

The Common Cause Principle

Frank Arntzenius

University of Southern California

1. Introduction

The common cause principle, roughly speaking, consists of the following 2 sub-principles:

i) If 2 events (or types of events, or facts, or conditions, or) are correlated, and the one does not cause the other, then there is a third event (type of event,) such that the 2 events are probabilistically independent given the presence or absence of the third event. That is to say: for every pair of correlated events that do not have direct causal links there is a screener off of that correlation.

ii) This screener off occurs before the correlated events.

Thus, assuming that the screener off is the cause rather than the effect, the common cause principle tells one that the cause occurs before the effects.

In this paper I will argue that the common cause principle is false. I will begin by giving some well-known objections to the principle, suggest some new objections, and then prove a more restricted version of the common cause principle.

2. Bell Experiments

It has been shown that the correlations between the results of spin measurements on pairs of electrons, prepared in a particular way, can not have a screener off, given some very plausible assumptions. If one furthermore assumes that these correlations could not be due to direct non-local causation, then it follows that for certain peculiarly quantum mechanical phenomena the common cause principle fails. See, for instance, Van Fraassen (1982). However it is contentious whether there could not direct non-local causation in this case.

3. Indeterministic Decay with Conservation of Momentum

Suppose that a particle decays into 2 parts, that conservation of total momentum obtains, and that it is not determined by the prior state of the particle what the momentum

of each part will be after the decay. By conservation, the momentum of one part will be determined by the momentum of the other part. By indeterminism, the prior state of the particle will not determine what the momenta of each part will be after the decay. Thus there is no prior screener off. This example is from van Fraassen (1980), page 29.

4. Similar Laws of Evolution

The bread prices in Britain have been going up steadily over the last few centuries. The water levels in Venice have been going up steadily over the last few centuries. There is therefore a correlation between bread prices in Britain and sea levels in Venice. However, there is presumably no direct causation involved, nor a prior screener off available, at least not one that we would be inclined to call a common cause. For, although it is plausible that given both the bread price and the water level at some time the bread price and water levels at any later time are statistically independent (given some intuitive guess regarding the physical chances involved in his example), such a specification of bread price and water level at a given time would not intuitively be considered a common cause. This example is from Sober (1988).

Note, however, that in the formulation of the common cause principle that I gave, I only demanded a prior screener off, and said nothing about how 'common' it should be. Thus at least some varieties of the common cause principle are not clearly violated by this example.

5. Deterministic Worlds

In this section I will argue that the temporal asymmetry of the common cause principle can not hold in a deterministic world, if one is liberal about what one counts as events (types of events, facts, conditions, ...). By a deterministic world I mean a world in which the complete state of the world at any given time determines the complete state of the world at any other time.

I take the following claim about events to be unproblematic: whether an event occurs at a given time is determined by the complete state of the world at that time. Thus I take it that for every event there is unique set of states such that the event occurs if and only if the state of the world is in that set. If one were liberal about what one counts as events one might also adhere to the converse: to every set of states there corresponds a unique event, namely the event that occurs at some time if and only if the state of the world is in that set at that time. Of course many such events will not be natural or simple events. However, *prima facie*, one would hope that the validity of the common cause principle does not depend on the claim that certain sets of states are too unnatural to count as the entities to which the common cause principle is supposed to apply. We will return to this issue in the next section, so let us for now accept a liberal notion of events.

Determinism entails that for any state at any time there is a unique state into which that state develops at any other time. Thus for every set of states at any time there is a unique set of states at any other time. Thus for any event *A* at any time *t* there is at any other time *t'* a unique event *A'* that is determined to occur at *t'* if and only if *A* occurs at *t*.

Events that determine each other will have conditional probability 1 upon each other, and hence have the same probabilistic relations with all other events. It therefore follows from determinism that for any screener off that occurs before some pair of correlated events, there is a screener off that occurs after the correlated events,

namely the event that occurs at that later time, iff the screener off occurs at the earlier times. Thus in deterministic worlds there is no asymmetry regarding screeners off, if one allows any set of states to count as a potential screener off. I have argued this at more length in Arntzenius (1990).

Of course the event that occurs if and only if the screening off event occurs at the earlier time will in general be a very complex and unnatural event, it will be some weird condition on the distribution of elementary particles throughout some large part of the universe. But, by determinism, we know that there is some such condition which occurs after the correlated events and screens them off, if there is a prior screener off.

Note that we can also be sure, given determinism and a liberal notion of events, that there will be a prior screener-off, at least if we allow partitions, rather than properties, to count as screeners off. For consider the case of 2 correlated events. There are 4 possible combinations of 2 those events occurring or not occurring. This corresponds to a partitioning of the statespace into 4 cells (4 sets of states). By determinism for every partition $\{C_i\}$ and any times t and t' , there is a partition $\{C'_i\}$ such that the state of the world is in cell C'_i at t' if and only if it is in cell C_i at t . Thus for any time t' there will be a partition with 4 cells which determines which combination of the correlated properties occurs. But, trivially, the properties are uncorrelated on each cell of such a partition, and thus such a partition is a screener off. Of course, there is no guarantee that one could, at any time, find a property (a 2-celled partition) which screens off the correlated properties. But why would one expect this in the first place? To demand that the screener off be a 2-celled partition, rather than a partition in general, would seem to be a strange demand for which I can see no justification.

6. Macroscopic Events

Given the above problem, one might conjecture that the common cause principle holds when one restricts oneself to some natural class of events. For every screener off at an earlier time there is a screener off at a later time, but that later screener off is bound to be some horrendously unnatural event. Thus the common cause asymmetry might hold for natural events. However, as I argued in Arntzenius (1990), the following example shows that the common cause principle will not hold in the class of macroscopic events, or in the class of directly observable events.

Cleopatra is throwing a big party, and wants to sacrifice around fifty slaves to appease the gods. She is having a hard time convincing the slaves that this is a good idea, and decides that she ought to give them a chance at least. She has obtained a very strong poison, so strong that one molecule of it will kill a person. She puts one molecule of the poison in each of a hundred goblets of wine, which she presents to one hundred slaves. Having let the molecules of poison move around in Brownian motion for a while she then orders the slaves to drink half a goblet of wine each. Let us now assume that if one consumes the poison, then death is preceded by an ominous reddening of the left hand and of the right hand. Then, the molecule being in the consumed half of the wine glass will be a prior screener off of the correlation between left hand reddening and right hand reddening. Assuming that death occurs exactly in the cases that the poison is swallowed, death will be a posterior screener off. If one restricts oneself to macroscopic events, there will only be a posterior screener off. If death is not strictly determined by the swallowing or non-swallowing of the poison, there will be no macroscopic screener off at any time. Thus, if microscopic events can have such macroscopic consequences, the common cause principle can not hold of macroscopic events.

This argument more generally suggests that the common cause principle can not hold of a class of events that has causes outside that class. This argument appears even more forceful for those who, like myself, believe that the only reason that we can acquire knowledge of microscopic events and microscopic laws, is precisely the fact that microscopic events, in certain situations, have effects upon observable events. But, if that is so, it seems that the only class of events for which the common cause principle could possibly hold is the entire class of microscopic events. But if the laws of microscopic physics are deterministic we are back at the previous problem.

7. Statespace Correlations

Suppose that someone gives you a list of pairs of observed values for 2 observables, each with possible values 1 and 0. Suppose that you notice in the list that if one observable has value 1, the other observable always has value 0, and if one observable has value 0 the other observable always has value 1. This constitutes a perfect correlation, and the common cause principle demands a common cause. However, suppose that you are then told that the observed system was a particle bouncing around in a box, and that the observables were 'presence of particle in left hand side of box', and 'presence of particle in right hand side of box'. In such a case one would not seek a common cause of the correlation, but agree that the correlations need no explanation other than a reference to the fact that any possible state in which the one observable has one value the other observable has the other value. Even if one had absolutely no idea as to what the earlier cause of the later position of the particle is, say because its development is indeterministic or chaotic, there would be no mystery regarding the correlation. Correlations that are entailed by the set of possible states that make the value attributions true, do not demand a prior screener off.

8. Equilibrium Correlations

Suppose that someone gives you a list of values for 2 observables with 2 possible values, and that they are correlated, though this time not perfectly correlated. Because the correlations are not perfect, and indeed all 4 of the possible value combinations occur in the list, the correlations can not be entailed by the truth sets of the values of the observables. Thus, even if we adjust for the previous counterexample, the common cause principle demands a prior screener off. Suppose you are now told, that, as above, the observed system was a particle bouncing around in a box, and that the observables were 'presence of particle in region A', 'presence in region B'. Suppose moreover that you are told that region A, and region B, though not completely overlapping, have nearly all area in common. Indeed suppose that the correlation that was observed is exactly the correlation one obtains for presence in these regions if one assumes that the particle has a uniform probability of being anywhere in the box, i.e. that its probability density is constant over the volume of the box. Then, one would presumably agree that the explanation of the correlation is the large amount of volume that the regions have in common, rather than some prior cause which makes each region more likely and which screens off presence in the one region off from presence in the other region. Thus equilibrium correlations, correlations that obtain when the systems in question are in (thermodynamic) equilibrium, do not demand a common cause explanation.

It should also be noted that the occurrence of equilibrium correlations is not rare, and often presupposed in common cause (and other) explanations. Consider the following example of a common cause explanation. (I think this example is due to Wesley Salmon, but I am not sure.) There is a correlation between the take-off time of airplanes, and the time clothes take to dry on nearby washing lines. The common cause explanation is that high humidity causes both long dry times and long take-off

times. However, note that this explanation presupposes that the humidity at the airport and at nearby houses is correlated. The humidity in the one area does not cause the humidity in the other area. Moreover, there is no apparent common cause of this correlation, at least none that do not beg the question by assuming other correlations. The explanation of this correlation is presumably that, in (approximate) equilibrium humidity in different areas is (approximately) identical. Indeed the world is full of approximate equilibrium correlations, most of which we demand no common causes for.

I have recently found that Malcolm Forster has already given the gist of the objection to the common cause principle described in this section. See Forster (1986).

9. Indeterministic Worlds

If the common cause principle can not hold for deterministic world, it is to be hoped that it can hold for indeterministic worlds. The most natural form of indeterminism is time-homogeneous Markovian indeterminism: the current state of the world determines probabilities for future states of the world, these probabilities are time-independent, and future states of the world are probabilistically independent of past states of the world given the state of the world at any intermediate time. Let me give an example which indicates that in such worlds the common cause principle typically will not hold.

Suppose a particular type of object has 4 possible states: S_1 , S_2 , S_3 and S_4 . Suppose that if such an object is in state S_i at time t , and is not interfered with (isolated), then at time $t+1$ it has probability $1/2$ of being in the same state S_i , and probability $1/2$ of being in state S_{i+1} , where we define $4+1=1$ (i.e. '+' represents addition mod 4). Now suppose we put many such objects in state S_1 at time $t=0$. Then at time $t=1$ approximately half of the systems will be in state S_1 , and approximately half will be in state S_2 . Let us define property A to be the property that obtains precisely when the system is either in state S_2 or in state S_3 , and let us define property B to be the property that obtains precisely when the system is either in state S_2 or in state S_4 . At time $t=1$ half of the systems are in state S_1 , and therefore have neither property A nor property B, and the other half are in state S_2 , so that they have both property A and property B. Thus A and B are perfectly correlated at $t=1$. These correlations are clearly not statespace correlations. Moreover, these correlations are not equilibrium correlations. For the equilibrium distribution is equal probability ($1/4$) for each state, and in such a distribution the properties are uncorrelated. Thus the common cause principle demands a prior event conditional upon which the properties are independent. However, there is no such screener off.

For, the full explanation of the correlation is the prior occurrence of S_1 . It entails the observed distribution given the laws of nature, and no less information will do so. However, in so far as it makes sense to talk of being a cause of events in this context, S_1 at $t=0$ is neither a cause of A at $t=1$, nor of B at $t=1$. For, A and B each have probability $1/2$ in case S_1 obtains with probability 1 at $t=0$, each have probability $1/2$ in case S_1 does not obtain at $t=0$ and S_2 , S_3 and S_4 obtain with equal probability at $t=0$, and each have probability $1/2$ in the equilibrium state. Thus S_1 does not appear to raise the probability of A, nor that of B, in any relevant sense. More importantly, S_1 does not screen property A off from property B, since $P(A \text{ at } t=1/S_1 \text{ at } t=0) \cdot P(B \text{ at } t=1/S_1 \text{ at } t=0) = 1/2 \cdot 1/2 = 1/4$, and $P(A \& B \text{ at } t=1/S_1 \text{ at } t=0) = 1/2$. Thus there appears to be no sense in which the correlation of properties A and B is caused, or explained, by a common cause. Note also that such a correlation between A and B will occur not only in the special case in which $P(S_1)=1$ initially. Indeed a correlation between A and B will occur following 'most' initial probabilities.

This example can be generalized. In order to have correlated properties we need at least 2 logically independent properties and hence at least 4 states. So let us restrict attention to time homogeneous Markov processes with at least 4 distinct states. Let us call such a Markov process normal if it has at least one equilibrium distribution (a distribution that develops into itself according to the transition probabilities of the Markov process), and there is at least one state S such that state does not immediately develop into an equilibrium distribution. One can show that for any normal Markov process there is at least one prior state S and a set of properties, such that if initially $p(S)=1$, subsequently there will be non-equilibrium correlations between those properties which are not screened off by S . For a proof, see the appendix. Thus, the common cause principle fails for any normal Markov process.

10. Macroscopic Events Again

Let me use an example, similar to the example of the previous section, to show that correlations can occur at a macroscopic level without a common cause at a macroscopic level, because of equilibrium correlations between the microscopic facts which, given the macroscopic state, determine the future macroscopic state. Thus I will again argue that the common cause principle can not hold at the macroscopic, or observable, level. This time it will not be because there is a microscopic common cause as in the example of section 6, but because the microscopic facts, which in addition to the macroscopic facts, determine the later macroscopic properties, are 'equilibrium correlated'.

Suppose we have water in a teacup that we are constantly stirring in the same direction, with a small ball floating on the surface of the water. Let us divide the surface of the water, which has the shape of a circle, into 4 quadrants 1, 2, 3 and 4. Let us call presence of the ball in quadrant i , state S_i . The flow in the teacup is roughly circular, but because of the stirring mechanism it is quite turbulent, so that the motion of the ball is very irregular. Consider now the motions of such a ball in 1 second in many observed cases. Suppose that in about half the cases such a ball crosses the boundary between 2 quadrants in the forwards direction, in about half the cases crosses no boundaries, that the ball very rarely crosses a boundary in the backwards direction, and very rarely crosses more than 1 boundary in the forwards direction. We then have roughly the transition probabilities which I gave in the previous section. Thus if we start a large number of balls in quadrant 1, then 1 second later there will be a correlation between being in S_2 or S_3 , and being in S_2 or S_4 . At the macroscopic level (states S_j), we have a Markovian indeterministic process, and there is no prior screener off of the correlation.

At the microscopic level, if we assume classical mechanics, the microscopic state of the molecules in the fluid and the ball, plus all the microscopic influences on those molecules, in conjunction with S_1 , will determine the future position of the ball. But, of course, the experiment tells us that, although there are microscopic antecedents which, in conjunction with S_1 lead to all 4 possible subsequent states, most of them lead to S_1 or S_2 , and hence to -A&-B or A&B. Thus the microscopic factors which, in conjunction with S_1 , determine whether A subsequently occurs and whether B subsequently occur, are correlated. Somewhat stretching the usage of the notion of equilibrium, one could say that they are equilibrium correlated.

If one thinks my example is not to be trusted because A and B are funny properties, then one can easily construct an example where A is the property of being in the top half of a box, and B is the property of being in the left hand side of the box, such that for a given initial macroscopic state, which does not screen off the later correlation between A and B, the microscopic antecedents of A and B are correlated.

One can show more something more general. Suppose that in a particular cell in statespace the microscopic antecedents of properties A and B are uncorrelated. Define property C as $(A \& B) \vee (-A \& -B)$. Simple algebra then shows that the microscopic antecedents of A and of C must be correlated in that cell, unless $p(A)=0$ or $p(B)=1/2$. Microscopic chaos does not and can not mean, that all microscopic properties of a system are uncorrelated. Thus any attempt to prove the common cause principle by assuming that all microscopic properties, or all properties not represented in a particular coarse-grained statespace, or all variables in addition to some given ones, are uncorrelated, rests on inconsistent premises. For instance, P. Horwich in Horwich (1987), page 74, makes such an assumption, while D. Papineau in Papineau (1985), and Peter Spirtes, Clark Glymour and Richard Scheines, in Spirtes et al (1993), chapter 3, come perilously close to making such an assumption.

11. Order Out of Chaos

When one lowers the temperature of certain materials, the spins of all the atoms of the material will line up in the same direction. Pick any two atoms in this structure. Their spins will be correlated. However, it is not the case that the one spin orientation caused the other spin orientation. Nor is there a simple common cause of each orientation of each spin. The lowering of the temperature determines that the orientations will be correlated, but not the direction in which they will line up. Indeed, typically, what determines the direction of alignment, in the absence of an external magnetic field, is a very complicated fact about the total microscopic prior state of the material and the microscopic influences upon the material. Thus other than virtually the complete microscopic state of the material and its environment there is no screener off of the correlation between the spin alignments.

Consider a fluid in a box in thermodynamic equilibrium. The directions of the motions of nearby molecules will be uncorrelated. Now heat one end of the box and cool the other end. In many cases convection currents will appear. In such cases the directions of the motions of nearby molecules will be highly correlated. But it is not the case that the motion of the one molecule causes the motion of nearby molecules. Nor will there be a simple common cause that screens the motions of nearby molecules off from each other. The external constraints determine that such correlations will occur, but they do not determine what the correlated motions will be. Thus the external constraints do not screen of the motions of nearby molecules off from each other. Indeed presumably nothing less than virtually the entire prior microscopic state of the fluid and the microscopic influences upon it will screen the motions of nearby molecules off from each other.

In general when chaotic developments result in ordered states there will be final correlations which have no prior screener off, other than the virtually the full microscopic state of the system and its environment. In such cases the only screener off will be a horrendously complex microscopic fact. But we have seen that allowing such screeners off, also poses a problem for the common cause principle, in that one will then also have posterior screeners off for deterministic systems.

One might think that these examples are more examples of equilibrium correlations. However, in a strict thermodynamic sense of equilibrium, such correlations will often not be equilibrium correlations. The convection current system is obviously not in thermodynamic equilibrium, since there will be a temperature gradient in the fluid, which is maintained by the external constraints. Indeed I. Prigogine, see Prigogine (1980), has argued that typically, such order will arise out of chaos, in systems which are in thermodynamic disequilibrium. Whether there is some other sense of equilibrium, according to which these correlations are equilibrium correlations is an interesting question. But

note that such a wider notion of 'equilibrium correlations' would only be useful with respect to an attempt to save the common cause principle, if it did not include the correlations for which we do want to demand the existence of a prior common cause.

12. A Restricted Common Cause Principle

Despite all of these objections, there must be something right about the common cause principle. For we use it all the time in everyday inferences, and it appears to be an indispensable part of the best method for inferring causal structure from statistical data in the social sciences (see Spirtes et al (1993)). Let me therefore indicate a particular restricted version of the common cause principle which is provably true. I do not claim that this is all that is true about the common cause principle. Indeed I hope that this is not so, but I do hope to show that at least in certain circumstances one version of the common cause principle is provably true.

Consider a large number of pairs of systems $\langle A, B \rangle$ at time t that have developed according to deterministic laws from time 0. Assume that the 2 members of each pair have been isolated from time 0 to time t , but that both members of each such pair have interacted with a third system C at some time between time 0 and time t . Let us furthermore assume, that at time 0 the states of the members of the pair are statistically independent. One can then prove that there must be a partition of the statespace of systems C , such that conditional upon the cells of that partition any property of systems A must be probabilistically independent from any property of systems B . For a proof see the appendix.

Moreover one can also show that if, for a given initial distribution of states of A and B , the probabilities of certain properties of A and B vary conditional upon these cells of C , then for 'almost all' initial distributions over these cells of C a correlation of these properties will ensue. Elliott Sober has pointed out to me that he and Martin Barrett have proved a theorem similar to the one of this section, namely 'Theorem 2' in Sober & Barrett (1992).

Let me state the above less precisely but more clearly: given the assumptions, one can prove that any correlation between properties of the 2 systems will be screened off by a prior 'property' (really: partition) of the 3rd system. Moreover, if the 3rd system has any effect upon the probability of each of a pair of properties of the 2 systems, then in almost all circumstances a correlation will ensue. Even more simply: any correlation of properties of the 2 systems has a prior common cause in the third system, and any such common cause will almost always produce a correlation. The common cause principle is provably valid in the specified circumstances, for the specified sets of properties.

Note that, since we are assuming determinism, there will of course also be a later screener off. But, in general this will be a property of all 3 systems, rather than system C alone. Thus, in this theorem something is said about the extent to which the prior screener off will be common, or unified, and the posterior one will not be common or unified. The prior screener off will be a property of system C , whereas, in general, the posterior one will not, and indeed in general will spread out over any systems with which systems A , B and C interact after their own interaction.

13. Conclusions

Question 1. Does any correlation have a prior screener off?

Answer. Yes, if one assumes determinism, and allows any weird partition to count as a potential screener off. But then there will also be a posterior screener off. There is no guarantee that in any particular more coarse grained statespace one will find a screener off. Moreover, there is plenty of guarantee, as outlined in the paper, that in any coarse grained statespace there will be correlations for which there is no prior screener off. However, if one knows that given some partition, the additional factors, which determine whether each of the correlated properties occurs, are statistically independent, then one knows that such a partition is a screener off. Experience and statistical mechanics may help us in guessing what such a partition might be. In particular, as the previous section indicates, knowledge that the states of certain systems are initially independent can lead to useful versions of the common cause principle.

Question 2. Suppose we a prior common cause of properties A and B in the following sense: $p(B/C)$ is unequal to $p(B/\text{not-}C)$, and $P(A/C)$ is unequal to $P(A/\text{not-}C)$, where C occurs before A and B. Should we then expect a later correlation between A and B?

Answer. Not necessarily. It depends whether the factors which in addition to C determine whether A and B occur are correlated or not. Even when it concerns initial states, and we have some variety of initial chaos, many such properties will be correlated. Experience and statistical mechanics may help us in guessing which ones will be correlated initially, and which ones will not.

Appendix

I. Proof that normal Markov processes fail the common cause principle

The proof I give below is for discrete time, discrete statespace Markov processes. I expect that an analogous proof can be given for continuous time and/or continuous statespace Markov processes.

Consider a discrete statespace with at least 4 states S_i with a probability distribution p over it. Let states S_1 and S_2 have probabilities p_1 and p_2 according to p , and let p_1/p_2 be their ratio according to p . For any ordered pair of states $\langle S_i, S_j \rangle$ such a ratio will be defined unless $p_j=0$. Of course, the set of all the probabilities will determine all the ratios, but it is also the case that the set of all ratios will fix all the probabilities, given that the probabilities have to sum to 1. For start by conjecturing some arbitrary probability for some state that has non-zero ratios with respect to the other states. The ratios between this state and the others will fix all the other probabilities (undefined ratios mean that the other state has probability=0). The end result, however, may not sum to 1. So now rescale so as to sum to 1. Then you will have determined the probabilities from the ratios. Thus, if 2 probability distributions agree on all the ratios of probabilities, they are the same probability distribution.

Now, suppose that p_1/p_2 has a different value according to a probability distribution p' , from what it has according to a distribution p . Define property A to hold exactly if the state of the system is S_1 or S_2 , and define property B to hold exactly if the state of the system is S_1 or S_3 . In that case $p(B/A)=p_1/(p_1+p_2)$. It then follows trivially that p and p' must also differ regarding the value of $p(B/A)$. Thus any difference in ratios entails a difference in some (in fact many) conditional probabilities for some (in fact many) pairs of properties. Moreover, if $p(B/A)$ has a different value according to 2 different probability distributions then the correlation between A and B is different according to the 2 different probability distributions. For any normal Markov process

there is, by definition, a state S such that if at time $t=0$ we have that the probability $p(S)=1$, then at time $t=1$ the probability distribution p is a non-equilibrium distribution. Since p is a non-equilibrium distribution, there must be some set of ordered pairs of states, such that their ratios according to p are not equal to their ratios according to any of the equilibrium probability distributions. Hence there must be some set of properties, such that the correlations amongst that set of properties are not the same in p as they are in any equilibrium distribution. Since p is immediately preceded by $p(S)=1$, it follows that even the finest prior partition does not screen off those non-equilibrium correlations. The fullest account of the prior 'causes' of the non-equilibrium correlations do not screen these correlations off. Thus for any normal Markov process the common cause principle fails.

II. Proof of a restricted common cause principle

Let us assume that the state of system A just after the interaction with C is determined by the states of systems A and C just before that interaction, and similarly that the state of system B just after the interaction is determined by the state of systems B and C just prior to the interaction. Moreover let us assume that the state of system C just after the interaction is determined by the state of systems A , B and C just before the interaction, while at all other times the development of each system is determined by its own state.

Now consider any property X of system A (or indeed any partition of the statespace of A), and any property Y of system B (or partition of the statespace of B). By determinism there must be some partition $\{C'_i\}$ of the statespace of C , such that the state of system A at $t=0$ and the cell of partition $\{C'_i\}$ that the state of C is in at $t=0$ together determine whether system A has property X at $t=1$, i.e. $P(X \text{ at } t=1 / \text{State}(A)=A_i \ \& \ \text{State}(C) \in C'_j \text{ at } t=0) = 1$ or 0 , for any state A_i and cell C'_j . Let me for notational convenience drop the obvious references to times, and write A_i for $\text{State}(A)=A_i$, and write C'_j for $\text{State}(C) \in C'_j$. Similarly, by determinism, there is a partition $\{C''_i\}$ such that $P(Y/B_i \ \& \ C''_i) = 1$ or 0 . Let $\{C_i\}$ be the partition one gets by using $\{C'_i\}$ to further partition $\{C''_i\}$. Then the above claims will continue to hold when one replaces $\{C'_i\}$ and $\{C''_i\}$ by $\{C_i\}$ in both claims.

By probability theory $P(X \ \& \ Y / C_i) = \sum_{j,m} P(A_j \ \& \ B_m) \cdot P(X \ \& \ Y / C_i \ \& \ A_j \ \& \ B_m)$.

By initial independence $P(A_j \ \& \ B_m) = P(A_j) \cdot P(B_m)$.

By the claims above $P(X \ \& \ Y / C_i \ \& \ A_j \ \& \ B_m) = P(X / C_i \ \& \ A_j \ \& \ B_m) \cdot P(Y / C_i \ \& \ A_j \ \& \ B_m)$.

Thus we have $P(X \ \& \ Y / C_i) = \sum_{j,m} P(A_j) \cdot P(B_m) \cdot P(X / C_i \ \& \ A_j \ \& \ B_m) \cdot P(Y / C_i \ \& \ A_j \ \& \ B_m)$.

But that equals $\sum_{j,m} P(A_j) \cdot P(B_m) \cdot P(X / C_i \ \& \ A_j) \cdot P(Y / C_i \ \& \ B_m)$.

By probability theory this equals $P(X / C_i) \cdot P(Y / C_i)$.

Thus we have proved that $P(X \ \& \ Y / C_i) = P(X / C_i) \cdot P(Y / C_i)$.

Thus we have shown that there is a partition of the statespace of C which screens off properties (or partitions) X and Y .

References

- Arntzenius, F. (1990), "Physics and Common Causes", *Synthese* 82: 77-96.
- Forster, M. (1986), "Unification and Scientific Realism revisited", in *PSA 1986*, Vol 1: 394-405.
- Horwich, P. (1987), *Asymmetries in Time*. Cambridge: MIT Press.
- Papineau, D. (1985), "Causal Asymmetry", *British Journal for the Philosophy of Science*, Vol 36: 273-289.
- Prigogine, I. (1980), *From Being to Becoming*. San Francisco: W.H. Freeman.
- Sober, E. (1988), "The Principle of the Common Cause", in *Probability and Causality*, J. Fetzer (ed.). Dordrecht: Reidel, pp 211-229.
- Sober, E. and Barrett, M. (1992), "Conjunctive Forks and Temporally Asymmetric Inference", *Australasian Journal of Philosophy*, Vol 70, No 1: 1-23.
- Spirtes, P. et al (1993), *Causation, Prediction and Search*. Springer Verlag.
- Van Fraassen, B. (1980), *The Scientific Image*. Oxford: Clarendon Press.
- Van Fraassen, B. (1982), "The Charybdis of Realism: Epistemological Implications of Bell's Inequality", *Synthese* 52: 25-38.