1

Introduction and overview

No correlation without causation. This is, in its most compact and general formulation, the essence of what has become *Reichenbach's Common Cause Principle*.

More explicitly, the Common Cause Principle says that every correlation is either due to a direct causal effect linking the correlated entities, or is brought about by a third factor, a so-called Reichenbachian *common cause* that stands in a well-defined probabilistic relation to the correlated events, a relation that explains the correlation in the sense of entailing it.

The Common Cause Principle is a nontrivial metaphysical claim about the causal structure of the World and entails that *all* correlations can (hence, should) be explained causally *either* by pointing at a causal connection between the correlated entities *or* by displaying a common cause of the correlation. Thus, the Common Cause Principle licenses one to infer causal connections from probabilistic relations; at the same time the principle does not address whether the causal connection holds between the correlated entities or between the correlated entities or between the correlated entities or between the causal connection holds between the correlated entities or between the common cause and the elements in the correlation.

While the technically explicit notion of common cause of a probabilistic correlation within the framework of classical Kolmogorovian probability theory is due to Reichenbach (1956), the Common Cause Principle was articulated explicitly only later, especially in the works by W. Salmon (see the "Notes and bibliographic remarks" to Chapter 2). The chief aim of this book is to investigate the Common Cause Principle; in particular, the problem of to what extent the Common Cause Principle can explain probabilistic correlations.

The Common Cause Principle has been discussed extensively both in the philosophy of science literature and in papers on foundations of physics, especially in the past thirty years. There seems to be consensus among philosophers of science that the principle is *not* universally valid – the literature is full of alleged CAMBRIDGE

2

Cambridge University Press 978-1-107-01935-5 — The Principle of the Common Cause Gábor Hofer-Szabó , Miklós Rédei , László E. Szabó Excerpt <u>More Information</u>

Introduction and overview

counterexamples: correlations that are claimed to exist between causally unrelated events that do not admit common causes. The counterexamples run from simple, everyday situations such as the correlation between bread prices in England and the water levels in Venice, Italy (both having been on the increase in the past two centuries) (Sober, 2001), to correlations arising from conserved classical physical quantities such as momentum (Cartwright, 1988), and correlations predicted by quantum theory [such as the Einstein–Podolsky–Rosen (EPR) correlations between spins] (Fraassen, 1982a).

But how can a pure existential claim be falsified so easily as these counterexamples seem to suggest? After all, the Common Cause Principle only states that the presence of correlations entails the existence of a common cause (if the correlated entities are causally independent), but it is completely silent about the further nature of the hypothetical common causes: neither their spatiotemporal properties nor any other features are prescribed by the Common Cause Principle. How could one then be so sure that the common causes of those simple correlations are definitely ruled out?

The answer is: one *cannot* so easily be sure. The main message of this book is that assessing the status of the Principle of the Common Cause is a very subtle matter requiring a careful investigation of both the principle itself and the evidence for/against it provided by our best scientific theories.

Specifically, the arguments from the above-mentioned counterexamples to the failure of the Common Cause Principle are too quick. What makes this perilous speed possible is in part the ambiguity and vagueness of the counterexamples in question: almost invariably, the probabilistic framework in which the counterexamples would be well-defined is *not* specified explicitly; this has the consequence that the problem of validity and falsifiability of the Common Cause Principle does not get a conceptually and technically sharp formulation.

By insisting on an explicit specification of the probabilistic model of the situation in which the problem of presence (or lack) of a common cause of a correlation can be meaningfully discussed is meant the specification of a classical (Kolmogorovian) probability measure space (X, S, p), where X is the set of elementary events, S is a Boolean algebra of certain subsets of X representing (general) events, and p is a (additive, in some cases a countably additive) probability measure on S. How to create a model of a random phenomenon in terms of a Kolmogorovian probability space is a nontrivial, nonmathematical questions, but without having set up such a model of a concrete situation explicitly, one cannot meaningfully discuss any probabilistic problem.

Specifically, the notions of correlation and of common cause are meaningful only within the framework of a given probability space: Given a classical probability space (X, S, p), events $A \in S$ and $B \in S$ are defined to be probabilistically

Introduction and overview

(positively) correlated with respect to the probability measure p if

$$p(A \cap B) > p(A)p(B) \tag{1.1}$$

3

Event $C \in S$ is called a Reichenbachian common cause of the correlation (1.1) if it satisfies four probabilistic conditions formulated in terms of the probability measure p: both C and its negation screen-off the correlation between A and B, and both A and B are more probable on condition C than they are on condition of the absence of the common cause C (see Definition 2.4 for more details). Note that this definition of common cause presupposes that events A, B, and C belong to the same Boolean algebra S, for if this were not the case then the probabilistic requirements would not be defined. This trivial observation has nontrivial consequences for the problem of falsification of the Common Cause Principle: the Principle states that if A and B are causally independent and correlated in the sense of (1.1), then there has to exist a common cause C of the correlation; however, the Principle does not require the common cause C to belong to any specific Boolean algebra; in particular, the common cause C need not belong to the S in which the correlated events A and *B* have been found. As a consequence, one cannot declare the Common Cause Principle invalid by displaying a particular probability space (X, S, p) that contains causally unrelated correlated events but no common cause of this correlation: all one is justified to say in this situation is that the probability space (X, S, p) is *com*mon cause incomplete. In other words, one can, in principle, argue that there might exist common cause events explaining the correlation, but the probabilistic model (X, \mathcal{S}, p) is just too meager to contain them. Such an argument is only maintainable, however, if one can show that the following holds: there exists a larger space (X', S', p'), a consistent extension of (X, S, p), which is rich enough in events to contain a common cause that explains the correlation in (X, S, p). This condition is necessary (although not sufficient) to defend the Common Cause Principle against an attempt to falsify it on the basis of displaying a common cause incomplete probability space. Can this necessary condition always be shown to hold?

Chapter 3, "Common cause extendability of probability spaces," investigates this problem. After an explicit definition in Chapter 2 of the notions of Reichenbachian common cause and of extension of a probability space, the notion of common cause incompleteness and common cause extendability of classical probability spaces are defined in Chapter 3. A classical probability space will be defined to be common cause extendable with respect to a given correlation if there exists an extension of the probability space is called *strongly* common cause extendable if for any type of common cause one can have an extension containing a common cause of the given type (Definition 3.7). (The type of a common cause is specified in Definition 3.6.) Are probability spaces common cause extendable? It is shown in

4

Introduction and overview

Chapter 3 that *every* classical probability space is *strongly* common cause extendable with respect to *any* given correlation (hence with respect to any *finite* number of correlations) (Proposition 3.9). Consequently, one can always defend the Common Cause Principle by claiming that correlations may have "hidden" common causes, "hidden" in the sense of not being accounted for in the event algebra of the probability space that predicts correlation between certain events. This is not to say that the extendability result should be interpreted as *proof* that the Common Cause Principle is valid – whether the common cause events in the extended probability space can be interpreted as representatives of empirically discernible "real" events is a question that needs careful scientific scrutiny. (Chapter 10, "Where do we stand?," discusses this point further.)

(Strong) common cause extendability of a classical probability space with respect to a finite set of correlations however, does not entail that the extension is (strongly) common cause *closed* (*complete*) in the sense of containing (every type of) common cause of every correlation it predicts: the extended probability space may very well contain correlations between events that do not belong to the Boolean algebra of events of the original probability space. (Indeed it *must* contain such correlations: for instance the correlations between C and A and C and B.) Therefore, it is not at all obvious that probability spaces exist that are (strongly) common cause closed. When can classical probability spaces be common cause closed? This problem is the topic of Chapter 4, "Causally closed probability theories." It is shown in this chapter that common cause closedness is not impossible mathematically - not even if the probability space has a finite number of events - but common cause closedness is not typical either. Chapter 4 gives a complete characterization of common cause closedness in terms of the measure theoretic atomicity properties of the probability measure spaces: It is shown that a probability space is common cause closed if and only if it contains at most one measure theoretic atom (Proposition 4.18). This result is then used to show that every classical probability space is not only common cause extendable, but common cause completable with respect to any set of correlations: every classical probability space can be extended into a common cause closed one (Proposition 4.19). It is not known if the strong version of this proposition also holds (Problem 4.20).

It will also be argued in Chapter 4 that common cause closedness is too strong a notion, however: in view of the Common Cause Principle, it is more natural to ask if a probability space (X, S, p) is *causally closed* with respect to a causal *independence* relation R_{ind} defined between elements of S – in the sense of containing a common cause of every correlation between elements A and B that are correlated *and* are causally independent, $R_{ind}(A, B)$. On what conditions on the probability space (X, S, p) and on R_{ind} is (X, S, p) causally closed? This problem is also analyzed in Chapter 4. It is shown that under weak and reasonable assumptions

Introduction and overview

on the causal independence relation $R_{ind}(A,B)$, causal closedness with respect to $R_{ind}(A,B)$ is possible even if the set of random events is finite. There are a number of open questions concerning causal closedness (Problems 4.13 and 4.14) though.

Chapter 5, "Common common causes" raises the problem of whether different correlations can, in general, have the *same* common cause (a so-called *common* common cause). On the basis of common cause extendability of probability spaces demonstrated in Chapter 4, one might think that different correlations can always have a *common* common cause in a sufficiently large probability space; however, it is shown in Chapter 5 that this intuition is wrong because the assumption that different correlations can have a *common* common cause entails certain conditions expressed in terms of the probabilities of the events involved (Proposition 5.4), and these necessary conditions can be violated by two pairs of correlated events in a simple probability space (Proposition 5.5 and its proof). In fact, we will show that given any correlation in any probability space, the probability space can be extended in such a way that the extension contains another correlation with the property that these two correlations cannot have the same common cause (Proposition 5.6). The upshot of this analysis is that different correlations cannot always have an explanation by a single common cause, no matter how refined a picture of the World one creates in terms of events in probability spaces. No necessary and sufficient conditions are known, however, that ensure the existence of common common causes of different correlations in general (Problem 5.8).

Another possible strategy one can follow in trying to explain correlations in common cause incomplete probability spaces is to take the position that the correlation is brought about not by a single common cause, but by a number of partial common-cause-like events. This idea is developed in Chapter 7 "Reichenbachian common cause systems." First, the notion of the Reichenbachian common cause is generalized to the notion of a Reichenbachian common cause system. A Reichenbachian common cause system is a partition of a Boolean algebra in such a way that any two elements of the partition behave like a Reichenbachian common cause and its negation (see Definition 7.1). The cardinality of the partition is called the size of the Reichenbachian common cause system. Reichenbach's original definition of common cause can then be viewed as a Reichenbachian common cause system of size 2. It is shown in Chapter 7 that if a correlation is not strict (not maximal), then one can, in principle, explain the correlation by a Reichenbachian common cause system of any finite size. It is an open problem whether this is also possible with a Reichenbachian common cause system of (countably) infinite cardinality (Problem 7.7); it is conjectured that this is possible. It also is not known whether strict (maximal) correlations also can be explained by an arbitrarily large (finite) Reichenbachian common cause system (Problem 7.8); it is conjectured that this also is possible. Chapter 7 also investigates the relation between different

5

6

Introduction and overview

Reichenbachian common cause systems of a given correlation. It is shown that Reichenbachian common cause systems possess a certain rigidity and uniqueness: there exists at most one, single common cause system in any *linearly* ordered subset of the partially ordered set of all partitions of a Boolean algebra, where the partial ordering is the finer–coarser relation between partitions (Proposition 7.5). This chapter closes with formulating natural definitions of causal closedness in terms of common cause systems (Definition 7.9); only very few results are known, however, concerning causal closedness with respect to common cause systems of cardinality greater than 2 (Problem 7.11 and "Notes and bibliographic remarks").

Chapter 6 entitled "Common cause extendability of nonclassical probability spaces" investigates the problem of common cause extendability of nonclassical probability spaces (\mathcal{L}, ϕ) where a general orthocomplemented, not necessarily distributive lattice \mathcal{L} takes the role of the Boolean algebra and where ϕ is an additive (countably additive) generalized probability measure on \mathcal{L} . Special examples of such nonclassical probability spaces are the quantum probability spaces $(\mathcal{N}, \mathcal{P}(\mathcal{N}), \phi)$ where $\mathcal{P}(\mathcal{N})$ is the orthomodular lattice of projections of a von Neumann algebra \mathcal{N} and ϕ is a countably additive probability measure on $\mathcal{P}(\mathcal{N})$ (a normal state on \mathcal{N}). A particular case of quantum probability spaces is $(\mathcal{B}(\mathcal{H}), \mathcal{P}(\mathcal{H}), \phi)$, where \mathcal{H} is a (finite or infinite dimensional) complex Hilbert space, $\mathcal{B}(\mathcal{H})$ is the von Neumann algebra of all bounded operators on \mathcal{H} , and $\mathcal{P}(\mathcal{H})$ is the projection lattice of all closed linear subspaces of Hilbert space \mathcal{H} (Hilbert lattice). [The Appendix contains a concise review of the main mathematical notions related to nonclassical (quantum) probability spaces, including von Neumann algebras.] Since the notion of (Reichenbachian) common cause was defined in classical probability spaces, to raise the problem of common cause extendability of nonclassical probability spaces the notion of common cause needs to be specified in terms of general probability spaces. This can in principle be done in several ways; we opt for a conservative strategy by requiring the common cause to be compatible with the correlated events that are also assumed to be compatible (see Definition 6.1 of common cause in general probability spaces and the references in the "Notes and bibliographic remarks" for Chapter 6 for other conceivable but less attractive options). The definition of common cause in general probability spaces is followed by the formulation of the problem of common cause extendability of general probability spaces - along the lines of the classical case (Problem 6.2). Remarkably, this problem remains entirely open. It is proved in Chapter 6, however, that quantum probability spaces $(\mathcal{N}, \mathcal{P}(\mathcal{N}), \phi)$ can always be common cause extended with respect to all the correlations predicted by a single quantum state (Proposition 6.3). The notions of common cause closedness and causal closedness also can be defined in nonclassical probability theories (Definitions 6.4 and 6.5) in close analogy with the classical counterparts of these concepts,

Introduction and overview

and one can give a characterization of common cause closedness of nonclassical probability spaces in terms of the measure theoretic atomicity property of these spaces: we show that under some additional conditions, general probability spaces with one single measure theoretic atom are still common case closed (Proposition 6.15) and that probability spaces with two measure theoretic atoms are *not* common cause closed (Proposition 6.12).

The positive results on common cause extendability of both classical and quantum probability spaces entail that it is always possible in principle to explain correlations by common causes; in other words, that a necessary (but not sufficient) condition for the explanation always holds. This entails that in order to falsify Reichenbach's Common Cause Principle, one has to impose some further conditions on the common cause, conditions that go beyond the four probabilistic relations originally formulated by Reichenbach. The extra conditions should be inferred from the special features of the situations modeled by probability spaces. Chapters 8 and 9 investigate in detail whether correlations between spatiotemporally *localized* events can always be explained by *properly localized* common causes.

Chapter 8 "Causal closedness of quantum field theory" recalls first the special local correlations predicted by local relativistic quantum field theory as a consequence of the violation of Bell's inequalities in local relativistic quantum field theory. Because the correlations predicted by this theory are between local observables pertaining to spacelike separated spacetime regions, which are regarded causally independent by the Special Theory of Relativity, one would like to see a properly localized common cause of these correlations. "Properly localized" here means: localized in a spacetime region that lies within the intersection of the backward light cones of the spacelike separated spacetime regions that contain the correlated observables. This localization is, however, not the only one that is theoretically possible: Both stronger and weaker localizations of the common cause are feasible and, accordingly, there are in principle three distinct, nonequivalent ways in which local, relativistic quantum field theory can be compatible with the Common Cause Principle. Definition 8.11 specifies the corresponding three notions of causal closedness of local quantum field theory and the problem is raised then whether local quantum field theory is causally rich enough to contain "strongly," "properly," and "weakly" localized Reichenbachian common causes of the spacelike correlations it predicts. Surprisingly, the problem of existence of properly localized common causes is completely open (Problem 8.13). It turns out, however, that if the Local Primitive Causality axiom holds in local relativistic quantum field theory (Definition 8.3), then there exist Reichenbachian common causes localized in the union of the causal pasts of the spacelike separated spacetime regions containing the correlated observables (Proposition 8.14). In this weak

7

8

Introduction and overview

sense at least, local relativistic quantum field theory respects Reichenbach's Common Cause Principle. It is an immediate consequence of the violation of Bell's inequality for algebras pertaining to complementary wedge regions in local relativistic quantum field theory that strongly localized common causes of spacelike correlations do not exist in general (Proposition 8.12). This chapter also investigates the problem of the status of the Common Cause Principle in lattice quantum field theory – a discrete version of quantum field theory in which computations can be carried out more easily due the fact that the local observable algebras have finite dimension. It is proved in this chapter that in lattice quantum field theory, even the weak Common Cause Principle does not hold: Discrete quantum field theory contains correlated projections localized in algebras pertaining to spacelike separated discrete points for which there exist no common cause at all in any local algebra - no matter where the local algebra is situated in the lattice. It is shown, however, that if one weakens the notion of common cause by allowing it to not commute with the correlated projections (Definition 8.26), then weakly localized common causes do exist in lattice quantum field theory as well (Proposition 8.30).

Chapter 9 "Reichenbach's Common Cause Principle and EPR correlations" investigates the problem of whether one can in principle provide an explanation of the famous EPR correlations in terms of common causes. The common cause extendability results in the previous chapters entail that the EPR correlations *can* in principle be explained by common causes – if no conditions are imposed on the common causes in addition to the standard Reichenbachian ones. In the case of EPR correlations, however, one has extra information, both about the probabilities of certain events in a probabilistic model of the EPR correlation experiment and about the spatiotemporal (hence causal) structure of the correlation experiment. One has to take into account this additional information when defining a common cause explanation of the EPR correlations, and this leads naturally to imposing some additional probabilistic constraints on the common causes. There are two sorts of extra requirements: "locality" and "no-conspiracy."

Extra care must be exercised, however, when formulating these locality and noconspiracy conditions because, first, in view of the distinction between common causes and common cause *systems*, one has to be careful about whether one is looking for a common cause or for a common cause *system* when seeking an explanation of correlations – while a common cause for a correlation may not exist, a common cause *system* might.

Second, the EPR correlations involve more than one pair of correlated events but the Common Cause Principle only concerns a *single* correlation. This is important to realize because it was seen in Chapter 5 that common causes are not *common* common causes – nor are, therefore, common cause systems *common* common cause systems in general. Clearly, the weakest question one can ask in connection

Introduction and overview

with a set of correlations is whether a *separate* common cause *system* exists for *each* correlation in the set. Thus, Chapter 9 will ask whether a *separate* common cause system exists for *each* pair of correlated events occurring in the EPR situation in such a way that these separate common cause systems satisfy the additional locality and no-conspiracy requirements.

Matters are further complicated by the fact that both locality and no-conspiracy come in two forms: surface locality [conditions (9.23)-(9.24)] and hidden locality [conditions (9.147)–(9.148)] on the one hand, and weak no-conspiracy (Definition 9.6) and strong no-conspiracy (Definition 9.7) on the other. "Surface locality" expresses that the probability of the outcome in one wing of the EPR correlation experiment is the same no matter in which direction the measurement is carried out in the other wing of the EPR experiment. The hidden locality condition requires that, given a pair of correlated outcomes of an EPR correlation experiment, the probability of the outcome in one wing of the EPR correlation experiment be the same no matter in which direction the measurement is carried out in the other wing of the EPR experiment *if* the hypothetical common cause of this correlation also has happened. The weak no-conspiracy condition requires that any choice of the direction in which a measurement is decided to be carried out in any wing is probabilistically independent of the hypothetical common cause explaining the correlation in the chosen direction. The strong no-conspiracy condition requires that any Boolean combination of choices of measurement directions in the two wings of the EPR experiment are probabilistically independent of any Boolean combination of any of the hypothetical common causes.

It is shown in Chapter 9 that the EPR correlations *can*, in principle, be explained by common causes that satisfy the surface locality and the weak no-conspiracy conditions (Proposition 9.10). The significance of this proposition is that the EPR correlations (hence quantum mechanics) do *not* constitute *strictly empirical evidence* against the Common Cause Principle. The other major result in this chapter is that there do *not* exist separate common cause systems for each EPR correlation in a four element set that are both *hidden* local and *strongly* nonconspiratorial (Proposition 9.16). The proof of this latter proposition is based on deriving an "approximate Clauser–Horne-type inequality" [Equation (9.239)] from the assumption of the existence of a separate, hidden local and strongly nonconspiratorial common cause system for each of the four correlations and showing that the inequality is violated by correlations in some directions.

The closing Chapter 10, "Where do we stand?," summarizes the main points of the analysis of the Common Cause Principle carried out in the previous chapters. We take the epistemologically moderate position in this final chapter (in fact, in the whole book) that general metaphysical principles such as the Common Cause Principle cannot be verified or falsified conclusively; rather, one should aim at a

9

10

Introduction and overview

careful assessment of the status of such principles in light of the evidence provided by our best scientific theories. The evidence available at this time and presented in the present book is mixed and can be summarized as follows: The formal analysis shows that it is *always* possible in principle to explain correlations in terms of (possibly "hidden") common causes if one does not require of the common causes to have features in addition to the properties specified in their definition given by Reichenbach. Relativistic quantum field theory is not a disconfirming evidence for the Common Cause Principle in spite of the abundance of spacelike correlations it predicts; what is more, quantum field theory might even turn out to be a confirming evidence for the Principle if it can be shown to be causally complete in the sense of providing *local* common causes for the spacelike correlations – the problem of causal completeness of quantum field theory is open at this time. The EPR correlations are not disconfirming evidence either if the hypothetical common causes explaining the correlations are required to be surface-local and weakly conspiratorial only; however, these correlations are not compatible with a more heavily metaphysical Common Cause Principle that requires the common causes explaining the EPR correlations to satisfy hidden locality and strong no-conspiracy as well.