5. MARKOV PROPERTIES AND QUANTUM EXPERIMENTS

Few people have thought so hard about the nature of the quantum theory as has Jeff Bub, and so it seems appropriate to offer in his honor some reflections on that theory. My topic is an old one, the consistency of our microscopic theories with our macroscopic theories, my example, the Aspect experiments (Aspect et al., 1981, 1982, 1982a; Clauser and Shimony, 1978; Duncan and Klempner, 1998) is familiar, and my simplification of it is borrowed. All that is new here is a kind of diagonalization: an argument that the fundamental principles found to be violated by the quantum theory must be assumed to be true of the experimental apparatus used in the experiments that show the violation.

The chief principle I have in mind is essential in causal inference in macroscopic problems, and is used almost without notice in experimental and observational studies in economics, epidemiology, biology, physics, everywhere. The *Causal Markov Condition (CMC)* is the following property:

Consider any system $S = \langle G, Pr \rangle$, including a set V of variables whose causal relations are represented by a directed acyclic graph G having members of V as vertices. A directed edge, $V_1 \rightarrow V_2$ in G represents the proposition that there exists a set A of values for $V \setminus \{V_1, V_2\}$ such that V_1 covaries with V_2 upon an intervention fixing $V \setminus \{V_1, V_2\}$ and randomizing V_1 . Let V be *causally sufficient*: there is no variable X not in V such that if G were expanded to include X , there would be two vertices in V with edges from X directed into them. For any variable V in V , let $\text{Par}(V)$ be the set of vertices in V that have edges directed into V , and let $\text{Des}(V)$ be the set of edges that are endpoints of directed paths from V . Let Pr be a joint probability distribution on all possible assignments of values to variables in V such that for all vertices V_1, V_2 in V , and for all such assignments of values, if V_2 is not a member of $\text{Des}(V_1)$, then V_1 is independent (in measure Pr) of V_2 conditional on $\text{Par}(V_1)$. Then S satisfies the Causal Markov Condition.

Abstract as it may be, the condition is merely a reasonably rigorous generalization of Hans Reichenbach's "(1956) screening off" conditions for causal relations. Causally sufficient, feed-forward deterministic systems satisfy the condition if their exogenous causes are independent in probability.

* Dept. of Philosophy, Carnegie Mellon University, Pittsburgh PA 1521-3890, USA and Florida Institute for Human and Machine Cognition, University of West Florida, Pensacola, FL 32052, USA; E-mail: cg09@andrew.cmu.edu

A second principle is *Faithfulness*: All conditional independence relations in a system satisfying the Causal Markov Condition are consequences of that condition applied to the graph of the system.

One way to view the experiments that demonstrate the inconsistency of quantum theory with the Bell inequalities is that they show that one or both of these conditions must fail as universal causal principles: feed-forward systems exist that cannot be made causally sufficient consistent with CMC and Faithfulness. There are many diagnoses in different terms. David Bohm, Bub's teacher, would perhaps have said that that is because no system is causally sufficient; other commentators might locate the problem with the assumption of a joint probability distribution, and so on. I wish merely to point to the curiously valid, almost Wittgensteinian logic, that gets us to the inconsistency.

Instances of assumptions of the CMC and of Faithfulness could be traced through the details of the experimental set up, runs and data analyses of the Aspect experiments. But it has been a long time since I was any kind of physicist, and I would inevitably misrepresent details and confuse even the readers of clearest mind, and there are details of sensor behavior and sensitivity that complicate without clarifying. So I will pass on the details and consider instead a very simple idealization of the phenomenon, due to N. David Mermin (1985, 1990).

Consider two detectors I and II that are spatially separated. Each detector has three settings, $S = 1, 2$ or 3 . Further each detector has a red bulb R and a green bulb G. Pairs of particles are emitted from a source and enter the two detectors. There is no other physical connection of any kind we know of between the detectors (Figure 5.1).

The detectors behave this way: (1) when both detectors are set to same value, no matter which, they both show red or they both show green. Red and green occur with equal frequency; (2) when the two detectors are set to any two *different* values, they show the same color, both red or both green, $1/4$ of the time—again, red and green occur with equal frequency in this case, and different colors $3/4$ of the time—each

FIGURE 5.1.

Table 5.1

Left indicator setting	Left indicator light color	Right indicator setting	Right indicator light color	Probability of the two light colors given the settings
1	Red	1	Red	1
1	Green	1	Green	1
2	Red	2	Red	1
2	Green	2	Green	1
3	Red	3	Red	1
3	Green	3	Green	1
1	Red	2	Red	1/8
1	Green	2	Green	1/8
1	Red	2	Green	3/8
1	Green	2	Red	3/8
2	Red	1	Red	1/8
2	Green	1	Green	1/8
2	Red	1	Green	3/8
2	Green	1	Red	3/8
1	Red	3	Red	1/8
1	Green	3	Green	1/8
1	Red	3	Green	3/8
1	Green	3	Red	3/8
3	Red	1	Red	1/8
3	Green	1	Green	1/8
3	Red	1	Green	3/8
3	Green	1	Red	3/8
2	Red	3	Red	1/8
2	Green	3	Green	1/8
2	Red	3	Green	3/8
2	Green	3	Red	3/8
3	Red	2	Red	1/8
3	Green	2	Green	1/8
3	Red	2	Green	3/8
3	Green	2	Red	3/8

combination of colors (I green, II red; I red, II green) equally often. We can show the whole story about the probabilities with a tedious but clear table (Table 5.1).

The thing to notice immediately is that, no matter how we set the two detectors, the colors the detectors show will not be independent in probability. If both detectors are set at the same value, the probability that Detector II is red is 1 conditional on Detector I being red, and vice versa. If both detectors are set at different values, the probability that Detector II is green given that Detector I is red is three times the probability, on that same condition, that Detector II is red. Notice further, that someone at Detector I cannot use his settings of the detector to send signals or communications to someone at Detector II via the color that shows up at Detector II. For despite the fact that no matter how the detectors are set, the colors are correlated, the color at Detector II is independent in probability of the setting at Detector I.

Table 5.2

State	1,2	2,1	1,3	3,1	2,3	3,2
RRR	Same	Same	Same	Same	Same	Same
RRG	Same	Same	Differ	Differ	Differ	Differ
RGR	Differ	Differ	Same	Same	Differ	Differ
GRR	Differ	Differ	Differ	Differ	Same	Same
RGG	Differ	Differ	Differ	Differ	Same	Same
GRG	Differ	Differ	Same	Same	Differ	Differ
GGR	Same	Same	Differ	Differ	Differ	Differ
GGG	Same	Same	Same	Same	Same	Same

Mermin puts the problem this way. The only explanation (he says) for the first six rows of the probability table is that the particles each have internal states that specify their response to each state of a detector. The internal states of each particle specify what color it will activate for each of the three settings of the detector. Since there are 2 possible colors for each detector setting, and three settings, there are 8 possible internal states for each particle. If and only if (Mermin says) both particles have the same internal states will the colors of the two detectors agree when they have the same setting, for all 3 possible settings. So the states of the particles have to be perfectly correlated, the same. If one particle will make a detector go red on setting 1, red on setting 2, and green on setting 3, so will the other. So the question becomes: *is there a probability distribution over these possible internal states of the two particles that, consistent with their perfect correlation, agrees with probability table?* There is not. In particular, there is no way to assign probabilities to the particle states so that when the settings of the detectors are different, the detector colors agree less than 1/3 of the time. Let's do another table (Table 5.2). The columns indicate the settings of the two detectors when they are different, and the entries indicate for each state and pair of settings whether the colors of the detectors are the same or different.

In each row the fraction of cases in which the colors are the same is 1/3 or more. No matter what the relative frequency of the various particle states may be, if the detectors are set at any pair of distinct settings, the colors must be the same at least 1/3 of the time, but in the data for the experiment, for such settings the colors are the same only 1/4 of the time.

So what does this have to do with Markov Assumption and so forth? Two things. On the one hand, the conclusion of the example, while not inconsistent with the Markov Assumption, is inconsistent with the conjunction of the Markov Assumption and the claim that the state of the particle is the only causal connection between the detectors. On the other hand, while Mermin's reasoning is perfectly correct, his argument depends on using the Markov Assumption. I will represent Mermin's account of his experiment as a causal graph, like this (Figure 5.2).

The causal diagram and the Markov Assumption explain why the setting of Detector I cannot be used to send a signal to Detector II via the color that appears at Detector II—there is no causal pathway from Setting of Detector I to Color for

FIGURE 5.2.

Detector II, or vice-versa, so the two variables must be independent. And the causal diagram explains why the colors at the two detectors are correlated: they have a common cause. Nonetheless, there is something very wrong. There is no causal pathway from Color for Detector I to Color for Detector II, or in the other direction. There is no common cause of detector colors other than Particle State. Since Color for Detector II is not an effect of Color for Detector I, and vice versa, the Markov Assumption says they if the causal graph above is correct, the detector colors should be independent of one another *conditional* on Particle State. Indeed, that is exactly what Mermin's particle states do imply. For example, given that the particle state is RRR, then Detector I is red and Detector II is red: no matter the settings and neither detector provides any information about the other detector not already entailed by the particle state. If the particle state is RGR, then no matter how Detector I is set, the color in Detector I gives no further information about color that will appear at Detector II. (The setting chosen for Detector II provides further information about the color that will show up for Detector II when the particle is in the RGR state, but that is beside the point.) But Mermin's argument shows that these particle states cannot be made consistent with the assumed observed frequencies of colors in each combination of settings shown in Table 5.1. So there are logically just three alternatives (1) Mermin has sneaked in some extra assumption somewhere, or (2) the Markov Assumption is false for this case, or (3) there is no causal explanation of the correlations of the detector colors. Perhaps more than one of these alternatives is true.

Mermin has certainly sneaked in some assumptions—all of them instances of the Markov Assumption—and the fact that he does not make them explicit may indicate that the Markov Assumption is so fundamental to our reasoning about experiments that we use it automatically, without notice. For *there is* a common cause explanation of the probabilities in Table 5.1. Here is the idea, first noted by Suppes and Zanotti (1981) in a more general case: Change the particle states so that they no longer just specify a color for each of the three settings of a detector. Now they specify a color

FIGURE 5.3.

for each detector setting *and a setting for each detector*. Instead of 8 internal states of the particle, we now have 48 internal states of the particle. The particle state now uniquely determines the color at each detector. Given the (new) particle state, the color at either detector provides no further information about the color at the other detector, because there is no more information to provide. We can give another causal diagram (Figure 5.3).

The Markov Assumption is satisfied. (Alternatively, the particle states can influence the interventions, which influence the detector settings.) Why doesn't Mirmin allow this? Because he thinks, quite reasonably, that the particle states do not cause the detector settings. Why not? Because he thinks the human act of setting the detectors (or a machine act of randomly setting the detectors) is an *intervention*, a cause that is not influenced by any feature of the system and that fixes the value of the Detector setting while leaving all of the conditional probabilities of other variables unchanged. (Similar reasoning applies to the idea that the detector settings influence the particle state.)

Ok, take out the causal influence of the particle states on the detector settings, but leave the 48 states of the particle and their probabilities just as before:

Now we can still account for the correlations in Table 5.1, and the particle state is still a common cause of the detector colors, condition on which the detector colors are independent—the Markov Assumption is satisfied. Why doesn't Mirmin allow that? Because the causal diagram in Figure 5.4 and the probabilities assumed for the particle states are jointly inconsistent with the Markov assumption in another way—each detector setting is dependent in probability on the particle state (and vice-versa), but there is no causal pathway or common cause relating the detector setting variables to the particle state. Supposing there is another common cause beside the particle state that also influences the colors won't help things—the same argument goes through, it's just more complicated. However, we do things, we do not have a causal explanation of the experiment consistent all the way through with the Markov assumption.

FIGURE 5.4.

Mirmin—and we—reason about his imaginary experiment using the Markov Assumption and the notion of an intervention, and yet the experiment allows of no causal explanation consistent with the Markov Assumption. The example is a simplification of what goes on in real experiments to test remote correlations predicted by a consequence of the quantum theory, Bell's theorem. In quantum experiments, we pull ourselves *down* by our bootstraps.

Now there is an obvious solution to the problem: the color at one or both of the detectors influences the color at the other detector.

This is a popular solution, and the reason why the problem is often said to be about "locality" or the phenomenon is said to exhibit "non-locality." Often the non-locality solution is implicitly motivated by the idea that the correlations between the colors must have a causal explanation.

Since the detectors can be far enough apart, and the color measurements close enough in time that the theory of relativity prohibits a signal from being sent from one detector to another, the solution has a problem. The problem is this: Suppose before the experiment, the guy at Detector II tells the guy at Detector I how Detector II will be set. Then, if the causal story above is correct, by adjusting the settings of Detector I the first guy can send signals to the second guy, who will figure them out from the color that shows up at Detector II. It works this way. There is in Figure 5.5 a causal pathway from setting of Detector I to the color at Detector II. The pathway must create an association between the two, and associations are all that is needed for communication, for sending a signal. The Faithfulness assumption says a direct causal connection creates an association—and the very point of the non-locality hypothesis is to create such an association between the colors. (Consistently with the Markov Assumption the association cannot be the effect of a common cause—for reasons we have already reviewed.) The setting of Detector I influences the color at Detector I, so we have a sequence of causal links—and correlations or associations—between Detector I and the color at Detector II. Now, a causal linkage of one variable with

FIGURE 5.5.

a second linked with a third need not always create an association between the two variables, even if it is the only pathway connecting the variables (as in this case between Detector I and the color at Detector II). For example, suppose variable A has three values and variable B has three values (say, b1, b2 and b3) and variable C has two values, and the probabilities for two of the values (b1 and b2) of B depend on the value of A, (but the third value, b3, of B does not depend on the value of A) and the probability of values of C depends on whether B has value b3 or one of the values b1, b2, but doesn't depend on which of the values b1 or b2 B has. Then interventions that vary A will not create any association with C. Despite the fact that A influences B, and B influences C, A does not influence C: causation is not transitive. *But if B has only two values, the causal relations must be transitive, and A must be associated with C.* That is exactly the situation in the Mermin's thought experiment. Hence relativity can be violated. Having the influence go in both ways doesn't help; the argument still works.

The argument doesn't depend on any philosophical niceties about what "causation" means, and it doesn't depend on any details of the physics. It depends on the assumption that the settings of the Detectors are interventions, and the hypothesis that the "non-locality" relation creates an influence between the colors. So, if relativity is true and the statistics drawn from the Aspect and similar experiments are sound, causal non-locality is a non-starter.

The upshot is this: real experiments with associations analogous to those of Mermin's thought experiment create associations that have no causal explanation consistent with the Markov Assumption, and the Markov assumption must be applied, implicitly or explicitly, to obtain that conclusion. You can say that there is no causal explanation of the phenomenon, or that there is a causal explanation but it doesn't satisfy the Markov Assumption. I have no trouble with either alternative. It is not a truth of logic that all experimental associations have a causal explanation, and it is

not a truth of logic that all causal relations satisfy the Markov Assumption. That's up to Nature. But I do have this problem: *why, then, does the Markov Assumption work with our experiments on middle sized dry and wet goods, with climate, and rats and drugs, and so much else?*

I have no definite answer. I would suggest looking in these banal directions. First, among properties of middle sized objects, Aspect-like associations are extremely small, so the properties of systems are nearly deterministically related, or would be if all significant causes of variation were accounted for; second, when system are not causally sufficient, we make them nearly so when we can by redefining variables, by conditioning on variables with unexplained associations, and other devices; third, insofar as macroscopic frequencies are generated as "strike ratios" from deterministic processes, as proposed long ago by Hans Reichenbach in his doctoral thesis and more recently by Michael Strevens (2003), we should expect the Causal Markov Condition to hold necessarily. And finally, there are proofs that under continuous measures on the parameters of various families of probability distributions, the Markov Condition implies that the Faithfulness condition holds almost always (Spirtes et al., 2000).

REFERENCES

- A. Aspect, P. Granger and G. Roger, *Physical Review Letters* 47, 460, 1981.
- A. Aspect, P. G. Granger and G. Roger, *Physical Review Letters*, 49, 91, 1982.
- A. Aspect, J. Daligard, and G. Roger, *Physical Review Letters*, 49, 1804, 1982a.
- J. F. Clauser and A. Shimony, *Reports on Progress in Physics*, 41, 1881, 1978.
- A. J. Duncan and H. Kleinpoppen, in *Quantum Mechanics versus Local Realism*, F. Selleri (ed.), Plenum, New York, 1998.
- N. D. Mermin, *Physics Today* 38, 38-47, 1985.
- N. D. Mermin, *Boojums All the Way Through*, Cambridge, Cambridge University Press, 1990.
- P. Spirtes, C. Glymour, and R. Scheines, *Causation, Prediction and Search*, MIT Press, Cambridge, MA., 2000.
- M. Strevens, *Bigger Than Chaos*, Cambridge, MA Harvard University Press, 2003.
- P. Suppes and M. Zanotti, When are probabilistic explanations possible?, *Synthese* 48, 191, 1981.