


Centre for the Philosophy of Natural; and Social Science
Contingency and Dissent in Science
Technical Report 07/08

The Principle of Common Cause and Indeterminism: A Review

Iñaki San Pedro and Mauricio Suárez



Series Editor: Damien Fennell

The support of The Arts and Humanities Research Council (AHRC) is gratefully acknowledged. The work was part of the programme of the AHRC Contingency and Dissent in Science.

Published by the Contingency And Dissent in Science Project
Centre for Philosophy of Natural and Social Science
The London School of Economics and Political Science
Houghton Street
London WC2A 2AE

Copyright © Iñaki San Pedro and Mauricio Suárez 2008

ISSN 1750-7952 (Print)
ISSN 1750-7960 (Online)

All rights reserved.

No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, without the prior permission in writing of the publisher, nor be issued to the public or circulated in any form of binding or cover other than that in which it is published.

The Principle of Common Cause and Indeterminism: A Review*

Iñaki San Pedro⁺ and Mauricio Suárez[‡]

Editor's Note

How and under what conditions can one infer causal relationships from correlations? This paper contributes to the ongoing project work on the strengths and limits of methods of causal inference by critically reviewing work on the principle of the common cause, with a particular focus on the cases of inferring causal relationships from correlations in fundamentally indeterministic situations.

Abstract

We offer a review of some of the most influential views on the status of Reichenbach's Principle of the Common Cause (PCC) for genuinely indeterministic systems. We first argue that the PCC is properly a conjunction of two distinct claims, one metaphysical and another methodological. Both claims can and have been contested in the literature, but here we simply assume that the metaphysical claim is correct, in order to focus our analysis on the status of the methodological claim. We briefly review the most entrenched or classical positions, including Salmon's 'interactive forks', van Fraassen's scepticism, and Cartwright's generalisation of the fork criterion. We then go on to review the results of the 'Budapest school' on the existence of formally defined screening-off events for any correlation —by means of the ideas of probability space extensibility and (Reichenbachian common cause) completability. We distinguish the Budapest doctrine clearly from any of the classical conceptions, and thus present an overall framework for discussions of causal inference in quantum mechanics. The framework, however preliminary, is essential work for a thorough assessment of the conditions under which PCC may be a reliable tool for causal inference in a genuinely probabilistic (indeterministic) context.

1. Reichenbach's Principle of the Common Cause (PCC)

The discussion of common causes in the contemporary literature has its origins in Reichenbach's work.¹ Reichenbach stated that apparently unrelated events that nonetheless take place correlatively—in a prescribed sense—underlie a common cause:²

If an improbable coincidence has occurred, there must exist a common cause.

This quote expresses a 'metaphysical' or 'ontological' claim, since it suggests that 'common

* We gratefully acknowledge support from the Spanish Ministry of Education and Science (HUM2005-07187-C03-01), and the Education Department of Madrid's Autonomous Community (S2007-HUM/0501).

⁺ Universidad del País Vasco, UPV-EHU

[‡] Universidad Complutense de Madrid.

1 Cf. (Reichenbach, 1956).

2 (Reichenbach, 1956, p. 157).

causes' are provided whenever 'improbable coincidences' occur. We here shall refer to it as Reichenbach's *Postulate of the Common Cause* (PosCC).

This claim *per se* says nothing about what specific probabilistic relations, if any, common causes must obey. This is filled in by a conceptually separate criterion —also introduced by Reichenbach— which we shall refer to as Reichenbach's *Criterion for Common Causes* (CritCC).³ In Reichenbach's own words:⁴

In order to explain the coincidence of A and B , which has a probability exceeding that of a chance coincidence, we assume that there exists a common cause C . [...] We will now introduce the assumption that the fork ABC satisfies the following relations:

$$\begin{aligned} p(A \wedge B / C) &= p(A / C) \cdot p(B / C) \\ p(A \wedge B / \neg C) &= p(A / \neg C) \cdot p(B / \neg C) \\ p(A / C) &> p(A / \neg C) \\ p(B / C) &> p(B / \neg C) \end{aligned}$$

We are not here following Reichenbach's notation precisely. Instead we use $p(X / Y)$ to represent the conditional probability of X on Y (with $X = A, B$ and $Y = C, \neg C$); and $X \wedge Y$ to represent the joint event "X and Y". It is then clear that the last two expressions represent correlations, to be expected in an unbiased run of the relevant experiments, between event types A and C , and B and C . But, since the main question at stake in most theories of causal inference concerns the legitimacy of inferences from probabilistic dependencies to causal relations, we may not assume without begging the question that C is a positive relevant cause of either A or B .

As for the first two probabilistic conditions, they express a restriction on the postulated common cause C , introduced specifically by Reichenbach to account for common causes. They require, specifically, that when the presence (or the absence) of the common cause is taken into account by conditionalising on it, the correlated events A, B are rendered probabilistically independent. The common cause C is then said to *screen-off* the correlation $Corr(A, B)$.⁵

In what follows, we will refer to the conjunction of the Postulate of the Common Cause (PosCC) and the Criterion for Common Causes (CritCC) as *Reichenbach's Principle of the Common Cause* (PCC). Thus whenever Reichenbach's Principle of the Common Cause (PCC) is

³ The distinction between these two notions is first introduced in (Suárez, 2007).

⁴ (Reichenbach, 1956, p. 159).

⁵ And indeed the first two expressions in the quotation above can be shown to be equivalent to what is commonly known in the literature as the *screening-off* condition, which is probabilistically expressed as $p(A / B \wedge C) = p(A / C)$.

invoked it involves *both* claims above. This is the standard usage of the term, but properly distinguishing the two claims that are involved is crucial, as we shall argue, for the assessment of the status of PCC as a whole.⁶

Although both the Postulate of the Common Cause (PosCC) and the Criterion (CritCC) are intended as causal claims they have a completely different philosophical significance. The Postulate (PosCC) is a metaphysical claim that states the existence of common causes. The Criterion (CritCC), by contrast, is a methodological claim which, although complementary to (PosCC), is logically independent from it. In particular, while the Postulate informs us about the ontology of the possible causal structure underlying the correlation between *A* and *B*, the Criterion aims to provide the tools for an adequate statistical characterisation of such causal structure. The four probabilistic relations that define a conjunctive fork *ABC* aim to provide an adequate characterisation of the statistical relevance of the common causes postulated by the first (metaphysical) claim, but they do not themselves express anything causal. In other words the Criterion could be true even were the Postulate false, in which case it would just define formally the set of “screening-off” events, or screener-offs, devoid of any genuine causal significance.⁷

Tradition has often followed Reichenbach in supposing that PCC, and CritCC in particular is explanatory. As Reichenbach writes:⁸

When we say that the common cause *explains* the frequent coincidence we refer not only to this derivability of [correlation between *A* and *B*] but also to the fact that relative to the cause *C* the events *A* and *B* are mutually independent.

It has in addition often been supposed that the explanatory power of the Criterion for Common Causes (CritCC), and screening-off in particular, is if anything grounded upon the inference to the hidden but real common causes promoted by the ontological Postulate of Common Causes (PosCC).⁹ But once the Criterion and the Postulate are clearly distinguished, as we do here, there

6 There is a certain amount of confusion in the literature regarding the terminology and a clarifying note is perhaps in order. In some cases the expression ‘Reichenbach’s Principle of the Common Cause’ is used to refer just to the four probabilistic relations, i.e. Reichenbach’s characterisation of the postulated common causes, which we here refer to as the Criterion for Common Causes (CritCC). In other occasions it is used to refer to what we here call the Postulate of Common Causes (PosCC). But more often still it is taken to apply to the conjunction of both. Our terminology clears up any lingering confusion, by distinguishing clearly between two distinct commitments: the *Criterion* and the *Postulate*. The term *Principle of the Common Cause* (PCC) is then reserved for the strongest form of commitment to the conjunction of both.

7 Of course, the converse is also possible. Trivially, the Postulate may be true even if the Criterion were false. For instance, there could be no way to reliably infer any causal conclusions whatever from probabilistic relations grounded upon statistics, in the form of the Criterion, or any other form, and this still would not settle the issue as to whether common causes exist.

8 Cf. (Reichenbach, 1956, p. 159).

9 Salmon (1984) and Cartwright (1989) are perhaps the most salient examples. This explanatory order, which ranks

seems to be no reason to suppose that the explanatory power of the Criterion need depend on anything other than itself. Certainly, as is clear in the quote above, Reichenbach himself seems to have thought that the mere satisfaction of screening off is explanatorily efficient.¹⁰ We will here generally follow suit and agree that there is a form of explanation that bequests the Criterion with explanatory power regardless of whether the Postulate is true or false.

2. Indeterminism and Reichenbach's Criterion for Common Causes

Screening-off is well known not to be a *sufficient* condition on common causes: not all screener-offs are common causes. For instance, common effects of two separate causes, lying in the proper future of both, may also screen off but they are definitely not common causes.¹¹ And there are many other significant cases.¹² However, following Reichenbach, screening-off has often been taken to be a *necessary* condition on common causes. This would at least allow negative causal inference, identifying those events which are not common causes by means of violations of the screening-off condition. In other words, if the Criterion is necessary for common causes then we would be able to use violations of the screening-off condition to discard the events that are not common causes. Otherwise violations of the screening-off condition would not provide any useful information about the underlying causal structure. This is a key to both the classical and contemporary discussions, and several examples have been employed to suggest that some common causes might violate screening-off.¹³

The most interesting and powerful arguments against screening-off as a necessary condition for common causes involve genuinely probabilistic causes. In particular, counterexamples to Reichenbach's Criterion for Common Causes (CritCC) typically consider correlations between events which occur in tandem —as a result perhaps of some conservation law—, both as an effect of an in principle (patently) obvious common cause. One such example was first proposed by van

the *Postulate* ahead of the *Criterion*, also follows from the form of 'causal realism' that Van Fraassen (1982a) has attacked. Unlike Van Fraassen we are not causal sceptics —we appreciate that the explanatory power of the *Postulate* is fully causal, while any explanatory power that the *Criterion* might have on its own would fall short of such a standard. But there are different modes of explanation, and the kind of explanatory power that the Criterion enjoys on its own seems to us legitimate too. In this paper we argue that more work is needed in order to assess the explanatory power of different 'causal' assumptions.

10 As we will argue later, the autonomous explanatory power of screening-off may provide methodological motivation to apply Reichenbach's Criterion for Common Causes whenever possible, even though it will be shown to patently fail as a necessary or sufficient condition for causation.

11 (Reichenbach, 1956, p. 159).

12 For a discussion see (Suárez, 2007, section 2).

13 The validity of Reichenbach's Principle of the Common Cause (PCC) is tightly linked to that of the Causal Markov Condition (CMC). It is usually acknowledged that the CMC holds for deterministic causes but it is controversial whether it does too for genuinely probabilistic ones. The examples we discuss here have indeed figured in the intense debate on the status of the CMC, which is not our aim in this paper to review. For our views on the topic see (Suárez and San Pedro, 2007).

Fraassen¹⁴ as an argument against Salmon’s defence of conjunctive forks.

The example consists of a particle that collides with an atom. As a result of the collision the atom emits two new particles. Suppose for simplicity that the angle with which each particle is emitted can only take two values, each with probability 1/2. That is to say, PARTICLE 1 may be emitted either at angle θ or at angle θ' , each with probability 1/2. PARTICLE 2, on the other hand, may be emitted either at angle $-\theta$ or at angle $-\theta'$, also with probability 1/2 in each case.

Now because of the conservation of linear momentum, if PARTICLE 1 is emitted at angle θ (expressed as 1_θ), PARTICLE 2 must be emitted at angle $-\theta$ (expressed as $2_{-\theta}$), and conversely. More precisely, due to conservation of momentum the corresponding angles at which the particles are emitted are *perfectly correlated*¹⁵:

$$\begin{aligned} p(1_\theta / 2_{-\theta}) &= p(2_{-\theta} / 1_\theta) = 1, \\ p(1_{\theta'} / 2_{-\theta'}) &= p(2_{-\theta'} / 1_{\theta'}) = 1. \end{aligned}$$

A common cause λ may now be postulated such that, if present, the particles are emitted at angles θ and $-\theta$ respectively. Otherwise the particles are emitted at θ' and $-\theta'$. A final feature of the example is that if the postulated common cause is deterministic, the joint probabilities of the two particles factorise, i.e. the screening-off condition is satisfied.¹⁶ If the postulated common cause is purely probabilistic, however, screening-off need no longer be satisfied. Let’s look into the argument in a bit more detail.

Take the deterministic case first. A deterministic common cause must obey:

$$\begin{aligned} p(1_\theta / \lambda) &= p(2_{-\theta} / \lambda) = 1, \\ p(1_{\theta'} / \neg\lambda) &= p(2_{-\theta'} / \neg\lambda) = 1. \end{aligned}$$

14 See (van Fraassen, 1982b). This and other arguments by van Fraassen —such as (van Fraassen, 1982a)— against common causes are however motivated by his attempt to reject Reichenbach’s Principle of the Common Cause (PCC) altogether. Van Fraassen does not distinguish clearly between the Postulate and the Criterion, but having shown the Criterion to fail, he seems to reject the Postulate as well. In light of what we argue here, it would seem that he is unnecessarily throwing the baby out with the bathwater. There is no need to reject the metaphysical claim as a consequence of the failure of the methodology, nor does the presumption that the metaphysical claim is false amount to a proof that the Criterion must go. See Section 4.3 and Table 2 for details of van Fraassen’s position.

15 The expressions below do not intend to provide a formal definition of *perfect correlation*. However, it is easy to check that they correspond to the special case of maximal (perfect) correlation. In what follows, ‘perfect correlation’ will refer to events conforming to expressions of the kind below.

16 In the case of perfect correlations, the postulated (two-valued) common cause must then be deterministic, if it is a screening-off common cause. This is shown, for instance in (van Fraassen, 1982a) and, more recently in (Graßhoff, Wüthrich and Portman, 2005). The original result is proved by (Fine, 1982), who shows that if there exist a screening-off hidden variable for a perfect correlation then there exists as well a deterministic hidden variable model for it, and *vice versa*. In other words, there is no conceptual room for indeterminism when common causes are meant only for perfect correlations (or perfect anti-correlations).

And for joint probabilities, we have

$$p(1_{\theta} \wedge 2_{-\theta} / \lambda) = 1,$$

$$p(1_{\theta'} \wedge 2_{-\theta'} / \neg\lambda) = 1.$$

It is then straightforward to prove that the corresponding probabilities factorise (since all probabilities equal one):

$$p(1_{\theta} \wedge 2_{-\theta} / \lambda) = p(1_{\theta} / \lambda) \cdot p(2_{-\theta} / \lambda),$$

$$p(1_{\theta'} \wedge 2_{-\theta'} / \neg\lambda) = p(1_{\theta'} / \neg\lambda) \cdot p(2_{-\theta'} / \neg\lambda),$$

Two other screening-off conditions may be written replacing the common cause λ by its negation in the first equation above and the $\neg\lambda$ by the common cause in the second. In that case all are zero probabilities and screening-off is again trivially satisfied. Finally, the $1/2$ probabilities for the occurrence of each of the events separately are in this case simply reproduced by assuming that the common cause occurs with probability $1/2$, i.e. $p(\lambda) = 1/2$.

The issue turns out to be entirely different if λ is a genuinely probabilistic cause. In this case the occurrences of 1_{θ} and $2_{-\theta}$ ($1_{\theta'}$ and $2_{-\theta'}$) are still perfectly correlated, exactly as in the deterministic case. Recall as well that the observed probabilities are

$$p(1_{\theta}) = p(2_{-\theta}) = p(1_{\theta} \wedge 2_{-\theta}) = 1/2,$$

$$p(1_{\theta'}) = p(2_{-\theta'}) = p(1_{\theta'} \wedge 2_{-\theta'}) = 1/2.$$

But, since λ is a genuinely probabilistic cause, the probability for the occurrence of 1_{θ} ($1_{\theta'}$) given that the common cause is present is now, say, r (r'), i.e. it is different from one: $p(1_{\theta} / \lambda) = r$ (or $p(1_{\theta'} / \lambda) = r'$). We do not need to know at this point what the probability $p(\lambda)$ of the common cause is.¹⁷ What matters is that the restrictions imposed by the conservation of momentum, which entail that the events 1_{θ} and $2_{-\theta}$ ($1_{\theta'}$ and $2_{-\theta'}$) are perfectly correlated, ensure that the following probabilities obtain:

¹⁷ The result we are aiming for is not dependent on that number. Even in the case the common cause happened to be present in all cases, i.e. $p(\lambda) = 1$, the result above would obtain, as we will explain.

$$p(1_\theta / \lambda) = p(2_{-\theta} / \lambda) = p(1_\theta \wedge 2_{-\theta} / \lambda) = r,$$

$$p(1_{\theta'} / \lambda) = p(2_{-\theta'} / \lambda) = p(1_{\theta'} \wedge 2_{-\theta'} / \lambda) = r'.$$

Since the joint probabilities are now equal to the marginal probabilities (and all assumed to be different from 1 or 0) it is easy to check that the corresponding screening-off conditions are now violated:

$$p(1_\theta \wedge 2_{-\theta} / \lambda) \neq p(1_\theta / \lambda) \cdot p(2_{-\theta} / \lambda),$$

$$p(1_{\theta'} \wedge 2_{-\theta'} / \lambda) \neq p(1_{\theta'} / \lambda) \cdot p(2_{-\theta'} / \lambda),$$

where, again, replacing the common cause λ by $\neg\lambda$, yields the corresponding two expressions for the negation.

This example shows that in the case of genuinely indeterministic systems there are plausible common cause explanations for certain correlations —those arising from the conservation of a quantity in particular— which do not fulfil the screening-off conditions required by Reichenbach. In other words, screening-off is not a necessary condition on common causes in general.

3. Is There Need for a New Common Cause Criterion?

The conclusion in the previous section poses further interesting questions. One might conclude that if the Criterion (CritCC) is neither necessary nor sufficient for common causes, then we ought to abandon the Postulate (PosCC) as well.¹⁸ Yet, one might insist on retaining the Postulate come what may. This opens up three different logical possibilities. First, we may impose further restrictions on common causes so as to find a set of jointly sufficient conditions that will enable us to identify common causes straightaway. Or we may precisely conversely weaken the Criterion (CritCC) on common causes in order to establish a merely necessary condition. We would hence restore the possibility of negative causal inference, by identifying those events which are definitely not common causes. The third alternative would be to keep Reichenbach's *Criterion* for Common Causes (CritCC), as it stands, and explore further the nature of those events that conform to it, i.e. figure out what all screener-offs may have in common, and whether what they have in common has any residual causal character. We might need to employ additional causal information in this endeavour, and if so, the third approach would turn out to be problematic as a method for inductive

18 And indeed this seems to be the moral van Fraassen wants to draw from his example of the 'bombarded atom'.

causal inference from statistics alone.

3.1. Strengthening the Conjunctive Fork Criterion

The first option invites us to suppose that the *Criterion* embodies some minimal features of common causes, which are insufficient to characterise them entirely. So if we are to provide a sufficient condition for common causes we must require that these fulfil stronger conditions, well beyond those expressed in the Criterion. Such further conditions may enable us to identify directly the right common causes from statistical (or probabilistic) information alone.

This option, however, does not seem very promising. It is hard to see what general condition we could add, other than temporal order, and this will only get us around some of the counterexamples. But proceeding by just imposing special further conditions on common causes designed to avoid specific problems would be *ad hoc*. And if the further conditions turned out not to be generalisable, we could hardly claim to have obtained a more constraining set of necessary and sufficient conditions. However, we may need not worry much about the screening-off condition not being sufficient for common causes. For the fact that there exist screener-offs which cannot be regarded as common causes of a correlation does not entail that there exists no common cause for the correlation—even of the screening-off variety. The example cited by Reichenbach of common effects of separate antecedent causes is a very good illustration. Thus we may conclude that strengthening conjunctive forks, whether or not it is possible or plausible, is not really needed for our purposes.

3.2. Weakening the Conjunctive Fork Criterion

The second alternative is to weaken Reichenbach's criterion in the hope that a necessary condition for common causes may finally be found. This is the option that *prima facie* seems preferable in light of examples such as van Fraassen's 'bombarded atom' in the previous section. In particular, it seems quite reasonable to ask how the non-screening-off common causes suggested by such examples may be characterised, if at all, in terms of probabilistic relations. If Reichenbach's Criterion can not do this, then perhaps we should turn to a more appropriate probabilistic characterisation. Provided such characterisation may be turned into a distinct criterion of its own—perhaps more general in scope than (CritCC)—, we may ask whether it is an appropriate necessary condition on common causes. If this is the case, we will then have a tool available for rightfully dismissing certain events as definitely not common causes—i.e. those events which violate our 'generalised' criterion.

This was basically the option chosen by Salmon, for instance, when introducing his *interactive forks* in order to characterise a further set of common causes that failed to satisfy screening-off. We do not need to review Salmon's interactive forks here since they will not play any major role in the foregoing discussion. It is enough for our purposes to stress that interactive forks constitute an important weakening of the conditions on common causes. Moreover Salmon's interactive forks, despite implying logically weaker conditions on common causes, do not constitute a necessary condition for common causes either. They just characterise a further set of common causes, but need not exhaust the concept. So we are back with a dilemma: either we find a further weakening of the conditions on common causes, or we stick to our initial Criterion, and the screening-off condition that it embodies, and push it as far as it will go.

Our view is that from the point of view of causal inference, the only really powerful reason for a further weakening of Reichenbach's *Criterion* would be a demonstration that the new, weaker, criterion reaches further—in the sense of providing a more effective characterisation of common causes in probabilistic terms. This seems to be the underlying motivation for Nancy Cartwright's generalisation of Reichenbach's Criterion for Common Causes.

Cartwright¹⁹ believes that (CritCC) is an inappropriate characterisation of common causes, particularly those that are genuine indeterministic. (She extracts this diagnosis from examples not dissimilar in structure to van Fraassen's 'bombarded atom'). Cartwright thus proposes to generalise Reichenbach's Criterion, in order to find a genuinely necessary condition on all kinds of common causes, whether deterministic or not. This is to be achieved also by weakening the conditions on common causes, but *only as much* as required for genuinely probabilistic common causes. So Cartwright's response to the counterexamples to screening-off is more informative than Salmon's. Salmon builds a new category of common cause (the "interactive fork") which relinquishes any probabilistic conditions whatever, and embraces a fully ontological reading in terms of interacting causal processes. Cartwright aims to keep a minimal connection between common causes and probabilistic relations. It remains to be seen whether Cartwright's generalisation of CritCC is indeed a more appropriate characterisation of common causes. We turn to this question in the next section.

3.3. Indeterministic Common Causes and the Conjunctive Fork Criterion

Cartwright's generalisation of the conjunctive fork criterion relies on causal modelling techniques, which need not be discussed here. It will be enough for our purposes to discuss and assess Cartwright's arguments to the effect that CritCC is not in general an appropriate criterion for common causes—specifically indeterministic common causes. These are the arguments that

19 Cf. (Cartwright, 1987).

ultimately led her to propose her generalised criterion.²⁰

We mentioned that the source of Cartwright's concerns lies in examples such as van Fraassen's 'bombarded atom'. In such examples a deterministic common cause would not constitute a problem for the screening-off condition, but it cannot account for the correlations. And Cartwright does seem to endorse the view that deterministic common causes do in general satisfy screening-off.²¹ This would include deterministic instances of Salmon's interactive forks for perfect correlations, i.e. perfect forks.

However, postulating a genuinely probabilistic common cause for the 'bombarded atom' example results in a violation of screening-off. Cartwright puts the blame, not only on the genuinely stochastic character of the common cause, but also on the fact that the postulated common cause *operates under a constraint* (conservation of momentum):²²

But in this case it is not reasonable to expect the probabilities to factor, conditional on the common cause. Momentum is to be conserved, so the cause produces its effects in pairs. [...] Clearly the conjunctive fork criterion is not appropriate here. That is because it is a criterion tailored to cases where the cause operates independently in producing each of its effects: whether one of the effects is produced or not has no bearing on whether the cause will produce the other.

It is important to stress that for Cartwright the violation of screening-off is a consequence of two facts, namely that the common cause is genuinely indeterministic *and* that the common cause operates under a constraint. Cartwright nowhere seems to suggest that any of these facts separately accounts for the violation of screening-off. She does, however, urge us to draw the following lesson from a superficially similar example:²³

Lesson: where causes act probabilistically, screening-off is not valid. Accordingly, methods for causal inferences that rely on screening-off must be applied with judgement

20 Cartwright's generalisation of CritCC results from a study of the dependencies that should be expected in a three variable causal model when the common cause variable operates under a constraint to produce its effects in pairs, or in tandem. These are mainly considerations about the coefficients that relate the three events in the causal model (derived from the fact that the two effects occur in tandem). In particular, the coefficients relating each of the effects and the common cause are not independent. They may even happen to be the same —e.g. whenever the common cause operates in order to produce both effects at once (although in a genuinely stochastic manner). For the full formal details of Cartwright's generalisation the reader is directed to (Cartwright, 1987), where the generalised criterion was first formulated. See also (Cartwright, 1989) for further details.

21 See our remarks on Footnote 16.

22 Cf. (Cartwright, 1987, p. 184).

23 The example in question is the famous 'Cheap-but-Dirty /Green-and-Clean' example, which was first discussed in (Cartwright, 1993). We shall go back to this example in Section 4 in order to motivate the idea of completability. The quote below is from (Cartwright, 1999a, p. 8), where the factory example is also discussed.

and cannot be relied on universally.

So let us suppose for the sake of the argument that Cartwright’s generalisation of (CritCC) is intended to account for indeterministic common causes that operate under a constraint. Is it the case that examples such as van Fraassen’s justify a generalisation of CritCC? In other words, is it really the case that the Criterion (CritCC) cannot account for indeterministic common causes that operate under a constraint? In order to answer these questions let us first go back to some of the key features of the ‘bombarded atom’ example.

A crucial feature of van Fraassen’s example is that the effects (of the atom being bombarded) are *perfectly correlated*. This is also what Cartwright refers to with the expressions ‘effects in pairs’ or ‘tandem effects’. So the examples are quite specific in that they all involve both genuinely indeterministic causes, *and* perfectly correlated effects. But as regards genuine indeterminism, perfect correlation is a very special particular case. For it can be shown that screening-off common causes —that is Reichenbachian Common Causes— of perfect correlations are *deterministic* common causes.²⁴ To formalise it in precise terms:

$$\text{PCORR} \wedge \text{SO} \rightarrow \text{DC} \quad (1)$$

where PCORR stands for ‘perfect correlation’, SO for ‘screening-off’ and DC for ‘deterministic common cause’.

But, since we are here interested in genuinely indeterministic common causes, we must consider the negation of the expression above instead. That is:

$$\begin{aligned} \neg \text{DC} &\rightarrow \neg(\text{PCORR} \wedge \text{SO}) \\ &\rightarrow (\text{PCORR} \wedge \neg \text{SO}) \vee (\neg \text{PCORR} \wedge \neg \text{SO}) \vee (\neg \text{PCORR} \wedge \text{SO}) \end{aligned} \quad (2)$$

We then see that indeterministic common causes entail at best one of three logically possible options: (i) that the (indeterministic) common cause *does not* satisfy screening-off and its effects *are* perfectly correlated; (ii) that the (indeterministic) common cause *does not* satisfy screening-off and its effects *are not* perfectly correlated; and finally (iii) that the (indeterministic) common cause *does* satisfy screening-off but its effects *are not* perfectly correlated.

²⁴ We are here assuming a two-valued common cause variable V , with $q \in \{q_1, q_2\}$ such that $Vq_1 = C$ and $Vq_2 = \neg C$. See our remarks in Footnote 16.

Van Fraassen’s example refers explicitly to a purely indeterministic common cause of a perfect correlation. So this fits in with our first option i.e. (i) the effects of the (indeterministic) common cause *are* perfectly correlated and the common cause *does not* satisfy screening-off.²⁵ It is hence not surprising at all that the screening-off is violated there. Thus Cartwright is completely right when she claims that “in this case it is not reasonable to expect the probabilities to factor, conditional on the common cause. Momentum is to be conserved, so the cause produces its effects in pairs.”

The claim corresponds to the first part of our quote from Cartwright above. On the face of it, the second part of the quote, endorses a strictly speaking stronger claim. In particular, we are told that the conjunctive fork criterion is not appropriate in this example “because it is a criterion tailored to cases where the cause operates independently in producing each of its effects: whether one of the effects is produced or not has no bearing on whether the cause will produce the other.” Yet, the third logically possible option we derived above entails that, as long as the correlations are not perfect, the conjunctive fork criterion, i.e. screening-off, or (CritCC), may still hold for indeterministic common causes.²⁶ To make the point sharper, what matters most is whether the correlations we want to explain are perfect or not. And a common cause, operating under a constraint —such as conservation of momentum— may well produce imperfect correlations too. This may best be shown by means of the following example.

Let us consider a slight modification of van Fraassen’s ‘bombed atom’ example such as the atom splits in three fractions (instead of two) after being bombarded, which move away at angles θ_1 , θ_2 and θ_3 , each with a given probability. As in the original example, the values that θ_1 , θ_2 and θ_3 may take are also bound by momentum conservation.²⁷ But the correlations $Corr(\theta_1, \theta_2)$, $Corr(\theta_1, \theta_3)$ and $Corr(\theta_2, \theta_3)$ between the three angles need not be perfect any more. That is to say

25 According to our logical implication above, a non-screening-off common cause is the only option left when we insist on both genuine indeterminism and perfect correlation (and this is exactly what goes on in the ‘bombed atom’ example). Just by conjoining equation (2) and PCORR we obtain:

$$\neg DC \wedge PCORR \rightarrow \neg SO$$

26 Again, this can be made clearer if we take the conjunction of expression (2) and $\neg PCORR$, yielding:

$$\neg DC \wedge \neg PCORR \rightarrow (\neg PCORR \wedge \neg SO) \vee (\neg PCORR \wedge SO),$$

27 Conservation of momentum would imply that the sum of the three angles is 2π , or in other words, that the sum of the vectors representing the three trajectories is zero.

$$\begin{aligned}
p(\theta_1 / \theta_2) \neq 1 &\neq p(\theta_2 / \theta_1), \\
p(\theta_1 / \theta_3) \neq 1 &\neq p(\theta_3 / \theta_1), \\
p(\theta_2 / \theta_3) \neq 1 &\neq p(\theta_3 / \theta_2).
\end{aligned}$$

What this shows is that although the ‘conservation’ constraint is sufficient to generate correlations —and to determine their strength— it is not sufficient on its own for the common cause to produce its ‘effects in pairs’, i.e. to generate perfect correlations. So whether perfect correlations arise out of the operations of a common cause operating under a constraint is at least dependent on the number of effects that the common cause ‘generates’, so to speak.

Hence the second part of Cartwright’s quotation above, taken on its own would overstate the link between screening-off and genuinely indeterministic common causes operating under a constraint. For, while such a claim is true for common causes that produce their ‘effects in pairs’, it does not seem correct to say that it generally is true —since probabilistic causes may operate under a constraint without yielding perfect correlations.²⁸

The upshot of this discussion is that examples such as van Fraassen’s ‘bombarded atom’ and the like that involve genuine indeterminism only show that, as a general rule, screening-off is violated if the common cause produces perfectly correlated effects —and whether this is due to the common cause operating under a constraint or not does not seem crucial on its own. The question is then how far we can go with the original Criterion in characterising the remaining cases —i.e. indeterministic causes that only generate imperfect correlations.

3.4. Screening-off and Indeterministic Common Causes are not Incompatible

Our analysis so far has shown that indeterministic common causes may satisfy screening-off, as long as they do not generate perfect correlations. In terms of our explanatory agenda, we may summarise the situation, and the implications of expression (2) as follows. A non-perfect correlation may receive more than one common cause explanation: one in terms of a Reichenbachian Common Cause —which screens-off the correlation— and another in terms of a non-screening-off common cause. A recent school in causal inference, which we review in the next section, goes as far as claiming that that a Reichenbachian Common Cause explanation may be provided for any correlation.

28 To be fair to Cartwright, she seems aware of the problem —her ‘generalised conjunctive forks’ are wide enough to encompass all cases of indeterministic causes, whether they generate perfect correlations or not. See particularly (Cartwright, 1993). But the point is that the reasons she advances for rejecting screening-off as a necessary condition conflate two of the three logical possibilities that open up in response. And this shows in Cartwright’s characterisation of her own generalised criterion for common causes. See (Cartwright, 1987, chapter 6).

To sum up, we may distinguish at least two different kinds, so to speak, of common causes. We have, on the one hand, *Reichenbachian* common causes (RCC), which screen-off their corresponding correlations. These may be either of the deterministic or indeterministic kind. Then there are also indeterministic common causes that operate under a constraint to produce perfectly correlated effects. These are necessarily *non-Reichenbachian* common causes —since they do not fulfil Reichenbach’s Criterion for Common Causes, but a more general criterion at best, such as Cartwright’s generalisation of the fork criterion (see Figure 1).²⁹ Table 1 summarises these facts about the applicability of Reichenbach’s Criterion in deterministic and indeterministic cases.

It is clear that the *Criterion* is valid in very limited circumstances. But note that even Cartwright’s generalised criterion turns out to be unnecessary for common causes. We need still further causal information about the system at hand to know whether this criterion is applicable.³⁰ This favours a pluralism regarding common causes, whereby different kinds of common causes may be appropriate in each particular domain. This has methodological bite, since it entails that information is required regarding what *type* of common causes are appropriate before it can be decided what criterion to apply. The attitude is in line with Cartwright’s celebrated slogan ‘no causes in, no causes out’.

COMMON CAUSES

	Deterministic	Indeterministic
PCORR	CritCC	¬CritCC
non-PCORR	inapplicable	¬CritCC or CritCC

Table 1: The status of Reichenbach’s Common Cause Criterion (CritCC) for perfect correlations (PCORR) and standard non-perfect correlations (non-PCORR).

²⁹ Indeed, different types of common cause may perhaps be distinguished among non-Reichenbachian common causes. Those conforming to Salmon’s interactive forks are different than those characterised by the generalisation of the conjunctive fork advanced by Cartwright (1987). The point remains that the *Criterion* may be a reliable tool under the appropriate set of circumstances, regardless of whether it is defining or only a necessary condition for common causes in general.

³⁰ Suárez (1997, chapter 4) shows that there are relevant cases of perfect correlation where Cartwright’s generalised criterion fails.

4. Common Cause Completability

So far we've been arguing against any sufficient or necessary conditions on common causes in general. Certainly Reichenbach's Criterion for Common Causes (CritCC) —and more particularly screening-off— is neither sufficient nor necessary. Moreover, we find little reason to think that either stronger or weaker characterisations of common causes may be found to provide us with either sufficient or necessary conditions for common causes in general. The best philosophy of causation literature nowadays seems to us to rightly give up on the question: “what is a common cause?” and to direct efforts instead into the most efficient methodologies for causal inference in different domains. From this point of view, what matters is to first determine the kind of common cause that we are seeking, then figure out what the best characterisation may be in a particular domain, then apply the characterisation very selectively. Curiously enough, this seems to us to come as some vindication for Reichenbach's original criterion. Since there is no criterion for common causes in general (i.e. no criterion that would succeed where Reichenbach's failed), we may just as well take the original criterion and analyse its proper domain of applicability. For we have seen that for a certain kind of common causes, the Criterion does work. And no other Criterion does any better work in this particular domain, nor more generally in every domain. Since there is not much hope of a final theory of common causes that could replace the original criterion wholesale, the methodological and explanatory power of the criterion remains intact within its proper domain of applicability.

So, we propose then to return to Reichenbach's Criterion (CritCC), but this time from a purely methodological point of view —and to use it as a lever for investigating its proper domain of applicability. We may then ask not so much whether the criterion is true or false in general, but rather what its boundaries of applicability might be.³¹ In other words, we urge the third option that was logically left open to us in section 3.2., namely: to investigate the validity of CritCC for indeterministic causes that yield non-perfect correlations. This allows us to apply Reichenbach's *Criterion* as a methodological guide in domains where we lack —or are unsure about— our causal intuitions, as it is the case of quantum physics. We can do this not because we know that CritCC must apply there, but rather because we have, methodologically speaking, nothing at all to lose in trying out.

³¹ Thus we agree with Williamson (2005) and adopt CritCC as merely a default methodological rule for causal inference.

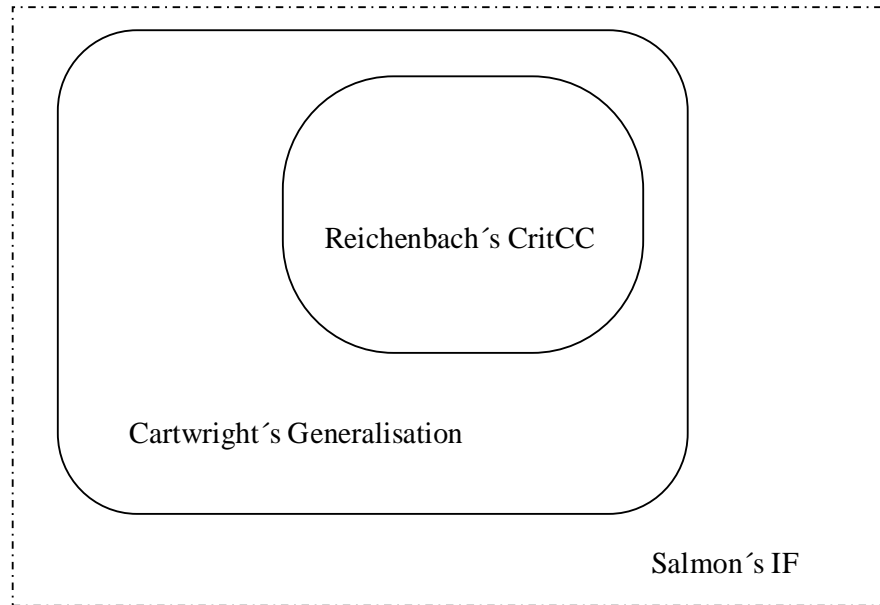


Figure 1: Reichenbachian common causes (RCC), Cartwright's generalisation and Salmon's Interactive Forks (IF).

4.1. Intuitive Motivation for the Existence of Reichenbachian Common Causes

We have learned that for some correlations there might exist different causal explanations, some in terms of screening-off common causes, some not. But we have not so far established whether some of these correlations only admit non-screening-off common causes —so far we have only suggested that perfect correlations in indeterministic contexts seem good candidates. We might then ask whether this is so, i.e. whether there are correlations that cannot be explained at all in terms of screening-off common causes. Put it differently, are there screening-off common cause explanations available to us for *any* given correlation? Or, yet in other words, is it the case that every time a correlation is observed, a common cause C may be 'found'³² which explains the correlation by screening it off?

Let us consider once again the potentially problematic examples, such as van Fraassen's 'bombarded atom', or the notorious 'Cheap-But-Dirty/Clean-And-Green' factory example, due to Cartwright.³³ The 'Cheap-But-Dirty/Clean-And-Green' factory example has been widely discussed in the literature, particularly in the context of debates over the Causal Markov Condition. What is interesting about this example is that the causal structure is presupposed. Hence there can be no doubts that the common cause is precisely what we are told it is in the example. It is further assumed that all relevant causal connections have been taken into account, which amounts to saying

³² The relevant sense of 'found' is not unproblematic —as we discuss in the last section of this work.

³³ Cf. (Cartwright, 1993).

that the system as described is causally complete.³⁴ We present here a simplified version of the example—which shall however be enough to display the main features we would like to discuss.³⁵

Suppose that a chemical factory (the ‘Cheap-But-Dirty’ factory) produces a chemical A through a genuinely probabilistic process. In such a process, the probability of actually getting the product A is 80%. The process misfires on 20% of the occasions. Suppose moreover that whenever the chemical A is produced, another (pollutant) chemical B is also produced as a by-product. The probability for the pollutant B to be produced is then 80% as well. The production process of the factory thus clearly correlates the production of the chemical A with the production of the pollutant B .

Environmental concerns arise when the pollutant is detected. Everything points to the process for the production of the chemical A in the ‘Cheap-But-Dirty’ factory as responsible for the production of the pollutant B as well. In other words, the process seems to originate in a ‘common cause’ C of both the chemical A and the pollutant B . The factory management, however, would not concede that, and defend their innocence with an argument that relies on the screening-off condition. If the process C for the production of the chemical A in the ‘Cheap-But-Dirty’ factory were to be the cause of the production of the pollutant B as well, then conditional on the process, the probability that the chemical and the pollutant are produced together would factorise, i.e. $p(A \wedge B/C) = p(A/C) \cdot p(B/C)$.

But what is really going on, we are told, is that even if the process is running, i.e. even if the common cause C is present, the cause only ‘fires’ 80% of the times. There is no reason why C should screen-off the correlation in this case. This is best seen if we look at the probability space (S, p) , formed by the Boolean algebra S containing the three events and all their possible conjunctions, and the probability measure p that assigns probabilities to each of the elements of S . That is

$$S = \{A, B, C, A \wedge B, A \wedge C, B \wedge C, A \wedge B \wedge C\},$$

and the probability measure p assigns the following probabilities:

34 Completeness is a strong assumption, often unwarranted in the practice of causal inference. But Cartwright employs her example in order to illustrate a conceptual possibility; so contrary to many of her critics, we do not believe that her argument for the possibility of non-screening-off common causes hinges on the validity of completeness in this particular case. However, while admitting with Cartwright that a non-screening-off common cause is *possible* in this scenario, we would like to argue that a failure of completeness supports the view that it might not be *the only* possible cause. The difference is thus that while critics of Cartwright have invoked a failure of completeness in order to rebut her argument for the possibility of non-screening common causes, we use it to back up a form of causal pluralism that accepts the conclusion of Cartwright’s argument, but aims to go further.

35 The detailed description of the example may be found in (Cartwright, 1993, 1999a, 1999b).

$$\begin{aligned}
p(A) &= p(B) = 0.8 \\
p(C) &= 1 \\
p(A \wedge B) &= 0.8 \\
p(A \wedge C) &= p(B \wedge C) = 0.8 \\
p(A \wedge B \wedge C) &= 0.8
\end{aligned}$$

The above assumes first that (S, p) contains *all* causal influences, i.e. is causally complete, as required in the example. We may also assume for the sake of simplicity that the common cause C is always present. Thus, since C only produces its effects 80% of the time, the probabilities for $A \wedge C$ and $B \wedge C$ are both 0.8. If, in addition, we assume (again for simplicity) that in the absence of the cause none of the chemicals A or B are produced, we get that the probability that A will occur is 0.8, and similarly for B . The same applies to $A \wedge B \wedge C$. The actual numbers are not so important. The important fact is that the model reproduces the example's features. In particular, we have that

$$\begin{aligned}
p(A \wedge B / C) &= 0.8 \\
p(A / C) &= 0.8 \\
p(B / C) &= 0.8
\end{aligned}$$

It becomes now clear that screening-off is not satisfied:

$$p(A \wedge B / C) \neq p(A / C) \cdot p(B / C).$$

However, C is by construction the common cause of both A and B . So we can conclude that *the common cause C does not screen-off the correlation*. We have emphasised the article 'the' in the sentence above to stress that C is *the only* possible candidate for common cause in the probability space (S, p) . Recall that this follows from the assumption that the probability space (S, p) is causally complete, i.e. that (S, p) contains all possible causes of A and B .

We have already seen that Cartwright's response to this problem is to weaken the criterion that characterises our common causes within (S, p) . This is certainly possible. The ensuing common cause is indeterministic and fails to obey screening-off. However, there is logically another possible move, perhaps as natural as Cartwright's own, that also respects the intuition that a common cause explanation should be found. In a sense, the alternative is 'just a matter' of rewriting the italicised sentence above as "*a common cause C does not screen-off the correlation*". The indefinite article suggests that there might exist, under the right conditions, *another* event C' which is also a common

cause of the correlation and which does screen it off.

The ‘just a matter’ however is not such a simple matter. It requires us to drop the assumption that the probability space (S, p) contains all relevant causal variables. In other words, we need to assume now that (S, p) is causally incomplete. But as we already pointed out, completeness is quite a strong assumption. There are few practical cases, if at all, in which completeness may be warranted. We rarely are in a position to know whether a cause is in fact a total cause of its effects, or just a partial cause. These considerations partly underlie the intuition that more detailed causal explanations are always possible, even in Cartwright’s example.

4.2. Extensibility and Common Cause Completability

One reason to search for alternative causal explanations is that different explanations may be required for different purposes. In some cases a given causal explanation might just not be good enough. From the explanatory point of view assumed in this paper, Reichenbach’s Criterion affords a particular form of explanation, independently of its ontological status. So it makes sense to look for Reichenbachian common causes, wherever they may be found, by expanding the original probability space and considering further variables. The most important recent attempt to pursue this road thoroughly is the work of the so-called ‘Budapest School’ on the extensibility of probability spaces for Reichenbachian common causes.³⁶

A probability space (S, p) that does not contain a screening-off common cause of the correlation $Corr(A, B)$, is said to be *Reichenbachian common cause incomplete*.³⁷ As pointed out in the previous section, one seems to have two alternatives when facing a Reichenbachian common cause incomplete probability space. We may opt, following Cartwright, for a weakening of the common cause criterion. This allows us to identify some variables already in the space as the common causes. Alternatively, we may seek a screening-off common cause, by rejecting completeness and expanding the original probability space to include further variables. The ‘Budapest School’ provides us with a definition of the *extension* of a probability space:

Definition 1 (Extension) *The probability space (S', p') is called an **extension** of (S, p) if there exist a Boolean algebra embedding h of S into S' such that $p(X) = p'(h(X))$, for all $X \in S$.*

36 Cf. (Hofer-Szabó, Rédei and Szabó, 1999, 2000).

37 This is just a little bit different from the original idea —see (Hofer-Szabó, Rédei and Szabó, 1999)— which referred to such probability spaces simply as ‘common cause incomplete’. We have introduced the term ‘Reichenbachian’ to make room for any non-Reichenbachian common causes —since we have already made clear that we also find these kinds of common causes perfectly acceptable (see Footnote 34).

Extensibility, as expressed above, allows then for the enlargement of the original probability space so that new events are included. Moreover, Definition 1 ensures that the extension operation is consistent with the old event structure (S, p) . Note that in extending a probability space (S, p) into (S', p') we are not only enlarging the set of events S —this is the role of the embedding h in Definition 1— but also the probability measure p needs to be altered. Thus the embedding h needs to be defined such that the initial probabilities and correlations are maintained under the new probability measure p' . In other words, correlations should stay invariant under the extension operation, that is $Corr(A, B) \equiv Corr_p(A, B) \equiv Corr_{p'}(A, B)$. That the extension operation is consistent in this sense is crucial for our purposes, since we are asking whether a Reichenbachian common cause of the *original* correlations exist in the extended space.

The definition of extension above is in principle applicable to any probability space. But it does not establish how many different extensions exist of a given probability space (S, p) . Nor does it tell us under what circumstances such an extended probability space (S', p') will contain the Reichenbachian common causes we are looking for. It does not even guarantee that there will be one extension that will include Reichenbachian common causes. In order to address the issue we need to introduce the idea of Reichenbachian common cause completability (RCC Completability):

Definition 2 (RCC Completability) *Let $Corr(A_i, B_i) > 0$ ($i = 1, 2, \dots, n$) be a set of correlations in (S, p) such that none of them possess a common cause in (S, p) . The probability space (S, p) is called **Reichenbachian common cause completable** with respect to the set $Corr(A_i, B_i)$ if there exists an **extension** (S', p') of (S, p) such that $Corr(A_i, B_i)$ has a Reichenbachian common cause C_i in (S', p') for every $i = 1, 2, \dots, n$.*

Completability, as defined above, is hence the key for successfully searching Reichenbachian common causes of any given correlation.³⁸ The question is now whether any *incomplete* probability space (S, p) can be extended such that is (Reichenbachian) common cause completable. In other

38 Note that although all definitions and propositions presented here refer to classical probability spaces, similar work can be developed for von Neumann spaces as well, giving completely equivalent results to those presented here for the quantum mechanical case. In fact, these results are largely motivated by the possibility to provide common cause explanations of EPR quantum correlations. On the other hand, these results refer only to a *finite* set of correlated events $i = 1, 2, \dots, n$. The question whether equivalent results may be obtained for an *infinite* set of correlated events is still open.

words, when is a probability space (S, p) Reichenbachian common cause completable?

Hofer-Szabó et al. show that an extension (S', p') may *always* be found for a Reichenbachian common cause incomplete probability space such that it contains (Reichenbachian) common causes for all the original correlations.³⁹ In other words:⁴⁰

Every classical probability space (S, p) is common cause completable with respect to any *finite* set of correlated events.

This result seems already quite promising for our own purpose. However, completable is relative to a specific set of correlations. In other words, even if Reichenbachian common cause completeness is achieved for a particular set of correlations —by extending the original probability space— the new extended probability space may be incomplete with respect to another set of correlations. We may then go on to ask whether there exists an extension of (S, p) such that the resulting probability space (S', p') is Reichenbachian common cause complete not only with respect to the original correlations in (S, p) but also with respect to all other correlations that might have arisen as a result of the extension operation. Such an extension is said to be a Reichenbachian common causally closed probability space.⁴¹ This is an interesting issue, which again involves the assumption of completeness. We will not broach it here, and will focus on the simplest form of Reichenbachian completable.

4.3. Indeterministic Reichenbachian Common Causes

The results in the previous section suggest that a Reichenbachian common cause explanation may be provided for any given correlation, as long as we drop an assumption of completeness about the causal structure of the original probability space where the correlation is observed. This may seem to restore our original, most powerful tool for causal inference. In particular, we can now test how Reichenbachian common cause explanations fare in any situation where correlated (and non-directly causally related) events are found. But of course, as pointed out in the previous section, a Reichenbachian common cause explanation may not be the only possible common cause explanation. There may be explanations that appeal to non-Reichenbachian common causes. How do these compare?

We partly addressed the question in Section 3. The ‘Budapest School’ would agree with Salmon

39 See (Hofer-Szabó, Rédei and Szabó, 1999) for the details of the proof.

40 (Hofer-Szabó, Rédei and Szabó, 1999, p. 384).

41 Cf. (Gyenis and Rédei, 2004).

and Cartwright as regards the validity of Reichenbach's Criterion in deterministic contexts. But their reasons are different. For, while Salmon and Cartwright would consider the case to be a particular instance of interactive forks and *generalised* forks respectively, proponents of the 'Budapest School' would take it that any appropriate extension would provide for such a deterministic common cause.

Indeterministic common causes however are a different issue, as we have seen. Let us consider again the 'Cheap-but-Dirty/Green-and-Clean' factory example. Recall that the motivation for such an example was to show that Reichenbach's Criterion for Common Causes did not hold for genuinely probabilistic causes. For such cases, Cartwright suggested accepting irreducibly probabilistic causation and to correspondingly weaken Reichenbach's criterion for common causes. An alternative in practice is to reject the completeness assumption, and continue the search for Reichenbachian common causes. The extensibility results in the previous section show that at least formally this is always possible to do. For what they show is that there is a way of formally extending any probability space so that a Reichenbachian common cause may ultimately be accommodated. Now what is clear is that in the particular case of the factory example such Reichenbachian common cause must be deterministic since the production of the chemical *A* and that of the by-product *B* are perfectly correlated. And we already know that a Reichenbachian, screening-off, two-valued common cause explanation of perfect correlations is also deterministic. However, extensibility results apply to imperfect correlations too, and there is no need to suppose in those cases that the ensuing common causes will be deterministic.

So far so good —Reichenbach's Criterion seems to be vindicated in practice. However, there are reasons why such arguments in favour of Reichenbachian common cause explanations are not compelling, even in the indeterministic case. The extensibility and common cause completability results we have reviewed so far are entirely formal. They teach us how to formally extend an original probability space so as to accommodate a variable that obeys the statistical characterisation for common causes endorsed by the Criterion. This is the only sense in which they may be said to restore Reichenbachian common cause explanations. But they do not tell us that such a Reichenbachian common cause exists in reality. That is, these results allow us to formally accommodate a screening-off variable within the probability space, but they do not guarantee that such variables represent anything at all in reality. This suggests that what we called 'common cause completability' —following the original terminology of the 'Budapest School'— is merely 'screening-off completability'. In other words, what the extensibility theorems show is merely that it is always possible to extend the original probability space in such a way that a *screener-off* will appear in the new expanded space. They neither tell us that such a screener-off will be found to represent a real event, nor that such events will happen to be genuine common causes. And we

already learnt that under no proposal on the table is the screening-off condition sufficient for common causes. So the formal results of the ‘Budapest school’ are indecisive when it comes to establishing or inferring *real* common cause structure for any correlations.

This seems to us an important and far-reaching objection to the claims by the ‘Budapest School’. It can be formulated in an alternative fashion as follows. The extensibility and completeness results do not guarantee a unique extension of the original probability space. On the contrary, it seems plausible to think that there can be in general several such extensions and therefore different putative Reichenbachian common causes of a given correlation. The physical interpretation of a Reichenbachian common cause resulting from an extension of a given probability space will thus depend on the choice of extension. The question to ask then is not simply whether there exists an extension of the original probability space such that it contains a screener-off for the correlation we wish to explain. Rather, we want to know whether there exists an extension of the original probability space such that it contains a *physically interpretable* Reichenbachian common cause of the correlation. This is far from a trivial issue, which the Budapest school has not even begun to address —but which is absolutely decisive for causal inference.

GENUINELY PROBABILISTIC CAUSATION (INDETERMINISM)

		CritCC	\neg CritCC	
				Interact. Fork
		General. Fork		
PosCC	‘Budapest School’	Cartwright	Salmon	
\neg PosCC		van Fraassen		

Table 2: The positions of van Fraassen, Salmon, Cartwright and the ‘Budapest School’ towards both the Postulate of the Common Cause (PosCC) and Reichenbach’s Criterion for Common Causes (CritCC) as regards to genuinely probabilistic (indeterministic) causation.

As a summary of our discussion, Table 2 contains the views so far reviewed as regards indeterministic common causes. It displays the main differences between Salmon, Cartwright and the ‘Budapest School’ regarding the status of Reichenbach’s Postulate (PosCC) and Criterion

(CritCC) for Common Causes. The ‘Budapest School’ maintains that the PosCC and CritCC hold in both deterministic and genuinely indeterministic contexts, including quantum mechanics. This is very much in opposition to van Fraassen who seems to reject both. On the other hand, Cartwright and Salmon accept the validity of the Postulate, and correspondingly the Principle (PCC), but only under a revision of (CritCC). They differ on the kind of revision they recommend since Salmon rejects screening-off altogether and defends interactive forks, while Cartwright defends a criterion weaker than Reichenbach’s but more restrictive than Salmon’s interactive forks. To sum up, with the exception of van Fraassen, all the other proposals discussed —Salmon’s, Cartwright’s and the ‘Budapest School’— accept the metaphysical claim in Reichenbach’s Principle of the Common Cause, i.e. the Postulate of the Common Cause (PosCC),⁴² but they differ on the methodological claim, which some find acceptable (the ‘Budapest School’) and some reject (Cartwright, Salmon).

5. Conclusion: Pluralism and Explanation

Reichenbach’s Criterion for Common Causes is neither a sufficient nor a necessary condition on common causes. While it is true that from the formal point of view ‘Reichenbachian common causes’, understood as formal variables satisfying screening-off conditions, can be accommodated within a given probability space for any given correlation, the question remains whether such formal variables represent anything physically real, and, if so, causally relevant. So the formal results are ontologically inert, since they do not guarantee the reality of any putative cause for any correlation. And the debate over causal inference in the natural and social sciences entirely turns on the reality of the causes that we supposedly infer to. It follows that the results of the ‘Budapest School’ are of limited relevance so far to the methodology of causal inference. However, these formal results may at least restore some confidence in well known methodological rules of thumb, such as screening off, and we have argued that this is worthy in its own terms, even if it were lacking ontological import. We see no reason not to adopt the Budapest extensibility and completeness results as rules of thumb, since Reichenbach’s criterion is no worse off than any other alternative. As a matter of methodological principle then, we propose that any correlation be taken to have a plurality of possible ‘causes’. Each of these causes is discoverable by different methods; and these methods may only be determined by context and background causal knowledge.

We also maintain that ‘Reichenbachian common causes’ (i.e. screener-offs), where and if they

⁴² And so ultimately they maintain the validity of the Principle (PCC). A further view that ought to be discussed is Sober’s (2001), who notoriously rejects the PCC on account of putative failures of the *Postulate*. It is not entirely clear to us that Sober maintains the *Criterion* —if he does then he would occupy the empty slot in the bottom left hand corner of the diagram.

obtain, are explanatory. Again we see no reason to give up on this general claim, originally made by Reichenbach himself. Explanation does not require contextual or methodological unanimity.⁴³ Obviously such ‘Reichenbachian’ explanations will not always be available, or might coexist with other causal explanations. So our view entails that explanations for correlations are not unique, but depend on the objective particular purposes of the explanation, and the context of the correlation. So we are embracing both methodological and explanatory pluralism.⁴⁴ One question that opens up then is how ‘Reichenbachian’ explanations fare in relation to other possible causal explanations. Is there a ranking in the explanatory power of different ‘common cause’ explanations? This is a question that would have even failed to make sense in the old times of ‘necessary and sufficient conditions’ for common causes. But in the present and rather different *zeitgeist* of ‘causal pluralism’ the issue becomes urgent —and we urge philosophers to pursue it.

⁴³ The fact that there can be different and even contradictory explanations in different contexts for the same phenomenon does not impugn their explanatory power. This presupposes a form of explanatory pluralism that seems to us to be becoming the norm. For one excellent defence, albeit applied to the notion of understanding, see De Regt and Dieks (2005).

⁴⁴ But emphatically not *ontological* pluralism á la Cartwright (1999b) —the contemporary manifesto for *dappled world* causal metaphysics.

References

- Cartwright, N. (1987) “How To Tell a Common Cause: Generalizations of the Conjunctive Fork Criterion”. In J. H. Fetzer (Ed.), *Probability and Causality*, pp. 181–8. Reidel Pub. Co.
- Cartwright, N. (1989) *Nature’s Capacities and Their Measurement*. Clarendon Press
- Cartwright, N. (1993) “Marks and Probabilities: Two Ways to Find Causal Structure”. In F. Stadler (Ed.), *Scientific Philosophy: Origins and Developments*, Yearbook 1/93, Institute Vienna Circle. Kluwer Academic Publishers
- Cartwright, N. (1999a) “Causal Diversity and the Markov Condition”. *Synthese*, **121**: 3–27
- Cartwright, N. (1999b) *The Dappled World: A Study of the Boundaries of Science*. Cambridge University Press
- Cushing, J. and McMullin, E. (Eds.) (1989) *Philosophical Consequences of Quantum Theory*. University of Notre Dame Press
- De Regt, H. and Dieks, D. (2005), “A Contextual Approach to Scientific Understanding”. *Synthese*, **144**: 137-170
- Fine, A. (1982) “Joint Distributions, Quantum Correlations, and Commuting Observables”. *Journal of Mathematical Physics*, **23**: 1306–10
- Graßhoff, G., Portman, S. and Wüthrich, A. (2005) “Minimal Assumption Derivation of a Bell-type Inequality”. *The British Journal for the Philosophy of Science*, **56**: 663–80
- Gyenis, B. and Rédei, M. (2004) “When can Statistical Theories be Causally Closed?” *Foundations of Physics*, **34**: 1285–303
- Hofer-Szabó, G., Rédei, M. and Szabó, L. E. (1999) “On Reichenbach’s Common Cause Principle and Reichenbach’s Notion of Common Cause”. *The British Journal for the Philosophy of Science*, **50**: 377–99
- Hofer-Szabó, G., Rédei, M. and Szabó, L. E. (2000) “Common Cause Completeness of Classical and Quantum Probability Spaces”. *International Journal of Theoretical Physics*, **39**: 913–9
- Reichenbach, H. (1956) *The Direction of Time*; Edited by Maria Reichenbach. Unabridged Dover (1999) republication of the original University of California Press (1956) publication
- Sober, E. (2001), “Venetian Sea Levels, British Bread Prices, and the Principle of the Common Cause”. *The British Journal for the Philosophy of Science*, **52**:331–46
- Suárez, M. (1997) *Models of the World, Data Models and the Practice of a Science: the Semantics*

of Quantum Theory. Ph.D. thesis, University of London, LSE

- Suárez, M. (2007) “Causal Inference in Quantum Mechanics: A Reassessment”. In F. Russo and J. Williamson (Eds.), *Causality and Probability in the Sciences*, pp. 65–106. London College
- Suárez, M. and I. San Pedro (2007), “EPR, Robustness and the Causal Markov Condition”, LSE Philosophy Papers PP/04/07
- van Fraassen, B. C. (1982a) “The Charybdis of Realism: Epistemological Implications of Bell’s Inequality”. *Synthese*, **52**: 25–38. Reprinted with corrections in (Cushing and McMullin, 1989, pp. 97–113)
- van Fraassen, B. C. (1982b) “Rational Belief and the Common Cause Principle”. In R. McLaughlin (Ed.), *What? Where? When? Why? Essays in Honour of Wesley Salmon*, pp. 193–209. Reidel
- Williamson, J. (2005) *Bayesian Nets and Causality: Philosophical and Computational Foundations*. Oxford University Press