very loud that, upon hearing it, anyone is struck deaf—totally and permanently. At the end of the minute we shall be completely deaf (any one peal being sufficient), but we shall not have heard a single peal! For at most we could have heard only *one* of the peals (any single peal striking one deaf *instantly*), and which peal could we have heard? There simply was no first peal. We are all familiar with various physical processes that are followed by what are called after-effects. We are now tempted to coin the barbarous neologism of a *before-effect*.[2]

He goes on to offer some more examples:

> A man is shot through the heart during the last half of a minute by A. B shoots him through the heart during the preceeding 1/4 minute, C during the 1/8 minute before that, &c. *ad infinitum*. Assuming that each shot kills instantly (if the man were alive), the man must be already dead before each shot. Thus he cannot be said to have died of a bullet wound. Here, again, the infinite sequence logically entails a before-effect. Consider now the following even more radical version of this paradox. A man decides to walk one mile from A to B. A god waits in readiness to throw up a wall blocking the man's further advance when the man has travelled 1/2 mile. A second god (unknown to the first) waits in readiness to throw up a wall of his own blocking the man's further advance when the man has travelled 1/4 mile. A third god … &c. *ad infinitum*. It is clear that this infinite sequence of mere intentions (assuming the contrary-to-fact conditional that each god would succeed in executing his intention if given the opportunity) logically entails the consequence that

> the man will be arrested at point A; he will not be able to pass beyond it, even though not a single wall will in fact be thrown down in his path. The before-effect here will be described by the man as a strange field of force blocking his passage forward.[3]

What, if anything, can be learned from these cases? Section one explores a case that is structurally similar to the above puzzle cases, though rather easier to think about. Having unpacked that puzzle case, we shall be in a position to defuse a number of the before-effect cases presented above. Section two explores the "more radical version of the paradox".

Section One

Consider a world where a series of walls are laid out on a two mile stretch of road in the following way: The road has two endpoints, A and B. At B, which is two miles from A, there is the surface of a wall which is a foot thick, the other surface being two miles plus one foot from A.[4] At the point between A and B that is one and a half miles from A, there is the surface of a wall which is half a foot thick, the other surface being one and a half miles plus half a foot from A. At the one and a quarter mile point, there is a wall that is quarter of a foot thick … and so on. There are thus infinitely many walls, such that for each wall there are infinitely many walls closer to A than that wall. For convenience, let us suppose that each wall has a number tag such that the wall at B is numbered '1', the wall next furthest from A is numbered '2' and so on.

A sphere made of material y is rolled from A. There are no objects standing betwen the sphere and the walls and hence the first mile stretch is empty. Let us stipulate that the stretch is on a slight incline, inviting the sphere to naturally roll towards B. It is not causally possible for it to burrow into the earth, nor for it to leave the ground. The motion of y-objects is continuous: if a y object is at p1 at t1 and p2 at t2 then it has travelled some continuous path connecting p1 and p2 between t1 and t2.[5] We stipulate further that each wall is impenetrable by y-constituted objects. In addition, each wall is rigid with respect to y objects: that is, upon contact with a y object, each wall will remain immobile with respect to the ground.[6] Let us also stipulate that there is nothing else in the vicinity—stampeding elephants, sphere interceptors, and so on—and nothing else by way of causal powers belonging to individual walls—extra repulsive forces and so on—that is importantly relevant to the behavior of the sphere or the walls.

Some may complain that the world is too distant to be worth being interested in. Actual walls do have extra repulsive forces, don't get to be rigid and impenetrable at any thickness and so on. Such a reaction is far too hasty. Distant worlds can often be either revealing or therapeutic with regard to our actual conceptual scheme. So let us press on. What happens?

The mixture of circumstances and causal powers—and in particular the impenetrability and rigidity of the walls—make true the following set of material conditionals:

C1 If the sphere hits wall 1, it does not proceed beyond the boundary of wall 1.
C2 If the sphere hits wall 2, it does not proceed beyond the boundary of wall 2.
...and so on.

Further, we have been told that in order to reach wall 1 the ball has to proceed beyond wall 2 and so on. Such facts make true the following set of material conditionals:

D1 If the sphere hits wall 1, it makes contact with wall 2 and proceeds beyond the boundary of wall 2.
D2 If the sphere hits wall 2, it makes contact with wall 3 and proceeds beyond the boundary of wall 3.
...and so on.

From these sets of material conditionals, we can deduce that

E1 The sphere does not hit wall 1.
E2 The sphere does not hit wall 2.

For, assume that it hits wall 1. We can then deduce a contradiction from C2 and D1. And so on.

Each wall has a natural number assigned to it and to no other wall. So it is true to say in the scenario that

P1: If the sphere hits a wall, then for some natural number N, the ball hits wall N.

If we claim further that

P2 The sphere hits a wall.

we will arrive at an inconsistency, since P2, P1 and E1, E2 ... are inconsistent. Sure enough, this is only an omega inconsistency, since P2, P1 and E1, E2 and so on do not comprise a finite list and no finite subset of that list is inconsistent. But omega inconsistency is bad enough, since an omega inconsistent set of statements can't all be true at the same time. Since P1, E1, E2 and so on are in effect built into the description of the world, we should reject P2. We should thus conclude that in worlds that satisfy the original description, the sphere does not make contact with a wall.

It also follows from the description of the world that

P3 If the ball proceeds beyond one mile, it makes contact with a wall.

(Note in this connection the importance of the continuity of motion requirement as well as the "no stampeding elephants" clause.)

So, by modus tollens we can conclude that the ball does not proceed beyond one mile.

(What if we stipulate in addition that the laws of nature in the world are such that the ball will proceed unless it makes contact with a wall? We will then have an inconsistent description which at once entails that the ball will not hit a wall and which also entails that the ball will hit a wall. There is no possible world satisfying that description and so the question as to what happens in such a world is illegitimate.)

We can thus deduce what happens: The ball does not proceed beyond a mile and it does not hit a wall. Does it stop dead at the mile mark or rebound? The description of the case offers no determinate answer to the question. For each world that satisfies the description, one or two of the following scenario's will obtain:

A There are laws of nature which determine what the sphere will do upon reaching the mile mark, in which case the sphere obeys those laws.

B There are no laws of nature which determine what the sphere will do upon reaching the mile mark, in which case the behavior of the sphere upon hitting the wall is causally undetermined.

The logic of the wall case is pretty straightforward. What can be learnt from it? There is a real danger of thinking one has learned too much. For it is tempting to suppose that one has learned from this case that one can conjure up action at a distance out of very mundane objects that do not, when finitely combined, ever act at a distance. It seems that we have taken a bunch of walls that act by contact and conjured up a scenario where a sphere is brought to a stop without contacting anything.

In discussing a similar case (one where a man undertakes to crash into a series of boards ordered rather as the walls are ordered), Benardete seems to offer just this diagnosis: "The infinite series of boards logically entails what we may describe as a field of force which shuts us out from further advance."[7] However, if we think carefully about contact, we shall see that no such conclusion can be reached.

Suppose space to be continuous. Suppose each thing to enjoy determinate location so that it determinately occupies a particular region which is in turn constitued by a particular set of points. There are two kinds of surfaces a thing can enjoy: A thing may have an open surface, enjoyed insofar as that surface occupies an open region (which is such that there is no outmost layer of points).

Alternatively, a thing may have a closed surface, enjoyed insofar as that surface occupies a closed region (which is such that there is an outmost layer of points).[8] What is contact? There are three cases to consider. I offer what I take to be the most intuitive gloss on ‘contact’ in each case:

A closed surface contacts an open surface insofar as there is no unoccupied space in between the two surfaces. Call this open-closed contact.

An open surface contacts an open surface insofar as there is no more than a line’s breadth of unoccupied space between them (the line can then be called the boundary of the two surfaces). Call this open-open contact.

A closed surface contacts a closed surface insofar as the outer skin of each overlaps. Call this closed-closed contact.

Consider the fusion of walls.[9] Call it Gordon. On reflection it is clear that the sphere contacts Gordon. Gordon has an open surface. When the ball stops proceeding at the one mile mark, there is no unoccupied space between the sphere and Gordon.[10] Contact occurs (which may be open-open or open-closed depending on the nature of the sphere’s surface). So the ball is stopped by contact: The ball hits something, though the thing that it hits is not one of the walls.[11]

(The case we have been describing thus provides no metaphysically exciting locus of action at a distance. If you want a surprising source of action at a distance, then impenetrability—even as between finite things—is the place to look. If closed surfaces are impenetrable to the extent of prohibiting even minimal overlap, then the causal power of impenetrability associated with closed surfaces is a power of acting at a distance. The limiting case of impenetrable point particles illustrates this nicely: By virtue of being impenetrable, point particles will constrain each other’s trajectory whenever they are headed towards each other, even though those point particles will never come into contact.)

Insofar as we were originally puzzled by the case of the walls, I believe it is because we are not clear about how to relate fusions to contact. The following principles hold in full generality:

If y is the fusion of x’s and the x’s are impenetrable, then y is impenetrable.
If y is the fusion of x’s and the x’s are immovable, then y is immovable.

Consider though what we may call “The Contact Principle”:

If y is the fusion of x’s and z contacts y, then z contacts one of the x’s.

That principle holds for the finite case. But it is false if the x’s are infinite in number. Once we are clear about this, there is no residual puzzle, nor anything further to learn about the wall case. It is clear what happens in worlds that satisfy the original description: At a mile, the ball makes contact with the fusion of walls, which is rigid and impenetrable. As a result, it does not proceed further. The ball does not, however, make contact with any wall.

Section Two

We are now in a position to defuse some of the before-effect cases. Consider the bullets penetrating the heart. The assumption was "each shot kills instantly". What does that mean? Does it mean that each bullet kills upon contact with the heart? That seems very strange. Let us say that the laws of nature are such that if a metal object penetrates 1/4 inch into the heart, then the person dies at that very moment. We need some such stipulation in order for the case to admit of rigorous reasoning. By parity of reasoning with the above, we can now say that in such a case, the fusion of the bullets will, upon penetrating the heart to 1/4 inch, kill the person. At the point at which the fusion penetrates 1/4 inch, no bullet will have so penetrated. So we should say that the fusion of the bullets kills the person without any bullet doing so. Any puzzlement will be removed once we recall the falsity of the contact principle.

Similar remarks apply to the gongs. Assume that each gong creates a sound wave—a sort of wall of sound. Then the sequence of peals will generate, it seems, an open ended series of walls of sound. Assume that contact of a sound wave is sufficient for deafening and we can happily say that the fusion of the walls of sound cause deafening upon contact with the ear. (If one insists that only contact with a sound wave and not with a fusion of sound waves can cause deafening, then the description will pertain to no possible world. That is like building into the wall case the requirement that only a wall can stop the sphere proceeding.)

What remains troubling is the "yet more radical version" of the paradox, one that cannot, it seems, be unpacked simply by thinking hard about contact. After all, in that case there is not a wall thrown down and so there is no fusion of walls with which to come into contact. To avoid distraction that might be caused by talk about Gods, let me rewrite the problem a bit:

There is an infinite series of assassins, each tagged with a natural number, no pair tagged with the same number, no number that isn't tagged to some assassin. Assassin 1 is disposed to attack Bob with a machete if Bob is still around at 2 pm. If he attacks, he will take half an hour to kill Bob. It is causally impossible for assassin 1 to attack Bob and fail to kill him within half an hour. Assassin 2 is disposed to attack Bob with a machete if Bob is still around at 1:30 and will take quarter of an hour to do it. It is causally impossible for assassin 2 to attack Bob and fail to kill him within quarter of an hour, and so on. Each assassin is unsurvivable as far as Bob is concerned. (Notice that the unsurvivability of a particular assassin attack corresponds to the impenetrability of a particular wall, while the fixed time threshold for an assassin's success corresponds to the rigidity of a wall.) For each time which is such that an assassin is disposed to begin attacking Bob at that time, there are infinitely many assassins which are disposed to attack Bob earlier.

To avoid the distractions of thinking of conscious beings, we can do yet more rewriting (those who prefer thinking about assassins to thinking about point particles might prefer to reckon with the above case rather than the one that follows). Make the assassins A-type point particles. Make Bob a B-type point particle. The A particles are laid out on a two mile line to the right of Bob. No. 1 particle is two miles away. No. 2 particle is one and a half miles away. No. 3 particle is one and a quarter miles away. And so on. Assassination is the transformation of a B particle into an A particle. This occurs by an A particle interpenetrating a B particle at t and at that time irradiating x radiation—this being what constitutes "attack". "Survival" here and in what follows is "remaining as a B-particle". The causal laws and state of the world are such as to make the following three claims true:

> Particle 1 will not move with respect to the point occupied by Bob before 2 pm.
> If at 2 pm Bob still exists as a B particle, then by 2.30, Bob will have been assassinated by particle 1.
> If at 2 pm Bob doesn't exist as a B particle, particle 1 will not move with respect to the point occupied by Bob between 2 pm and 2:30 pm.

and also to make the following three claims true

> Particle 2 will not move with respect to the point occupied by Bob before 1:30 pm.
> If at 1:30 pm Bob still exists as a B particle, then by 1:45 Bob will have been assassinated by particle 2.
> If at 1:30 pm Bob doesn't still exist as a B particle, particle 2 will not move with respect to the point occupied by Bob between 1:30 and 1:45.

and so on.

Let us also assume that the causal laws are such as to make true the following conditional:

> If Bob changes into a B-particle he will never change back into an A particle.

(The case as I have described does rely on the possibility of superluminous signals. As the number of N gets higher, assassin N will have to move quicker and quicker. Given that none are closer than a mile, and that the time required for assassination approaches 0 as N gets higher, infinitely many assassins would have to be able to move faster than the speed of light. Assuming with orthodoxy that the laws of nature are contingent, I don't see this as undermining the possibility of the scenario I have described. In any case, we can change things so as to not require superluminous signals by letting shorter distance compensate in the right way for speed of kill—so that the assassins are laid out in the mile immediately to the right of Bob so their fusion occupies an open line whose external boundary point is the point occupied by Bob.)

What happens? Once again, the logic of the case is pretty straightforward. It is part of the case that Bob can't survive any of the relevant attacks. We can thus generate the following material conditionals:

C1* If Bob is attacked by assassin 1, he does not survive the attack.
C2* If Bob is attacked by assassin 2, he does not survive the attack.
…and so on.

The case also requires that in order to be attacked by assassin 1, Bob has first to survive the attacks of each of assassin 2, 3 and so on. It thus requires the truth of the following material conditionals:

D1* If Bob is attacked by assassin 1, he is attacked by assassin 2 and survives the attack.
D2* If Bob is attacked by assassin 2, he is attacked by assassin 3 and survives the attack.
…and so on.

We can deduce:

E1* Bob is not attacked by assassin 1.
E2* Bob is not attacked by assassin 2.
and so on.

Each natural number is assigned to one assassin. So it it true to say in the scenario that

P*: If Bob is attacked, then for some natural number N, assassin N attacks Bob.

If we claim further that

P2* Bob is attacked by an assassin.

we will arrived at an inconsistency, since P2*, P* and E1*, E2* … are omega inconsistent. We should thus conclude that Bob is not attacked by an assassin.

It is also part of the description of the world that

P3 If Bob survives past 1 pm, he is attacked by an assassin.

So by modus tollens we can conclude that Bob does not survive past 1 pm.

(What if we stipulate in addition that the laws of nature in the world are such that Bob will survive unless he is attacked by an assassin? There is no possible world satisfying that description and so the question as to what happens in such a world is illegitimate.)

Is there anything we have learnt from this case? I suggested that the puzzlement associated with the case described in section one can be traced to a faulty principle that is all too readily relied on in our thinking, namely, that if z touches the fusion of x's it touches one of the x's. By parity of reasoning with the wall case, we can say that the fusion of the assassins cause the assassination of Bob, even though no individual assassin causes the assassination of Bob. But the question remains: "HOW does the fusion cause Bob to be assassinated?". The puzzlement resides in the fact that we think of the assassins as individually having to do something—whether it be swinging a machete or irradiating x radiation or…in order to produce a certain effect c. Yet the assassin fusion seems to accomplish effect c without doing anything at all. (*That* puzzle didn't arise in the wall case because the walls weren't required individually to be able to do anything in order to individually produce the relevant effect, namely stopping the motion of the sphere.) Its not as if there is a super-machete (or its radiatory correlate) that is used to assassinate. If x is the fusion of y's and the y's don't move with respect to z, x doesn't move with respect to z. So it follows that the fusion causally secures the assassination of Bob without even moving! Nor does the fusion need to undergo any other type of change at all in order to assassinate Bob. Our puzzlement thus relies, I suggest, on our tacit endorsement of the following principle relating fusions to their parts, which we can call the "Change Principle"

> If x is the fusion of y's and y's are individually capable only of producing effect e by undergoing change, then x cannot, (without the addition of some non-supervening causal power), produce effect c without undergoing change.[12]

This principle is mistaken, which I suggest is a big metaphysical surprise. By suitably combining things that need to change in order to produce a result, we can generate a fusion that can produce that result without undergoing change. Things that operate in relatively mundane ways can be combined to generate things that act in almost magical ways. If there were special laws of magic in operation for big things in force at a world, this would not be surprising. What is surprising is that getting together enough of the mundane things and suitably arranging them is all by itself logically sufficient to entail changeless causation. Just as the contact principle broke down in certain cases where y is a fusion of infinitely many x's, I suspect that the above principle breaks down in some such cases.[13]

The Contact Principle, in full generality, could be given up fairly readily on reflection. The Change Principle has a rather deeper hold on us. It seems to us scarcely thinkable that mundane causal powers—say that of killing with a machete—could combine so as to logically entail the causal power of producing some effect without the agent of the effect undergoing change. Nevertheless, surprising as this may be, the Change Principle should be rejected. The diagnosis is complete. The logic of each case is very much in order. And our puzzlement has been traced in each case to some faulty principle relating fusions to parts. Once we discard those principles, we will have no problem in accepting the required conclusions about what happens in each case. What we have exposed along the way are some natural mistakes about the related topics of contact and causality that beset our thinking.

Appendix

The reader may sense some family resemblance between the cases I have described and the famous lamp introduced by J.F. Thompson.[14]

> There are certain reading-lamps that have a button in the base. If the lamp is off and you press the button the lamp goes on, and if the lamp is on and you press the button the lamp goes off. So if the lamp was originally off, and you pressed the button an odd number of times, the lamp is on, and if you pressed the button an even number of times the lamp is off. Suppose now that the lamp is off, and I succeed in pressing the button an infinite number of times, perhaps making one jab in one minute, another jab in the next half-minute, and so on. … After I have completed the whole infinite sequence of jabs, i.e. at the end of the two minutes, is the lamp on or off? It seems impossible to answer this question. It cannot be on, because I did not even turn it on without at once turning it off. It cannot be off, because I did in the first place turn it on, and thereafter I never turned it off without at once turning it on. But the lamp must be either on or off. This is a contradiction.[15]

The case of Thompson's Lamp has a straightforward solution, noticed by Paul Benacerraf:

> …Thompson's instructions do not cover the state of the lamp at t1, although they *do* tell us what will be its state at every instant *between* t0 and t1 (including t0). Certainly, the lamp must be on or off (provided that it hasn't gone up in a metaphysical puff of smoke in the interval), but nothing we are told implies which it is to be. The arguments to the effect that it can't be either have no bearing on the case. To suppose that they *do* is to suppose that a description of the physical state of the lamp at t1 (with respect to its being on or off) is a logical consequence of its state (with respect to the same property) at times prior to t1.[16]

The answer in short is that the description underdetermines which state the lamp will be in. If the description captures all that is causally relevant in the world, the lamp will be in one of the states: on, off, though it is causally underdetermined which it will be. If the description does not capture all that is causally relevant, we need the extra information in order to figure out the state of the lamp at t1.

The Thompson's lamp case figures an open ended series of events whose limit is a point later than each event. The assassin's case figured an open ended series of hypothetical events (grounded in actual causal powers) whose limit is a point earlier than each hypothetical event. The hypothetical nature of the events and the orientation of the open ended series separate the two cases. But why does that make any significant difference?

It is important to notice that in the Thompson lamp case, the description does NOT entail anything about the state of the lamp at t1. By contrast, in the assassin case, the description DOES entail something interesting about the state of Bob at 1 pm: In particular, that he will be assassinated. Similarly, the case in section one (which we can think of usefully as an open ended series of hypothetical wall-hittings whose limit is a time earlier than the hittings) DOES entail something interesting about the state of the sphere at 1 pm—namely that it stops proceeding.

(Of course, there may be other kinds of indeterminacy. Suppose the walls have closed surfaces that alternate between positively and negatively charged surfaces. Suppose that it is a law that balls rebound quickly from positively charged surfaces, slowly from negatively charged surfaces. How will it rebound when it hits the fusion? The laws as described underdetermine how the sphere will behave when it hits an open surface where neither positive nor negative layers dominate.[17])

The puzzlement that arose in sections one and two had everything to do with the fact that the sequence described, and the causal powers ascribed to it, *did* entail something interesting about the state of the target object at the limit point. In the Thompson's lamp case, the observation that there is no state entailed by the description of the case (nor any prohibited) was pretty much therapy enough. By contrast, the therapy in sections one and two required an account of why we were even inclined to reject the limit description that *is* logically required by the description of the case.

Notes

[1]I am grateful to Jose Benardete, David Chalmers, Tamar Gendler, Trenton Merricks, Geoff McDonough, Ted Sider, Peter Van Inwagen, Dean Zimmerman and audiences at Rutgers University and the Mighty Midwestern Metaphysics conference, 1998, for helpful discussion.

[2]*Infinity: An Essay in Metaphysics* (Clarendon: Oxford, 1964), pp. 255–259.

[3]*Infinity: An Essay in Metaphysics*, pp. 259–60

[4]I am thinking of the walls as having closed surfaces. See later discussion for an explanation of the open/closed distinction.

[5]Of course, in a world where things are sometimes brought to a halt by contact with a body that does not in turn move at all upon contact, velocity will not be continuous.

[6]To be precise, we should add 'Also, for each wall: Upon contact of a thing of which the wall is a part with a y object, the wall will remain immobile with respect to the ground.' See footnote 8.

[7]*Op. cit.*, p. 258.

[8]Of course an object may have an open surface on one side and a closed surface on another.

[9]I have discovered that Peter Van Inwagen reached a similar conclusion to what follows in some informal correspondence with Allen Hazen.

[10]The relevance of footnote 4 should be apparent: without that extra clause, it is strictly speaking consistent to say that Gordon moves when contacted.

[11]Of course it also hits Gordon minus, where Gordon minus is the fusion of all the walls minus wall 1. And so on. Causal overdetermination in abundance!

[12]By a 'non-supervening causal power', I mean a causal power that doesn't supervene on facts about the causal powers of the x's together with facts about the number and distribution of x's. Clearly we can allow that big things could act changelessly even though little things didn't thanks to some extra law that said that big things generate, say, a magnetic force.

[13]A hybrid between the wall and the assassin cases that also serves as a counterexample to the change principle: Make it a law that walls with a closed surface will turn red upon contact with a y object. Make it a law that a blue wall with a closed surface will stay blue if if not contacted by a blue object. Assume each wall is blue and has a closed surface. Each wall's individually stopping a y object will thus involve its changing to red. The fusion has an open surface and, thanks to its blue parts is blue. It will not change color when it stops a y object.

[14]'Tasks and Super-Tasks,' *Analysis* 15, October 1954, pp. 1–13.

[15]'Tasks and Super-Tasks,' p. 5.

[16]'Tasks, Super-Tasks and the Modern Eleatics,' *The Journal of Philosophy*, Volume LIX, No. 24, p. 708.

[17]See also my discussion in section one concerning whether the ball will rebound or stop.