



Centre for Philosophy of
Natural and Social Science
LSE Philosophy Papers

EPR, *Robustness* and the Causal Markov Condition

Mauricio Suárez and Iñaki San Pedro

The London School of Economics
and Political Science
Houghton Street
London WC2A 2AE

Tel: 020 7955 7573
Fax: 020 7955 6869
email: m.steuer@lse.ac.uk



Editor

Max Steuer

Editorial Board


Roman Frigg

Alex Voorhoeve

John Worrall

PP/04/07

ISSN 1750-3604



EPR, *Robustness* and the Causal Markov Condition

Mauricio Suárez and Iñaki San Pedro *

19/08/07

Abstract

It is still a matter of controversy whether the Principle of the Common Cause (PCC) can be used as a basis for sound causal inference. It is thus to be expected that its application to quantum mechanics should be a correspondingly controversial issue. Indeed the early 90's saw a flurry of papers addressing just this issue in connection with the EPR correlations. Yet, that debate does not seem to have caught up with the most recent literature on causal inference generally, which has moved on to consider the virtues of a generalised PCC-inspired condition, the so-called Causal Markov Condition (CMC). In this paper we argue that the CMC is an appropriate benchmark for debating possible causal explanations of the EPR correlations. But we go on to take issue with some pronouncements on EPR by defenders of the CMC.

Contents

| | |
|--|----|
| 1. <i>Introduction</i> | 2 |
| 2. <i>EPR and Quantum Correlations</i> | 4 |
| 3. <i>Redhead's Robustness</i> | 5 |
| 4. <i>Healey on Robustness</i> | 8 |
| 5. <i>The Causal Markov Condition</i> | 10 |
| 6. <i>Robustness and the Causal Markov Condition</i> | 11 |
| 7. <i>EPR and the Causal Markov Condition</i> | 14 |
| 8. <i>Conclusions</i> | 22 |
| <i>Bibliography</i> | 24 |

* Research towards this paper has been funded by project HUM2005-01787-C01-03 of the Spanish Ministry of Education and Science. We would like to thank the members of its associated 2005 reading group on causal inference, as well as three referees for helpful comments and suggestions.

1 Introduction

Questions regarding the status of causation in quantum mechanics are as ancient as the discipline itself. The founding parents of quantum mechanics often identified causation with determinism and consequently understood the emergence of the fundamentally probabilistic quantum mechanics as the demise of a causal picture of the world. As a consequence quantum theory is often presented as non-causal.¹ The identification of causality and determinism was rather universal: even those who regretted the demise of a causal picture attempted to restore a causal understanding of quantum mechanics precisely by restoring determinism. For instance, David Bohm showed von Neumann's theorem against hidden variables to involve essentially questionable premises, thus paving the way for hidden variables. But while Bohm and von Neumann disagreed regarding the status of causation in quantum mechanics, they agreed that the fortunes of causation and determinism were essentially linked. Bohm's theory is in essence a programme to endow quantum mechanics with an underlying deterministic dynamics.

The identification causality = determinism (let us call it the "*c=d identity*") has continued in different, not always explicit, guises. For example Bell's theorem and the work leading up to it during the 1960's presupposes the notorious *factorizability* condition as a criterion of local causality. *Factorizability* is applicable to the correlations between measurement outcomes of spatially separated systems in EPR-like set-ups. Bell's theorem demonstrated that no "factorizable" theory can reproduce the quantum correlations. It is thus concluded that Bell's work shows that not only quantum mechanics but any other empirically indistinguishable theory would be non-causal in this sense. But philosophers have shown that Bell's theorem does not entail a departure from the *c=d identity*. Some brilliant work by philosophers of physics in the early 1980's showed that the *factorizability* condition implies determinism when applied to the EPR perfect anti-correlations.² So in the end it turns out that the rejection of local causality promoted by Bell's theorem also presupposes a rejection of determinism, and is hence compatible with the *c=d identity*.

Many physicists have continued to presuppose the *c=d identity*, sometimes unquestioningly so. Philosophers of science by contrast long ago started to work out the details of a stochastic view of causality. On this view causation is

1 See (Heisenberg, 1958) and (von Neumann, 1955).

2 See (Fine, 1982*b*, Fine, 1982*a*) and (van Fraassen, 1982). The original theorems are due to Suppes and Zanotti (1981).

essentially probabilistic association, and hence supposedly divorced from determinism. One of the earliest and most influential attempts is Hans Reichenbach's *The Direction of Time*, where the Principle of the Common Cause (PCC) is first stated. The programme gains its full and most developed expression in Patrick Suppes' epoch-making 1970 book, *A Probabilistic Theory of Causality*. In spite of the fact that these were both explicit attempts at building a stochastic theory of causality, it remains controversial just how much they depart as a matter of fact from the *c=d identity*. In particular regarding the PCC some philosophers have gone on to argue that the assumption of screening-off is only valid for deterministic, or quasi-deterministic common causes, but does not hold for probabilistic causes. Hence philosophers have for a very long time now considered that the *c=d identity* is controversial, although they have disagreed among themselves as to whether it should be rejected altogether, or weakened in some interesting sense.³

The disagreement over rejection vs. weakening goes a long way to explaining why the status of causation in quantum mechanics also remains controversial. A weak version of the *c=d identity* is at the heart of a condition that was widely discussed among philosophers of physics in the early 1990's in connection with the Einstein-Podolsky-Rosen correlations, namely Michael Redhead's *robustness*. The consensus reached then was by and large that *robustness* is too strong a condition on probabilistic causality. So the failure of *robustness* in the EPR set-up is uninformative, and a causal account of the EPR correlations remains an open option.

The current debate on causal inference has moved to a discussion of the Causal Markov Condition (CMC), a sophisticated version of the PCC for directed acyclic graphs.⁴ The condition employs a similarly weak version of the *c=d identity*, and remains equally controversial. But it has not been systematically applied to the EPR case, nor has the connection been made explicit to the *robustness* condition discussed in the early 1990's. Our main aim in this paper is to make an explicit link between CMC and *robustness* in the context of the EPR correlations. Thus we aim to show that the application of CMC to the EPR correlations is exactly as informative (or uninformative, depending on taste) as *robustness*. Both conditions hold or fail for the same types of systems. So a defender of the weak version of the *c=d identity* will find the failure of both CMC and *robustness* in the EPR correlations revealing of a striking failure of

3 And several philosophers have gone as far as to defend that causality and determinism in fact exclude each other. See (Hofer, 2004) for a recent example.

4 Cf. (Hausman and Woodward, 1999), (Cartwright, 2002) and (Steel, 2005).

causality in quantum mechanics —and there is a sense in which this result vindicates the founding parents’ suspicion that the probabilistic nature of quantum mechanics is what underlies the failure of causality. But those who are inclined to reject the *c=d identity* altogether are likely to draw rather the opposite lesson: the failure of CMC and *robustness* is precisely what is to be expected given the probabilistic nature of quantum mechanical causation. Although there is thus no essential superiority, in the context of the EPR correlations, to discussing CMC over *robustness* we aim to show that the application of CMC is sharper and less cumbersome. Thus we shall urge that the debate over the causal status of the EPR correlations is best continued in the new terms laid down by the Causal Markov Condition.

2 EPR and Quantum Correlations

Einstein, Podolsky and Rosen first introduced the so-called EPR thought experiment in 1935⁵ as an argument to suggest that the then young and emergent quantum theory did not provide a complete description of reality. In a later refined version presented by David Bohm, two *entangled* electrons are emitted from a source in opposite directions. The spin component of each of the electrons can be later detected (measured) when the electrons hit a fluorescent screen after having passed through an inhomogeneous magnetic field (produced by a Stern-Gerlach magnet).

Several features of this experiment are potentially relevant. First, we will denote by a and b the value of the spin variable of each electron which, in the singlet state, can be either ‘spin-up’ (\uparrow) or ‘spin-down’ (\downarrow) with probability 1/2. We can then denote the corresponding measurement outcome events on each particle as $\uparrow a$, $\downarrow a$, $\uparrow b$ and $\downarrow b$. Second, it is assumed that the state of the entangled electron pair is the *singlet* state:

$$\Psi = \frac{1}{\sqrt{2}}(|\uparrow_a\rangle|\downarrow_b\rangle - |\downarrow_a\rangle|\uparrow_b\rangle).$$

Third, it is assumed that measurement events at each wing of the experiment, such as $\uparrow a$, and $\downarrow b$, are space-like separated events, i.e. lie outside each other’s light cone. This is best represented in the diagram of Figure 1 as the statement that no time-like world-line can reach from b to a or *vice versa*. Under a

5 (Einstein, Podolsky and Rosen, 1935).




Figure 1: Spacetime representation of a typical EPR experiment.

conventional albeit controversial interpretation of special relativity, such events can not be causally connected.⁶

Quantum mechanics allows us to calculate conditional and joint probabilities for the different possible outcomes on both wings. When those calculations are performed on the entangled pair in the singlet state, correlations between these outcomes are derived. The EPR correlations between the different outcome events in both wings of the experiment can be succinctly expressed as:

$$p(a \wedge b) \neq p(a) \cdot p(b).$$

These are the EPR correlations, which have been often positively tested in experiment, and for which we would like to know whether they are the result of underlying causal processes, and which processes. An attempt to determine an answer to these questions was carried out in the late 1980's by the distinguished British philosopher of physics Michael Redhead.

3 Redhead's *Robustness*

Redhead introduced his *robustness* condition in 1987⁷ in order to argue that no *direct* causal relation could be established between the outcome events of an EPR-type experiment. The claim was part of Redhead's attempt at showing that

⁶ See (Maudlin, 1994) for a critical discussion.

⁷ (Redhead, 1987).

quantum mechanics and relativity can peacefully coexist. Under the presumption that only timelike related events can be causally connected, the measurement outcome events a and b in an EPR experiment can not be causally connected. In particular, Redhead suggested that the EPR correlations were not what he called *robust causal connections*. This in turn entitled him to discard direct causal links between EPR correlated events:⁸

“A stochastic causal connection between two physical magnitudes a and b pertaining to two separated systems A and B is said to be robust if and only if there exist a class of sufficiently small disturbances acting on B (A) such that b (a) screens off a (b) from these disturbances.

Denoting the disturbance action on B by d , then the first part of this condition can be rendered formally as

$$\exists D \left(\forall d \in D \left[p(a=\varepsilon_a | b=\varepsilon_b \wedge d) = p(a=\varepsilon_a | b=\varepsilon_b) \right] \right)$$

A similar condition can be written down for disturbances acting on A . The requirement of *robustness* as a necessary condition for a causal relation means that sufficiently small disturbances of either relata do not affect the causal relation.”

The intuition that underlies Redhead’s *robustness* is both simple and powerful, and is best brought out by a simplified version of the condition. First, let us simplify Redhead’s terminology by identifying physical quantities and the corresponding events. Typically A , B denote a quantity (variable), while a , b denote a value of the corresponding quantity. Hence A , B are capable of entering in causal relations, while a , b are capable of standing in probabilistic relations. However, for convenience, and without loss of generality, we will run them together. Thus a , b will denote indistinctly the quantities and their values; which is which should be clear from the context. We will then say in general that a stochastic causal link between two quantities a and b is *robust* if and only if the statistical relation $p(a|b)$ is invariant under small disturbances d acting on the putative cause b . In other words b is a robust cause of a if and only if $p(a|b \wedge d) = p(a|b)$. We can see that the intuition behind Redhead’s *robustness* is that it does

⁸ (Redhead, 1987, pp. 102-3).




Figure 2: Redhead's *Robustness* for EPR correlations.

not matter to the causal link between b and a how the putative cause b comes about, only that it does so (see Figure 2).

It is worth mentioning that initially Redhead apparently took *robustness* to be both necessary and sufficient for a causal link, but in response to criticism he weakened this to a necessary condition only.⁹ In any case *robustness* is understood to be at least a necessary condition on a causal link. So it becomes superfluous to speak of a *robust* causal link, since no link that fails to be robust can on this understanding be causal: there is no such a thing as a non-robust causal link. The double terminology points already to what will be the heart of the problem. For Redhead defines *robustness* as a statistical condition. Hence “robust causal link” is really a heterogeneous combination of a statistical condition and a causal relation. In stating that *robustness* is a necessary condition on a causal link Redhead is stipulating that the presence of the causal relation always necessarily implies the statistical condition. So there is a necessary statistical consequence of the existence of a causal relation. As we shall see the critics of *robustness* quickly pointed out that the statistical condition was not general enough to cover all kinds of probabilistic causes, but rather entailed a particular pseudo-deterministic assumption on the working of the cause. The situation is entirely analogous in the recent debate over the Causal Markov Condition.

⁹ For a discussion see (Healey, 1992b).

4 Healey on *Robustness*

The publication of Redhead's work on *robustness* attracted considerable attention and gave rise to an engaging debate on causation in quantum mechanics in general and in EPR in particular. One of the staunchest critics of *robustness* is Richard Healey, who discussed and criticised the condition at length in two papers published in the early 1990's.¹⁰ In these papers Healey cast doubts upon the validity of *robustness* as a criterion for causal inference. His arguments are designed to show that *robustness* is not a necessary condition in general for a causal link. Our thesis in this paper is that in the context of the EPR correlations the debate over the Causal Markov Condition recapitulates the debate over *robustness*, so it is worth reviewing the latter in a little detail.

Healey first pointed out that *robustness*, as defined by Redhead (see Section 3), can only be taken to be a necessary condition on *total* causes. In other words, *robustness* can only be a necessary condition on a causal link between b and a as long as no other causes are operating on a (see Figure 2).¹¹

In order to deal with cases in which b is only a *partial* cause of a , Healey introduced a new condition, which he called *internal robustness*:¹²

“A stochastic relation between two events h, k is *internally robust* just in case $p(h|k)$ is invariant under all (sufficiently small) modifications in the causal antecedents of k that leave k fixed and preserve independent causal antecedents of h .”

We may rephrase this condition in our terminology as follows. A stochastic causal link between a and b is *internally robust* if and only if the statistical relation $p(a|b)$ is invariant under small disturbances d which leave b and all the independent causal antecedents of a fixed. That is a stochastic causal link between quantities a and b is *internally robust* iff $p(a|b \wedge d \wedge c) = p(a|b \wedge c)$ and d does not causally affect c , where c is the set of all independent causal antecedents of a .

¹⁰ (Healey, 1992a, Healey, 1992b).

¹¹ We will not here assess this claim, since the aim of the paper is not to evaluate but to compare *robustness* and the Causal Markov Condition, and to show that they face similar difficulties and challenges.

¹² (Healey, 1992b, pp. 183-4).




Figure 3: Healey's *Internal Robustness*.

Healey finds both conditions problematic as criteria for causal inference: *robustness* is problematic because we are very rarely in a position to know that b is the total cause of a , and so a violation of *robustness* in practice will say nothing informative about whether or not there is a direct causal link between a and b . *Robustness* may fail because b does not cause a , but it may also fail because we are not accounting for a third partial cause c of a . More specifically in the quantum case it is impossible to know whether the measurement outcome event b in one wing is the only cause of the measurement outcome event a in the other wing, and hence impossible to know in advance whether a failure of *robustness* implies no causal relation between b and a or is simply due to the presence of third causes. Similarly for *internal robustness*: we are never in a position to know whether the small disturbances on b have in fact no causal effect upon some of the causal antecedents of a other than b . So as a criterion for causal inference *internal robustness* is just as unhelpful: a failure of *internal robustness* might mean that b is not even a partial cause of a , but it might also mean that there are other unaccounted partial causes of a besides b that are in turn effects of causes of b .

Redhead's response to these criticisms was to assert that "at some stage in the process of incorporating antecedents in the total cause, *robustness* must be rescued. Otherwise we would live in a 'marshmallow' world where the notion of cause would not, I believe, be appropriate."¹³ In other words, whatever our cognitive and epistemic limitations, a causal relation is properly causal only if robust in actual fact when all other causes have been accounted for. So in other

¹³ Cf. (Redhead, 1987, p. vi).

words Redhead’s most considered view is that while *robustness* is not helpful in general as a criterion for positive causal inference, its failure nonetheless allows some minimal negative causal inference. In the EPR case this allows him to say at least that a failure of *robustness* between the outcome events on both wings *b* and *a* definitely implies that *b* is not the *total* or only cause of *a*. Redhead identified robust causality with action at a distance, and distinguished it from what Abner Shimony¹⁴ called *passion at a distance*, a kind of nomic acausal stochastic link between variables that are ‘holistically’ implicated —whatever that might mean. The application of the *robustness* condition to the EPR set-up was designed to show that quantum phenomena exhibits *passion* rather than *action* at a distance. This in turn was argued to be enough to warrant peaceful coexistence with special relativity.

The critics of *robustness* did not rest content at this point however, but went on to argue against *robustness* as a necessary condition on causation in general.¹⁵ In other words, they came to dispute the very idea that causal links are Markovian in the way specified by either *robustness* or *internal robustness*. However, our aim in this paper is not to evaluate *robustness* and *internal robustness* as necessary conditions on total and partial causes; so we will not review this debate. Our main aim here is to more modestly show that *robustness* in the EPR case follows logically from the Causal Markov Condition. We consequently argue that discussion regarding causality in EPR is best conducted in terms of the CMC.

5 The Causal Markov Condition

The Causal Markov Condition (CMC) is inspired by the Principle of the Common Cause (PCC) and is a keystone and crucial assumption in the most powerful contemporary programs of causal inference. It is intended as a generalised version of the PCC and can be defined, following Hausman and Woodward, as follows:¹⁶

14 Cf. (Shimony, 1984).

15 Cf. (Healey, 1992*a*, Healey, 1992*b*) and (Cartwright and Jones, 1991).

16 See (Hausman and Woodward, 1999, p. 523). Note that Hausman and Woodward’s definition is distinct in some significant ways from the original in Spirtes, Glymour and Scheines (2000 [1993], p. 29) —see (Steel, 2006) for a discussion—. The distinction makes no difference to our argument, however, so we ignore it here —and instead stick to Hausman and Woodward’s definition for consistency—.

Causal Markov Condition (CMC): For all distinct variables X and Y in the variable set \mathbf{V} , if X does not cause Y then

$$p(X | Y \wedge \text{Par}(X)) = p(X | \text{Par}(X)),$$

where $\text{Par}(X)$ (read *parents of X*) is the set of all direct causes of X in \mathbf{V} .

The Causal Markov Condition is an extrapolation of the PCC to directed acyclic graphs. The PCC states that a common cause screens-off its effects from each other, as long as there are no direct causal links between these effects. The **CMC** more generally states that the parents of X ($\text{Par}(X)$) screen-off X from any other variable Y in the variable set \mathbf{V} that is not a direct causal descendent of X . In short: if X does not cause Y in \mathbf{V} , then $\text{Par}(X)$ will screen them off. The contraposition is rather useful in the EPR set-up: if the putative parent of one of the measurement outcome events, say a , does not screen it off from the other outcome event b then it follows that a does not cause b , or we have not identified the only putative parent. In the EPR scenario it is often assumed that (i) the two measurement outcome events can not be causally connected because of relativistic constraints, and (ii) that the only putative common parent of such measurement outcome events is the singlet state at the source. So the residual correlation between the events a and b which does not disappear when **CMC** is applied, must be accounted by some rather mysterious nomic and acausal mechanism.¹⁷ This of course is very much in line with Readhead's thought that underlying the EPR correlations are non-robust stochastic links that are unthreatening to special relativity. It will then not come as a surprise that there is a strong formal connection between the **CMC** and the *robustness* conditions.

6 Robustness and the Causal Markov Condition

We show in this section that *robustness* is indeed a consequence of applying the Causal Markov Condition to an EPR setting, given some additional assumptions. In fact we show this for both of Healey's conditions by simply applying the Causal Markov Condition to total and partial causes respectively.

6.1 Total Causes and the Causal Markov Condition

Let us first consider *robustness*. Let us suppose that b is the *total cause* of a . In

¹⁷ See (Hausman and Woodward, 1999, pp. 564-7) and (Hausman, 1999).

this case b is the set of *all parents* of a . That is:

(I) If b is the total cause of a , then $Par(a) = b$.

Let us first assume, following *robustness*, that there exist a small disturbance d on the putative parents of a , and let us substitute d in for the term Y in the expression of **CMC**:

(II) $\exists d: d = Y$.

Let us then assume that the measurement outcome event a is not a cause of the small disturbance, i.e. let us assume:

(III) a does not cause d .

And finally let us turn to the definition of the Causal Markov Condition **CMC** given in the previous section. By substitution it follows from (I), (II), (III) and **CMC** that:

$$p(a|d \wedge b) = p(a|b),$$

which is an explicit expression of *robustness*. Thus we have formally shown that under assumptions (II) and (III):

$$(TotalCause) \wedge (CausalMarkov) \Rightarrow (Robustness).$$

In other words under the assumption of the existence of independent disturbing causes, *robustness* is the consequence of applying the Causal Markov Condition to total causes.

6.2 *Internal Robustness, Partial Causes and the Causal Markov Condition*

Now let us turn to *internal robustness* and *partial* causes. Let us then suppose that b is a partial cause of a . It follows that there is a non-empty set of additional variables c that represent all other independent causal antecedents of a . Let us

simplify by bringing them all under an additional variable c in the causal graph. Then the complete set of parents of a in the graph is the union of c and b :

(I') If b is *partial cause* of a , then: $Par(a) = \{b, c\}$.

Let us assume, as before, the existence of a small disturbance d on the putative parents of a in place of Y in the expression of **CMC**:

(II) $\exists d: d = Y$.

And similarly, that the measurement outcome event a is not a cause of the small disturbance, i.e. that:

(III) a does not cause d .

By substitution, it follows from (I'), (II') and (III') and the **CMC** that:

$$p(a|d \wedge b \wedge c) = p(a|b \wedge c),$$

which is Healey's *internal robustness*. Thus we have formally shown that under the same assumptions (II) and (III):

$$(PartialCause) \wedge (CausalMarkov) \Rightarrow (InternalRobustness).$$

Under the assumption of independent disturbing causes, *internal robustness* is a consequence of applying Causal Markov to partial causes.¹⁸

6.3 Robustness updated

We have shown that *robustness* and *internal robustness* are consequences of applying the Causal Markov Condition to the measurement outcome events a and b . If b is taken to be a total cause of a then the **CMC** together with some special assumptions entails *robustness*. If on the other hand b is taken to be merely a partial

¹⁸ A referee pointed out that the role of total or partial cause in these proofs is to make sure that d can only cause a via b in the case of total cause, and via $\{c, b\}$ in the case of partial cause. Indeed that would be an alternative definition of Healey's terms.

cause of a then the **CMC** with the same assumptions entails *internal robustness*. So it seems that the intuition underlining Michael Redhead's conditions is as a matter of fact the Causal Markov Condition. And the contrary intuitions and arguments by their critics are conversely related to doubts regarding the Causal Markov Condition. The **CMC** backs up Redhead's *robustness*, so if **CMC** is false in general as many recent critics believe,¹⁹ then *robustness* is left without substantial justification. The failure of *robustness* in EPR established by Redhead would be without any consequences were it not backed up by the **CMC**.

Moreover we have shown that a failure of Redhead's conditions entails a failure of the **CMC** regardless of whether the putative link is taken to be a total or a partial cause. So the distinction between total and partial causes that seemed so important in the early 1990's now seems irrelevant. The Causal Markov Condition is what underlies Redhead's intuition regardless. Similarly Healey's subtle distinctions between kinds of *robustness* are now seen to be irrelevant for a proper assessment of the causal nature of the EPR correlations. The peaceful coexistence between quantum mechanics and relativity so sought after by philosophers in the early 1990's is to be achieved always at the cost of a violation of the **CMC**, regardless of the underlying causal structure. So philosophers of physics interested in the issue of coexistence would be well advised to turn to a careful and detailed analysis of the implications of the **CMC** to the EPR correlations. This is essentially the central claim of our paper, and we find it remarkable that it needs to be made. But indeed it does, for such an analysis has not yet been carried out. We can at best find the outlines in the very brief discussion of EPR in Hausman and Woodward²⁰, and in a recent paper by Daniel Steel²¹.

7 EPR and the Causal Markov Condition

It has been claimed (for example by Salmon²²) that many genuinely statistical phenomena violate the PCC. Most prominently the EPR correlations are supposed to provide a set of established correlations that can not be explained by either a direct cause or a common cause model under the strictures of PCC.²³

19 (Cartwright, 2002) and (Williamson, 2005).

20 (Hausman and Woodward, 1999).

21 (Steel, 2005).

22 (Salmon, 1984, ch. 7).

23 One of us has argued against this common lore (Suárez, 2007). However, these arguments do not vindicate the PCC as usually stated but a very different reformulation. We will not review this literature here, but instead refer the reader to

Yet an important part of Hausman and Woodward's defence of **CMC** is that EPR is no counterexample.²⁴ They do not claim that the **CMC** is satisfied by the EPR correlations, but rather that it is *inapplicable*: it is neither satisfied nor violated, simply inappropriate. The discussion interestingly brings out some crucial differences between on the one hand the PCC as usually understood and on the other the **CMC** and the *robustness* conditions. So we review it briefly here.

7.1 Causal Markov, Interventions and Modularity

The key difference between the usual statement of the PCC and the **CMC** is the assumption of invariance under intervention that, according to Hausman and Woodward underlies and motivates **CMC**. This is best expressed in the modularity condition:²⁵

Modularity (MOD): For all subsets **Z** of the variable set **V**, there is some non-empty range **R** of values of members of **Z** such that if one intervenes and sets the value of the members of **Z** within **R**, then all equations except those with a member of **Z** as dependent variable (if there is one) remain invariant.

Hausman and Woodward take this condition, in conjunction with a few others, to provide the grounds for the **CMC**. The set **V** is the set of variables in the causal graph, and the equations are the linear regression equations that characterise a causal system. Modularity as a condition on causal systems is then the thought that a relation between two quantities *a* and *b* is causal only if (i) it is possible at least in principle to intervene in order to set the values of *a* and *b* and their probabilities, and (ii) these interventions —as long as within a permissible range— leave intact the functional connections between the values of *a* and *b*, or their probabilities.²⁶

The statement of **MOD** is a conditional with an antecedent that may be false, so a truth-functional interpretation as a material implication would entail that

that paper.

24 See (Hausman, 1999) and (Hausman and Woodward, 1999).

25 Cf. (Hausman and Woodward, 1999, p. 545).

26 The qualification of values *or* probabilities is needed to account for probabilistic causality, which Hausman and Woodward define as deterministic causation of probabilities (Hausman and Woodward, 1999, p. 570).

MOD is true by default in all such cases. But the context of the discussion suggests that **MOD** is meant to be non-applicable in such cases. That is, if interventions are possible in some set **V** and equations do not remain invariant then modularity is false. But if, on the other hand, interventions are not possible for some subset **Z** of **V** then **MOD** is strictly speaking not false but non-applicable.

Hausman and Woodward’s strategy is to attempt to back up **CMC** by appeal to **MOD**. But some significant additional assumptions are required to show **MOD** and **CMC** equivalent, namely: (i) *causal sufficiency* i.e. that all common causes are included in the set **V**; (ii) the assumption that all correlations have causal explanation; and (iii) the assumption that there exist unrepresented causes which can play the role of interventions. There is no need to get into the details of the equivalence proof, although it is worth mentioning that it has been contested.²⁷ In this paper we assume for the sake of argument that the proof is valid, and that a failure of **CMC** entails a failure of either of these conditions.

This has consequences for the discussion of the EPR correlations as we shall see. It also helps to distinguish subtly *robustness* from the usual statement of the PCC. For the PCC makes no implicit or explicit reference to interventions. By contrast, the notion of ‘disturbance’ required by *robustness* is clearly akin to an intervention. Hence a system that allows no interventions at all on any of its variables even in principle (or countenances no small disturbances) might violate the PCC without violating *robustness*.²⁸ This is the line defended by Hausman and Woodward with respect to the EPR correlations.²⁹ Their argument is essentially that there is no possible way to intervene on either of the distant measurement events. Consequently, they argue, it is impossible in this set-up to evaluate the **CMC**: the EPR correlations can not be shown to be a counterexample. This is precisely the claim we take issue with in this paper.

7.2 Interventions in EPR

The main aim of this paper is to urge that the debate over possible causal explanations of the EPR correlations ought to move to a detailed discussion of the **CMC** and its presuppositions in the context of the EPR experiment. Thus we

²⁷ See e.g. (Cartwright, 2002).

²⁸ The observation is consistent with our results in the previous section, since we showed that **CMC** entails *robustness* but not that modularity entails *robustness* —the main difference is clear now—.

²⁹ See (Hausman, 1999) and (Hausman and Woodward, 1999).

oppose Hausman and Woodward's thought that **CMC** is inappropriate for the EPR correlations. On the contrary we believe it is an appropriate kind of condition to apply, but we just do not share the widespread intuition that **CMC** (or the **PCC**) necessarily fails for the EPR correlations. We argue instead that whether or not **CMC** holds depends very much on the details of the precise causal hypothesis under test. The question requires investigation and can not be brushed aside as speedily as Hausman and Woodward would like.

In a sense we believe we have already achieved this aim—it follows from the equivalence proof in Section 6—. Hence the paper so far may be taken as endorsement of the suitability of **CMC** for understanding the status of causality in quantum phenomena. However, in this final section we outline our disagreement with the particular conclusions that Hausman and Woodward draw concerning the EPR correlations.

Hausman and Woodward back up the **CMC** with **MOD**. So to evaluate their claims we must concern ourselves with whether interventions are possible in the EPR context, and what significance must be attached to this fact. Hausman and Woodward endorse the view that in the EPR set-up there are no distinct mechanisms in the wings of the experiment because in fact there are no different systems to speak of. Both entangled particles are just 'parts' of the same irreducible holistic or non-separable system.³⁰ Together with the fact that there is no way to control the outcome of the first measurement this indeed seems to entail that interventions to set the values of the outcome events *a* and *b*, separately or jointly, are impossible. They conclude that EPR is no counterexample to **MOD** or **CMC**, but rather that these conditions are inappropriate in this context.

However, note that **CMC** states nothing whatever regarding interventions. It neither requires nor disallows interventions. Hausman and Woodward justify **CMC** by appeal to **MOD**, and the latter condition certainly requires interventions. But **CMC** could in principle be justified by other means that do not require **MOD**, as long as some of the additional assumptions are forfeited. So, contrary to what Hausman and Woodward seem to claim, the applicability of

30 They refer extensively to an old paper by Skyrms that defends this view (Skyrms, 1984); it is worth mentioning that the literature on EPR has moved on a very great deal in the last two decades, particularly on the physics side. Quantum entanglement was not then the area of intense research among physicists that it has become now, and Skyrms' views were much more entrenched twenty five years ago than they are now among both physicists and philosophers.

CMC does not seem to turn on the applicability of **MOD** and the related availability of interventions.

We have already noted that both **MOD** and **CMC** are explicitly stated as conditions on either values *or* probabilities of variables in the variable set **V**. In cases of genuine probabilistic causation the only relevant factor are the probabilities of the variables, since the causal structure fails to determine the values themselves. And it is of course well known that deterministic hidden variables are ruled out for quantum mechanics by the Bell inequalities. Hence the EPR correlations are potentially a paradigm but subtle case of probabilistic causality.

Interventions are not impossible to set the *probabilities* of some of these outcomes in the appropriate circumstances. For notice that the experimenter controls the settings of the measurement apparatus that determine the direction of spin that gets measured on each wing. Let us refer to the two wings of the experiment and their corresponding particles as '1' and '2'. It is true that the spherical symmetry of the singlet state entails that the first measurement outcome in the laboratory rest frame always has probability one half, regardless of what direction one measures spin along. Suppose then that in that frame spin is measured on '1' first, and suppose the outcome corresponding to 'spin-up' is found. If this information is provided to the second experimenter on time to set the direction of spin measured on particle '2' she can then easily set her measurement device to definitely get the outcome corresponding to 'spin-down' with probability one (or indeed any other probability but zero). For any value of probability of 'spin-down' on particle '2' she can use quantum mechanics to calculate the appropriate direction of measurement and set her device accordingly.

So it turns out that interventions are possible in particular experimental EPR set-ups. Notice that the intervention does not just consist in choosing a frame; rather given any frame, an intervention is the setting of a polariser direction. In such set-ups the question is then whether **MOD** and **CMC** hold. We urge in this paper that this is the relevant question to ask for causal modellers of EPR; but we will not attempt a comprehensive answer here. The answer is complicated and depends on the details of the causal hypotheses under test.³¹ A brief and intuitive argument suggests that **CMC** may fail here. The EPR correlations are not screened-off by the creation event at the source. Similarly the value of the setting of the measurement device on '2' will not screen-off the outcome event in that

31 For a preliminary account see (Cartwright and Suárez, 2000).

wing from the outcome event in the distant wing. But this really says nothing about a direct causal link between the wings. And if **CMC** failed for indeterministic systems, as some authors argue, then a common cause structure underlying the direct cause link would also be possible, which means that **CMC** might fail for a and b too. However, this claim requires further investigation in the context of alternative causal hypotheses. For our purposes in this paper this is a side issue, since whatever the correct answer it will already show that **CMC** is applicable to the EPR correlations in spite of Hausman and Woodward's claims to the contrary.

7.3 Causal Markov and Other Interpretations

The argument we have just given shows that in any EPR experiment there always exists a subset of the relevant variables that are susceptible to intervention. This leaves open several causal accounts for the EPR correlations. The fate of the **CMC** very much depends on the details of each account. But we believe that a stronger claim can be made. So far we have been assuming the standard or orthodox interpretation of quantum mechanics. So we have assumed that the violation of the Bell inequalities in the EPR experiments is due to a failure of what is known as *outcome independence*, and correspondingly that the only possible causes of each measurement outcome event are the distant outcome event and the proximate measurement device setting event.

In other words we have assumed that it is meaningless to suppose that the setting events in each wing can be a causal influence upon the *distant* outcome events. But it is well known that on some interpretations of quantum mechanics this is essentially what happens. The paradigm case is Bohm's theory. On the account of the EPR correlations provided in Bohmian mechanics,³² the actual measurement outcome event on one wing has no influence upon the measurement outcome event on the other wing, because in Bohm's theory measurements simply reveal values that are already there, they do not bring these values into being. Yet the setting of the distant device does have a causal influence, since it affects the quantum wavefunction of both particles in configuration space, and thus causally affects the probabilities for outcomes in the distant wing. The distant setting does not determine the proximate outcome of course, which also depends on the initial wavefunction state and the initial complete state of both particles.

32 See (Cushing, 1985, pp. 82-95) for a very nice review.

Daniel Steel³³ has claimed that Bohm's theory shows that **CMC** can fail for deterministic systems. The claim is part of a larger argument in the debate over whether the **CMC** is satisfied only by deterministic, or more generally 'pseudo-deterministic', systems.³⁴ Steel argues that the key to the validity of the **CMC** is not whether the system is deterministic or pseudo-deterministic but rather whether there are exogenous variables that are inherently random and hence not susceptible to intervention. Bohm's theory is an important part of the argument because it is the only example that Steel provides of a deterministic system that does not satisfy **CMC**. Presumably, given Steel's argument, this must be the case because there are some exogenous random variables in Bohm's description of EPR. Steel's overall argument is in line with Hausman and Woodward's proof of the equivalence of **MOD** and **CMC**. But since it is not the aim of this paper to debate the general validity of **CMC** we will not assess the general argument. We are interested though in assessing the chances of **CMC** for the EPR correlations. And we conjecture that the fate of the **CMC** in EPR is extremely sensitive to both the details of the causal hypothesis under test *and* the interpretation of quantum mechanics that is adopted.

So does the Bohmian description of the EPR correlations violate **CMC**? Steel assumes that it does since it predicts the very same EPR correlations. As he writes:

[...] the EPR example is a problematic basis for the claim that the **CMC** is a more reliable assumption for deterministic than indeterministic systems for the simple reason that there is a fully deterministic (though heterodox) interpretation of quantum mechanics, namely Bohm's. Bohm's quantum theory predicts precisely the same non-local (and hence putatively non-causal) correlations in the EPR example as the standard, indeterministic interpretation. Hence it is far from clear that the blame for the (putative) counter-example can be laid at the door of indeterminism.

33 Cf. (Steel, 2005).

34 By pseudo deterministic system we mean a system with causes that do not fix the occurrence of all their effects, but that can nonetheless be "embedded in another more complete graph [...] in which the parents of the given effect are sufficient to fix the value of the effect". (Cartwright, 1999). For a discussion and a reference to the notorious *cheap but dirty factory* example of the presumed failure of **CMC** in indeterministic systems see (Cartwright, 1993).

In our view this makes the very mistake to suppose that the fate of the **CMC** is independent of the details of the causal hypothesis under test. There are two versions or interpretations of Bohm's theory: the minimal Bohm theory championed originally by Bell³⁵, and the causal interpretation defended by Dewdney et al.³⁶ and Holland³⁷. According to the minimal interpretation, particles' only primitive property is position, and there is no such thing as intrinsic "spin". Instead the theory manages to produce the same predictions as quantum mechanics for the motion of all particles going through a Stern-Gerlach apparatus simply by means of the influence of the guiding field upon the particle though the so-called "guidance condition".³⁸ The causal interpretation, by contrast, has it that particles are endowed with the intrinsic property of spin, which is understood to be causally reactive to the quantum potential.³⁹ In both cases the causal structure is rather different and hence there is no real reason to expect **CMC**'s fate to be the same as in orthodox quantum mechanics. On the contrary, we would like to argue that at least in its causal interpretation the Bohmian description of EPR definitely satisfies the **CMC**.

We have already noted that in the EPR experiment as described by Bohm the measurement outcome events do not cause each other, but the setting events have an influence upon the outcomes. On the minimal interpretation, the settings influence the quantum wave function in configuration space in such a way that the motion of the particles is correspondingly affected after interaction with their respective Stern-Gerlach measuring devices. However, since no intrinsic property of spin is hypothesized, no changes take place ahead of the particle's interaction with their respective measurement devices. So on Bohm's minimal theory, the settings causally influence the outcomes via the measurement process only.

On the causal interpretation by contrast, the settings have a direct and instantaneous causal influence upon both particles' spin values. Indeed the underlying determinism of the theory implies that, on this causal interpretation, the setting events are instantaneous partial causes of the values of spin of the distant particles, which are only later revealed by measurement, if there ever is one, on the distant wing. So, on this view, my setting the measurement device of particle '1' partially determines not just the probability of an outcome of a measurement on particle '2' —it actually partially determines its value—. The

35 (Bell, 1982).

36 (Dewdney et al., 1988).

37 (Holland, 1993).

38 The full details can be found in (Bohm and Hiley, 1993, ch. 10). See (Berkovitz, 2007, Section 5.3.1) for a brief review.

39 Dewdney et al. (1988, pp. 537-9); Holland (1993, chapters 10 and 11).

reason is that particles on Bohm's theory have well defined values of their dynamical variables at all times —so on the causal interpretation the EPR particles have a value of position and spin from the word go, as they are ejected from the source—. This value can change though at any time, and in the case of an entangled particles as in the EPR case, it might do so non-locally as a result of changes in the wavefunction. And the wavefunction is responsive not only to the values of the distant entangled particles but also to the features of the systems those distant particles interact with. Thus although essentially non-local, the causal Bohm theory is indeed also essentially causal, in the strong sense of the *c=d identity* that we mentioned in the introduction.

How do we evaluate **CMC** then? Since in Bohm's theory measurement outcomes a and b do not cause each other, we can apply **CMC** fully as follows:

Causal Markov Condition (CMC) for Bohm's theory: For measurement outcome events a and b , since a does not cause b then

$$p(a|b \wedge Par(a)) = p(a|Par(a)),$$

where $Par(a)$ is the set of all direct causes of a in \mathbf{V} .

There is absolutely no reason to suspect that in Bohm's theory this condition is false, in either the minimal or the causal versions of the theory. On the contrary, since in Bohm's theory the explicit causal antecedents of the measurement outcomes include the quantum wavefunction, the initial complete states of both particles (which includes their spin in the causal interpretation) and the distant settings, it follows that $Par(a)$ includes all these. And since these variables jointly determine the value of the outcomes a and b , they jointly determine their probabilities. So the **CMC** is trivially satisfied in the Bohmian description of the EPR correlations, as long as we include in the set V all those variables that according to the theory are effectively causal antecedents of the outcomes a and b .

8 Conclusions

Our aim in this paper has been to urge more research to be conducted on applying the Causal Markov Condition to the diverse models and interpretations of the EPR correlations. We hope to have shown that questions regarding the causal nature of explanations of the EPR correlations are best explored by means of a

detailed and careful analysis of the application of the **CMC**. This is the right framework to update the debate regarding Michael Redhead's *robustness* in the early 1990's and to generally conduct the debate. Despite claims to the contrary the answers are not trivial, and the **CMC** is in principle applicable to the EPR correlation phenomena. But questions remain as to whether **CMC** is satisfied by these phenomena. We conjecture that the answer to this question is highly sensitive to the details of the causal hypothesis under test. We have also claimed it to be sensitive to the interpretation of quantum mechanics that is adopted, a claim that we have supported by looking at the Bohmian description of the EPR experiment. Contrary to recent claims the Bohmian description of the EPR correlations satisfies **CMC**.

This suggests that the **CMC** is a generally valid background or methodological assumption for deterministic or pseudo-deterministic systems.⁴⁰ It remains to be seen whether it can be similarly assumed for indeterministic ones such as EPR on the orthodox interpretation of quantum mechanics. Concomittantly it also remains to be seen whether a causal understanding of indeterministic phenomena requires the **CMC**. Suppose that **CMC** fails for at least some of the main causal hypotheses for the EPR correlations under the standard or orthodox understanding. If **CMC** is not required for causation then even the weakest interpretation of the *c=d identity* will have been refuted. If on the other hand **CMC** is required for causation then quantum mechanical phenomena, on the orthodox interpretation at least, abandons causality as well as determinism, the *c=d identity* is retained, and the intuition of the founding parents is proved correct (for orthodox quantum mechanics at least). The questions are relevant, the stakes are high, and the answers should be informative.

40 *Modulo* the usually discussed exceptions such as nomic or non-causal inducers of correlations —see e.g. (Hausman, 1999)—, and accidental dependencies such as the one between British bread prices and the water level in Venice in Sober's famous example (Sober, 2001).

Bibliography

- Bell, J. S. (1982), “On the Impossible Pilot Wave”, *Foundations of Physics*, 12, pp. 989-999. Reprinted in Bell (1987), *Speakable and Unspeakable in Quantum Mechanics*, Cambridge University Press, pp. 159-168.
- Berkovitz, J. (2007), “Action at a Distance in Quantum Mechanics”, *Stanford Encyclopedia of Philosophy*.
- Bohm, D. (1952) “A Suggested Interpretation of Quantum Theory in Terms of Hidden Variables, I and II”. *Physical Review*, **85**: 166–193 and 369–396.
- Bohm, D. And Hiley, B. (1989) “Non-Locality and Locality in the Stochastic Interpretation of Quantum Mechanics”, *Physics Reports*, 172, 3, pp. 93 -- 122.
- Bohm, D. and Hiley, B. (1993) *The Undivided Universe*. Routledge.
- Butterfield, J. (1990) “Causal Independence in EPR Arguments”. In *Proceedings of the Bienal Meeting of the Philosophy of Science Association*, vol. I, pp. 213–25.
- Butterfield, J. (1992) “David Lewis meets John Bell”. *Philosophy of Science*, **59**: 26–43.
- Cartwright, N. (1990) “Quantum Causes: The Lesson of the Bell Inequalities”. In *Philosophy of the Natural Sciences: Proceedings of the 13th International Wittgenstein Symposium*. Hölderlin-Pichler-Tempsky.
- 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. (1999) “Causal Diversity and the Markov Condition”. *Synthese*, **121**: 3–27.
- Cartwright, N. (2002) “Against Modularity, the Causal Markov Condition, and Any Link Between the Two: Comments on Hausman and Woodward”. *British Journal for the Philosophy of Science*, **53**: 411–53.
- Cartwright, N. and Jones, M. (1991) “How To Hunt Quantum Causes”. *Erkenntnis*, **35**: 205–31.
- Cartwright, N. and Suárez, M. (2000) “A Causal Model for EPR”. Discussion Paper 50/00, LSE Centre for the Philosophy of the Natural

and Social Sciences.

- Chang, H. and Cartwright, N. (1993) "Causality and Realism in the EPR Experiment". *Erkenntnis*, **38**: 169–90.
- Cushing, P. (1994), *Quantum Mechanics*, University of Chicago Press.
- Dewdney, C., Holland, P., Kyprianidis, A. And Vigier, J. P. (1988) "Spin and Non-Locality in Quantum Mechanics", *Nature*, **336**, 8, pp. 536 – 544.
- Einstein, A., Podolsky, B. and Rosen, N. (1935) "Can a Quantum Mechanical Description of Physical Reality be Considered Complete? " *Physical Review*, **47**: 777–80.
- Elby, A. (1992) "Should We Explain the EPR Correlations Causally? " *Philosophy of Science*, **56**: 16–25.
- Elby, A. (1993) "Why Local Realistic Theories Violate, Nontrivially, the Quantum Mechanical EPR Perfect Correlations". *The British Journal for the Philosophy of Science*, **44**: 213–30.
- Fine, A. (1982a) "Hidden Variables, Joint Probability, and the Bell Inequalities". *Physical Review Letters*, **48**: 291–95.
- Fine, A. (1982b) "Joint Distributions, Quantum Correlations, and Commuting Observables". *Journal of Mathematical Physics*, **23**: 1306–10.
- Glymour, C., Spirtes, P. and Scheines, R. (1991) "Causal Inference". *Erkenntnis*, **35**: 151–89.
- Hausman, D. (1999) "Lessons from Quantum Mechanics". *Synthese*, **121**: 79–92.
- Hausman, D. M. and Woodward, J. (1999) "Independence, Invariance and the Causal Markov Condition". *British Journal for the Philosophy of Science*, **50**: 521–83.
- Healey, R. (1992a) "Causation, Robustness, and EPR". *Philosophy of Science*, **59**: 282–92.
- Healey, R. (1992b) "Chasing Quantum Causes: How Wild is the Goose? " *Philosophical Topics*, **20**: 181–204.
- Heisenberg, W. (1958) *Physics and Philosophy*. Harper & Row.
- Hoefer, C. (2004) "Causality and Determinism: Tension, or Outright Conflict? " *Revista de Filosofia*, **29**: 99–115.
- Holland, P. (1993) *The Quantum Theory of Motion*. Cambridge University

Press.

- Maudlin, T. (1994) *Quantum Non-Locality and Relativity*. Blackwell Publishing.
- Papineau, D. (1990) "Causes and Mixed Probabilities". *International Studies in the Philosophy of Science*, **4**: 79–88.
- Redhead, M. (1987) *Incompleteness, Nonlocality and Realism*. Oxford Calrendon Press.
- Redhead, M. (1989) "The Nature of Reality". *British Journal for the Philosophy of Science*, **40**: 429–41.
- Salmon, W. (1984) *Scientific Explanation and the Causal Structure of the World*. Princeton University Press.
- Shimony, A. (1984) "Controllable and Uncontrollable Non-Locality". In Kamefuchi *et al.* (Eds.), *Proceedings of the International Symposium: Foundations of Quantum Mechanics in the Light of New Technology*, pp. 225–30. Physical Society of Japan.
- Skyrms, B. (1984) "EPR: Lessons for Metaphysics". In P. A. French and T. E. Uehling, Jr (Eds.), *Causation and Causal Theories*. *Midwest Studies in Philosophy*, 9, pp. 245–55. University of Minnesota Press.
- Sober, E. (2001), "Venetian sea levels, British bread prices and the principle of the common cause", *British Journal for the Philosophy of Science*, **52**, pp. 1-16.
- Spirtes, P., C. Glymour and R. Scheines (eds.), (2000 [1993]) *Causation, Prediction and Search*, 2nd Edition, MIT Press.
- Steel, D. (2005) "Indeterminism and the Causal Markov Condition". *British Journal for the Philosophy of Science*, **56**: 3–26.
- Steel, D. (2006), "Comment on Hausman and Woodward on the Causal Markov Condition", *British Journal for the Philosophy of Science*, **57**, pp. 219 – 231.
- 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.
- Suppes, P. and Zanotti, M. (1981) "When are probabilistic explanations possible?" *Synthese*, **48**: 191–99.
- van Fraassen, B. C. (1982) "The Charybdis of Realism: Epistemological Implications of Bell's Inequality". *Synthese*, **52**: 25–38. Reprinted with

corrections in J. Cushing and E. McMullin (Eds.), *Philosophical Consequences of Quantum Theory*, University of Notre Dame Press, 1989.

von Neumann, J. (1955) *Mathematical Foundations of Quantum Mechanics*. Princeton University Press. First published in German in 1932 as *Mathematische Grundlagen der Quantenmechanik*, Springer.

Williamson, J. (2005) *Bayesian Nets and Causality: Philosophical and Computational Foundations*. Oxford University Press.