To: Other participants in the Workshop on Inferential Problems in the Analysis of Treatment Effects, Santa Fe Institute, 26-29 July 1997

From: Glenn Shafer

ON THE PUZZLE PRESENTED BY RICHARD GILL

Early in our meeting, Richard Gill presented a puzzle about causality in physics and challenged those of us who advocate general theories of causality to show how these theories handle the puzzle. At the end of the meeting, I asked that it be put on the record that I stood ready to respond to the challenge, although there was not enough time for me to present my response. Richard suggested I communicate it by e-mail. So here it is.

During his presentation, Richard mentioned work by Mermin, Maudlin, Peres, and others. I do not have the exact references yet. I have seen some similar puzzles in the past, but as I recall they were more complicated and hence harder to discuss.

1. THE PUZZLE

In the puzzle Richard presented, a physical process simultaneously emits three particles, each traveling at the speed of light in a direction orthogonal to the other two. Each particle passes an observer stationed a certain distance away, who makes a measurement on it, the result of which can be + or -. Quantum theory allows us to make two statements about the results of the three measurements:

Statement 1. For each observer, the result of her measurement is uncertain just before it is made. Quantum theory gives a probability for + that is strictly between zero and one.

Statement 2. Someone who is told afterwards which properties of the three particles were measured and the results of two of the measurements can predict the result for the third.

The puzzle is that these two statements seem susceptible of contradictory causal interpretations, at least when coupled with the following principle of causality:

Principle of Locality. The result of a measurement cannot be affected by an action or event that does not precede it in space-time. In other words, a measurement at space-time location b cannot be affected by an action or event at space-time location a unless you can get from a to b at the speed of light or less.

Statement 1 seems to say that the result of the measurement by each observer is indeterminate just before it is made. (This impression is reinforced if we allow each observer to decide exactly what to measure only just before the measurement, but this aspect of the story does not seem essential to this particular puzzle, and so I am not emphasizing it.) In other words, the result of each measurement is not determined by what happened in its space-time past. By the principle of locality, it is also not determined by anything that happens elsewhere in space-time. In particular, it is not determined by the results of the measurements made by the other two observers.

Statement 2, on the other hand, says that the result of one observer's measurement is determined by the results of the other two oberservers' measurements. Because of the symmetry among the three observers, we may want to avoid saying that two of the results cause the third, but we do want to say that there is some sort of causal interaction among the three.

According to Richard, many physicists and philosophers now lean towards resolving the puzzle by abandoning the principle of locality.

2. NATURE'S EVENT TREE

The understanding of causality that I advocate in "The Art of Causal Conjecture" is based on the concept of idealized prediction. By this I mean the best possible prediction, categorical or probabilistic, by an imaginary observer ("Nature") who observes everything that happens and knows all regularities that might be used to predict what will happen next. I formalize this by supposing that Nature has an event tree that expresses the possibilities she forsees. A node in this tree represents a situation Nature can be in or, equivalently, an instantaneous event (the event Nature arrives in that situation). Steps between nodes are non-instantaneous events. A human action can be considered an event, usually non-instantaneous, but of course most events, whether or not they are instantaneous, have nothing to do with people.

At first glance, one might not think that idealized prediction is adequate to serve the purposes to which the idea of causality is put. My book is devoted to showing that it is. Here are some examples of how the causal relations we need can be based on the idea of idealized prediction:

1. An instantaneous event A can make another instantaneous event B necessary. This means that Nature can predict B for sure when A happens.

2. Similarly, an instantaneous event A can leave another instantaneous event B possible. This means that Nature cannot rule B out when A happens.

3. A non-instantaneous event A (a human action, for example) can influence a variable X (some number that depends on how events come out). This means that A changes Nature's probabilities for X.

4. Two variables X and Y can be related causally in a great variety of ways, because events that influence X may also influence Y, and may do so in a variety of systematic ways.

Example 3 may come closest to what we mean most often by saying that an event is a cause. But I think it would be unwise to try to make it into a technical definition of "cause," because just what kind of change in X's probabilities we will want count as a cause will depend very strongly on the context of the conversation. It is better to leave "cause" for informal usage, replacing it in formal usage by a variety of precisely defined causal relations. As I said in my talk in the conference, there are many causal relations, among many different kinds of objects (instantaneous events, non-instantaneous events, ordinary variables, and Human variables), but no single one of them is important enough to be called "cause."

Many people feel that the concept of idealized prediction somehow misses the essence of causality, even if it is an adequate foundation for all the work causality can do in biostatistics, social science, engineering, and artificial intelligence. Certainly it does eliminate the enjoyable sense of mystery in much causal talk. But I am convinced that other more mysterious-sounding ways of explaining causality reduce to idealized prediction when we strip away what is non-empirical and simply superfluous in them.

Discussion at the conference made it appear that most participants were willing to explain causality in terms of human intervention. In David Freedman's words, X is a cause of Y if by wiggling X we can change Y. When I protested that people might not be around to wiggle X, Sander Greenland said that it is enough to imagine wiggling X. What does this mean? What are the rules for what we can imagine and what we cannot imagine? If we were allowed to imagine any intervention whatsoever, then this concept of causality would have little content. The only satisfactory way I have found to discipline our imagination in this matter is to say something like this:

We are allowed to imagine intervening to change the value of X to x in situation S if and only if x is possible, given all the information that could possibly be available at that time.

And this brings us back to Nature's event tree. It says that in S even Nature considers it possible for X to become equal to x.

Another favorite formulation at the conference was "counterfactuals." This idea has the same need for discipline. Of all the counterfactuals we can imagine, which have a reasonably definite meaning, and when is this meaning causal? I would suggest that "If A had happened..." has a definite causal meaning precisely in those cases where we have specified, explicitly or implicitly, an earlier situation S where A was possible even from Nature's point of view. And this brings us back again to Nature's event tree.

3. A "NATURE" FOR EACH SPACE-TIME TRAJECTORY

Relativity theory and quantum mechanics, although they have been in the saddle for nearly three-quarters of a century, still lie at some distance from our ordinary uses of causality. So I do not emphasize them when I explain the idea of Nature's event tree. The idea of Nature's event tree is easily extended, however, to accomodate them.

One important feature of quantum mechanics is the lack of a joint distribution for non-commuting observables, such as the position and momentum of a particle. This feature is evidently accounted for as soon as we allow Nature to have expectations that fall short of a full probability distribution in her event tree, and as I argued at the meeting, this generalization from probability is also useful and needed for everyday examples. So there is really nothing special here.

Richard's puzzle is more concerned with relativity theory. Here we do need to generalize the idea of Nature's event tree. The generalization is, I think, very natural and straightforward. Instead of a single event tree, we need many, one for each trajectory in space-time that might be followed by an observer of the all-observing type that I call Nature. Each such observer always knows everything that has happened in her own space-time past (in her past Minkowski cone, if I remember the terminology correctly). This means, of course, that two such observers know the same things and hence have the same probabilities when they happen to find themselves in the same place at the same time. But otherwise they have different knowledge and hence may have different probabilities. At the end of Chapter 2 of "The Art of Causal Conjecture" there is a picture that illustrates how their event trees might intersect; I got the idea from an article by Nuel Belnap.

So far, so good. What may be more revealing is the fact that causal relations, since they are merely relations in Nature's event tree, may be different along different space-time trajectories. Which of two events comes first may be different for "Natures" following different space-time trajectories, and if each event entails the other, this means that which "causes" the other may be different for different "Natures." The paradox is softened if we agree, as I have urged, not to use "cause" as a technical term, but undeniably, as we see from the puzzle Richard related to us, the physics goes beyond our ordinary experience.

In the puzzle, a "Nature" moving through space-time who hears first about the measurement by Observer A and then the measurement by Observer B will say that their results make the result of the measurement by Observer C necessary. A "Nature" who hears about the results in a different order will say something different. No problem. Since causal relations are nothing more than relations in "Nature's" tree and tell us nothing more than what "Nature" predicts, we have no difficulty in understanding that these relations vary with which "Nature" we are talking about. Fortunately, the applications of causality in medicine, social science, and artificial intelligence can make do with only one "Nature" and hence are simpler to talk about.

4. WHAT IS THE POINT?

The event-tree theory of causality casts little new light on physics. It is to its credit, however, that puzzles from physics give it so little difficulty. These puzzles teach us, I think, that we must abandon certain mystical ideas of causation and admit that causal structure is nothing more than the structure of ideal prediction. This is no problem for the event-tree theory, because this theory has made the admission already.