Venetian Sea Levels, British Bread Prices and the Principle of the Common Cause: A Reassessment

Iñaki San Pedro*

1   Introduction

An influential argument against Reichenbach's Principle of the Common Cause (RPCC), first proposed by Elliott Sober (1987, 2001), consists on an example which involves correlations between bread prices in Britain and sea levels in Venice. The following quotation summarises the spirit of the whole argument:

> Because both quantities have increased steadily with time, it is true that higher than average sea levels tend to be associated with higher than average bread prices. [...] we do not feel driven to explain this correlation by postulating a common cause. Rather, we regard Venetian sea levels and British bread prices as both increasing for endogenous reasons. [...] Here, postulating a common cause is simply not very plausible, given the rest of what we believe. (Sober, 2001, p. 332)

There have been different attempts to deal with examples of the kind of Sober's 'Venetian sea levels and British bread prices'. It is striking though that none of these make use of recent development as regards the formal structure of the probability spaces where the relevant events, corresponding correlations and potential common causes that might be involved in a causal explanation of such correlations are defined precisely. I refer, in particular, to the results by Hofer-Szabó et al regarding the existence of screening-off events for any correlation:

> [...] *every* classical probability space $(\mathcal{S}, \mu)$ is common cause completable with respect to *any finite* set of correlations [...] given *any finite* set of correlations in a classical event structure, one can *always say* that the correlations are due to some common

---
*Universidad Complutense de Madrid. (`inaki.sanpedro@filos.ucm.es`).

causes, possibly 'hidden' ones, i.e. ones that are not part of the initial set $\mathcal{S}$ of events. (Hofer-Szabó et al, 1999, p. 378)

Both quotations above seem to contain opposite claims and thus appear to convey two quite incompatible views. The aim of this paper is to put this two views into perspective and look at each of them in the light of the other. In particular, I shall investigate, in the light of the formal results by Hofer-Szabó et al, to what extent examples of the kind of Sober's constitute a thread to RPCC. This, in turn, shall provide a good analysis of the actual (physical) significance and range of applicability of the so-called *extendability* and common cause *completability* theorems. As a result, I will conclude that the two quotations above are to be taken to complement each other rather than as reflecting opposite incompatibles views.

The structure of the paper is as follows: Section 2 provides the mains of a formal approach to the main ideas behind RPCC, along with the main issues regarding its significance and philosophical status. Section 3 introduces the ideas of *extendability* and common cause *completability*. In Sections 4 and 5 I review Sober's argument against RPCC and discuss on what grounds it may be taken as a potential thread to the principle. Two recent reactions to Sober's example are reviewed here as well. Finally, I suggest in Sections 6 and 7 two alternative solutions to save RPCC from criticisms of the kind of Sober's. The paper closes with some concluding remarks.

2   Reichenbachian Common Causes

The idea of common cause goes back to Reichenbach and has its origins at the observation of apparently unrelated events that nonetheless take place simultaneously with a certain regularity:[1]

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

One way to proceed is to translate Reichenbach's original intuitions into a formal language so that the concept of common cause can be analysed in detail within the framework of classical probability theory.[2] The notion of correlation is central, as it is assumed to capture Reichenbach's intuitions regarding *improbable coincidences*. Let us then define correlation as follows:[3]

---

[1]Cf. (Reichenbach, 1956, p. 157).

[2]I follow here the work by Hofer-Szabó et al in the late 1990's and early 2000. See (Hofer-Szabó et al, 1999, 2000a,b) and (Rédei, 2002) for the main results of the program. We need to be aware that this 'formalisation' will only approximately capture some of the subtleties in Reichenbach's original intuitions. The whole program hinges however on the assumption that it is possible for the formal results achieved to be translated back into claims about the actual physical systems involved, and in particular into causal claims

[3]This definition is of positive correlation. A completely symmetrical definition may be given for negative correlations. Distinguishing between positive and negative correlations will not be important in what follows and positive correlations will thus be assumed.

Definition 1. *Let $(\mathcal{S}, p)$ be a classical probability measure space with Boolean algebra $\mathcal{S}$ representing the set of random events and with the probability measure p defined on $\mathcal{S}$. If A, $B \in \mathcal{S}$ are such that*

$$p(A \wedge B) - p(A) \cdot p(B) > 0, \tag{1}$$

then the events A and B are said to be (positively) correlated, and we write $Corr_p(A, B)$.

A *Reichenbachian* common cause is then defined as:

Definition 2 (Reichenbachian Common Cause). *An event C is said to be a Reichenbachian common cause if the following independent conditions hold:*

$$p(A \wedge B|C) = p(A|C) \cdot p(B|C) \tag{2}$$
$$p(A \wedge B|\neg C) = p(A|\neg C) \cdot p(B|\neg C) \tag{3}$$
$$p(A|C) > p(A|\neg C) \tag{4}$$
$$p(B|C) > p(B|\neg C), \tag{5}$$

where $p(A|B) = p(A \wedge B)/p(B)$ denotes the probability of A conditional on B and it is assumed that none of the probabilities $p(X)$ $(X = A, B, C, \neg C)$ is equal to zero.

With both the notions of correlation and Reichenbachian common cause above at hand Reichenbach's Principle of the Common Cause (RPCC) may be stated as follows:

Definition 3 (RPCC). *For any two (positively) correlated event types A and B ($Corr_p(A, B) > 0$), if A is not a cause of B and neither B is a cause of A, there exists a Reichenbachian common cause C of A and B, i.e. there exist an common cause event C such that it satisfies conditions (2)-(5).*

The definition above consists of two distinct independent claims. The first is a claim at the ontological level, regarding the existence of common causes, and the other at the methodological level which provides a concrete characterisation (through equations (2)-(5)) of the prior postulated common causes.

A proper distinction of these two claims is crucial for the assessment of the status of RPCC. In particular, since each of the two claims in RPCC is logically independent of the other, arguments aimed at criticising the characterisation of common causes through expressions (2)-(5) may very well leave untouched the metaphysical existential claim about common causes. In fact, while it is part of the received view that equations (2)-(5) do not constitute neither a necessary nor a sufficient condition for common causes, there have been prominent defenders of common cause explanation —or alternatively of some version of the Principle of the Common Cause— provided the notion

3

of common cause is characterised in an appropriate manner.[4] On the other hand, arguments devised to deny the very existence of common causes may be completely compatible with the claim that, whenever they exist, common causes are to satisfy equations (2)-(5).

Despite the controversies, endorsing RPCC may be motivated by at least two reasons. First, note that for Reichenbach the role of the principle as a whole, and of the screening-off condition in particular, is mainly explanatory. Reichenbach explicitly points out that the four statistical relations *explain* the correlations between A and B in two senses. First, he notes that the four relations entail that A and B are (positively) correlated, i.e. $Corr(A, B) > 0$. On the other hand, a common cause C satisfying these four relations explains the correlation by rendering A and B statistically independent.[5] The explanatory power of screening-off common causes by itself may thus be taken as a good reason to support RPCC.

Second, recent results show that, at least formally and under some qualifications, it is always possible to provide a Reichenbachian common cause (as defined above) for any given correlation.[6] These results build on the intuition that any probability space $\mathcal{S}$ which contains a set of correlations and which does not include (Reichenbachian) common causes of these, may be extended in such a way that the new probability space $\mathcal{S}'$ *does include* (Reichenbachian) common causes for each of the original correlations. These intuitions are formalised in so-called *extendability* and *common cause completability* theorems.

3 Common Cause Completability

It is not difficult to think of examples of probability spaces containing correlated events which do not feature however any event that conforms to the definition of Reichenbachian common cause. We shall call such probability spaces *Reichenbachian common cause incomplete* spaces.[7]

We seem to have two alternatives when dealing with Reichenbachian common cause incomplete probability spaces. Either we go for a weakening of the common cause criterion —this is for instance the case in both Salmon's 'interactive forks' and Cartwright's generalisation of the fork criterion—, or we may simply embark on the search for screening-off common causes, hoping that such events exist but have remained somehow 'hidden' to us all along. Here I shall only pay attention to the second alternative. So how

---

[4]Salmon (1984) and Cartwright (1987) advance perhaps the most influential proposals for a generalisation of Reichenbach's original criterion for common causes.

[5]Cf. (Reichenbach, 1956, p. 159).

[6]See (Hofer-Szabó et al, 1999, 2000b,a) for details.

[7]Reichenbachian common cause incomplete probability spaces are very common and, in fact, most examples aimed to rule out screening-off as a necessary condition for common causes exploit such incompleteness.

should this search be undertaken?

Note first that we need to be searching new events (i.e. the Reichenbachian common causes) which are not contained in the original (incomplete) probability space $(\mathcal{S}, p)$. Intuitively, we need some notion of extension than could be applied to our probability space. This is formally achieved as follows:[8]

Definition 4 (Extension). *The probability space $(\mathcal{S}', p')$ is called an extension of $(\mathcal{S}, p)$ if there exist a Boolean algebra embedding h of $\mathcal{S}$ into $\mathcal{S}'$ such that $p(X) = p'(h(X))$, for all $X \in \mathcal{S}$.*

Extendability allows then for the enlargement of the original probability space so that new events are included.

In a second step, we should be able to set up a procedure to enlarge our common cause incomplete probability space such that the new extended probability space *contains* common causes for the original correlations.[9] This intuition is formalised through the idea of Reichenbachian common cause completability:

Definition 5 (RCC Completability). *Let $Corr(A_i, B_i) > 0$ $(i = 1, 2, \ldots, n)$ be a set of correlations in $(\mathcal{S}, p)$ such that none of them possess a common cause in $(\mathcal{S}, p)$. The probability space $(\mathcal{S}, p)$ is called Reichenbachian common cause completable with respect to the set $Corr(A_i, B_i)$ if there exists an extension $(\mathcal{S}', p')$ of $(\mathcal{S}, p)$ such that $Corr(A_i, B_i)$ has a Reichenbachian common cause C_i in $(\mathcal{S}', p')$ for every $i = 1, 2, \ldots, n$.*

Completability is hence the key for successfully searching Reichenbachian common causes.

The question is now whether any incomplete probability space $(\mathcal{S}, p)$ can be extended such that it is turned into (Reichenbachian) common cause complete. We may ask, in other words, when is a probability space $(\mathcal{S}, p)$ Reichenbachian common cause completable?

Hofer-Szabó et al answer this question with the following proposition:[10]

Proposition 1. *Every classical probability space $(\mathcal{S}, p)$ is common cause completable with respect to any* finite *set of correlated events.*

The proposition shows that given a Reichenbachian common cause incomplete probability space an extension $(\mathcal{S}', p')$ may *always* be performed such that it contains (Reichenbachian) common causes for all the original correlations.

---

[8]Cf. (Hofer-Szabó et al, 2000a).

[9]Definition 4 ensures that the extension operation be consistent with the old event structure $(\mathcal{S}, p)$. In particular, the embedding h is defined such that correlations stay invariant under the extension operation, that is $Corr(A, B) \equiv Corr_p(A, B) \equiv Corr_{p'}(A, B)$. See (Hofer-Szabó et al, 2000a) for details.

[10]Cf. (Hofer-Szabó et al, 1999, p. 384).

Common cause completability hence constitutes a very powerful tool if we are to provide common cause explanations of generic correlations. It nevertheless faces its own problems, especially when it comes to the physical interpretation of either the enlarged probability space $\mathcal{S}'$ or the new common causes contained in it. In particular, it seems a fair criticism to the whole program to claim that common cause completability is merely a formal device, which is likely to lack physical meaning in many (perhaps too many) cases. I shall retake the issue later on and just point out for now that such criticisms may be successfully addressed.[11] On the other hand, the program hinges on two implicit assumptions which are far from uncontroversial. Namely, the assumption that common causes exist of any given correlation (except for those that result from direct causal influence), i.e. the metaphysical part of RPCC is assumed to be correct; and the assumption that such common causes are to be characterised by equations (2)-(5). Under these assumptions, and setting the issue of interpretability aside for now, common cause completability seems powerful enough a tool as to provide good methodological grounds in support of RPCC.

4 Venetian Sea Levels and British Bread Prices

Examples such as Sober's 'Venetian Sea Levels and British Bread Prices' (VSL & BBP) are devised to refute RCCP by criticising the metaphysical content of it. Sober however draws methodological consequences from the conceptual possibility (rather than the real existence) of correlations — e.g. between sea levels in Venice and bread prices in Britain— which cannot be accounted for in terms of common causes. (This, he argues, favours the use of the Likelihood Principle instead of RPCC). Although the intuitions behind Sober's argument are in my view quite strong, the argument seems to be at odds with the results we just discussed concerning common cause completability. I shall thus try to clarify in what follows whether there is such an incompatibility really. Let us first start with the details of Sober's example itself.

It is the case that the sea level in Venice (VSL) and the cost of bread in Britain (BBP) have been (monotonically) increasing during a given period of time. Table 1 displays the values of Venetian sea levels and British bread prices in accordance with Sober's actual example.

From the data displayed, we are told that 'higher than average values' of Venetian sea levels and those of British bread prices are correlated:[12]

---

[11]I point the reader to (San Pedro and Suárez, 2009) for a recent assessment of the significance of common cause completability, possible criticisms to it and possible strategies to avoid these.

[12]Cf. (Sober, 2001, p. 334). The appeal to 'higher than average values' rather than just 'values' is mainly motivated by criticism to an earlier version of the counterexample (Sober, 1987). There is no need to review such arguments here since they will not play any

Year (i)	1	2	3	4	5	6	7	8	$\langle$Year$\rangle = 4.5$
VSL	22	23	24	25	28	29	30	31	$\langle$VSL$\rangle = 26.5$
BBP	4	5	6	10	14	15	19	20	$\langle$BBP$\rangle = 11.625$

Table 1: Sober's Venetian sea levels and British bread prices data (Sober, 2001, p. 334).

As I claimed initially, higher than average bread prices *are* correlated with higher than average sea levels.

Let us denote the event 'the Venetian sea level in year i is higher than average' by the expression '$\text{VSL}_i > \langle\text{VSL}\rangle$'. (Similarly for bread prices in Britain '$\text{BBP}_i > \langle\text{BBP}\rangle$').

What Sober seems to have in mind when claiming that 'absolute values' of VSL and BBP are correlated is the following. On the one hand, the probability of observing a 'higher than average' Venetian sea level in year i can be seen (directly from the data displayed in Table 1) to be

$$p\left(\text{VSL}_i > \langle\text{VLS}\rangle\right) = 1/2.$$

Similarly, for British bread prices one has that

$$p\left(\text{BBP}_i > \langle\text{BBP}\rangle\right) = 1/2.$$

On the other hand, one may also calculate the joint probability of both:

$$p\left[(\text{VSL}_i > \langle\text{VSL}\rangle) \wedge (\text{BBP}_i > \langle\text{BBP}\rangle)\right] = 1/2.$$

These three probabilities entail that:

$$p\left[(\text{VSL}_i > \langle\text{VSL}\rangle) \wedge (\text{BBP}_i > \langle\text{BBP}\rangle)\right] \\ -p\left(\text{VSL}_i > \langle\text{VLS}\rangle\right) \cdot p\left(\text{BBP}_i > \langle\text{BBP}\rangle\right) > 0. \quad (6)$$

Thus, the argument goes on, 'higher than average values' of Venetian sea levels and British bread prices are in fact (positively) correlated.

The question is then how this correlation is to be explained away. Sober points out that there are three possible ways to go:[13]

_____

important role in the foregoing discussion. The important point is that Sober's later formulation stands. Sober also refers to 'higher than average values' as 'absolute values' and I shall use these two expressions indistinctly.

[13]These three possible explanations had been already suggested by Meek and Glymour after Yule (1926). See also (Sober, 2001, p. 332) and references therein.

(i) To postulate the existence of an unobserved common cause.

(ii) To claim that the data sample is unrepresentative.

(iii) To claim that the data arises from a mixing of populations with different causal structures and correspondingly different probability distributions.

Considering these three options in turn shows, according to Sober, that RPCC fails. The argument is as follows. First, Sober dismisses option (ii) by pointing out that the correlations in his example do not come out of an unrepresentative sample since data could be spread over a larger period of time and the correlations would still be there —I completely agree with this and I will also dismiss option (ii) altogether. Second, option (i) is false in the example *ex hypothesis*. Consequently, Sober takes option (iii) to provide the right (causal) explanation of the correlation. This constitutes a failure of RPCC.

5   Is VSL & BBP a Genuine Counterexample?

But does Sober's VSL & BBP really constitute a genuine counterexample to RPCC? In order to answer this particular question we need to address two further questions, I think. First, we need to know whether the VSL & BBP correlation is indeed genuine (as defined formally in Section 2). Second, in case it is so, we need to ask why is it that there is no possible common cause explanation of the VSL & BBP correlation. In other words, if the counterexample is to stand, we need to make sure that no common cause at all can be provided for the correlation. It is not enough to just assume this *ex-hypothesis*, in my opinion. I will tackle these two questions in turn

Kevin Hoover and Daniel Steel have both recently tried to diffuse Sober's example at some level, each with quite different arguments and each reaching different conclusions. The main issue these two reactions differ in is, in fact, the answer to our first question, i.e. whether the VSL & BBP correlations are indeed genuine correlations or mere associations of the sample.

As we have seen, the probabilities of 'higher than average values' of Venetian sea levels and British bread prices display what seems to be a probabilistic dependence —by means of expression (6) on page 7. However, if we are to conclude that 'higher than average values' of sea levels and bread prices are correlated, we need first make sure that the probabilities involved refer to the same and only probability space.[14] In other words we need to check that 'absolute values' of sea levels and bread prices are events of the very same probability space.

---

[14]This is explicitly required in the formal definition of correlation (Definition 1).

But nothing in the data set tells us the probability measure should be the same. In fact, the probabilities for each quantity are derived quite independently (from the relative frequencies of the corresponding measured sea levels and bread prices over a time span). Strictly speaking we should perhaps have initially written them as $p^1(\text{VSL}_i > \langle \text{VSL} \rangle)$ and $p^2(\text{BBP}_i > \langle \text{BBP} \rangle)$, i.e. as referring to two different probability measures p^1 and p^2. Similarly for the joint probability we should perhaps have written it relative to yet another probability measure, i.e.

$$p^3 \left[(\text{VSL}_i > \langle \text{VSL} \rangle) \wedge (\text{BBP}_i > \langle \text{BBP} \rangle) \right].$$

With this in mind, relation (6) should have been written:

$$p^3 \left[(\text{VSL}_i > \langle \text{VSL} \rangle) \wedge (\text{BBP}_i > \langle \text{BBP} \rangle) \right]$$
$$-p^1(\text{VSL}_i > \langle \text{VSL} \rangle) \cdot p^2(\text{BBP}_i > \langle \text{BBP} \rangle) > 0. \quad (7)$$

Now, the question whether the expression above reflects a correlation between sea levels and bread prices may be restated in terms of these three different probability measures, i.e. are p^1, p^2 and p^3 in fact one and the same probability measure?

The above is somehow related to Hoover's arguments in reaction to the VSL & BBP case.[15] Hoover distinguishes correlations from mere associations of the sample. Very succinctly, while associations are a property of the sample, correlations are a property of the probabilistic space used to model it. Hoover assumes that it is only correlations that can reveal 'real' properties of the system. In our case then, if the probability measures p^1, p^2 and p^3 would be different, expression (7) could only be said to reflect some degree of association between sea levels and bread prices, but not a correlation. In order for it to represent a correlation a consistent probabilistic model —with a single probabilistic measure, that is— must be construed such that the example's data may be embedded in it. This is not the case in Sober's example, in Hoover's view. He thus concludes that the VSL & BBP scenario does not constitute a counterexample to the RPCC.

Hoover's case might find support in Sober himself. For, as it is claimed in the original argument, each data series belongs to different causal structures —this was option (iii) in Sober's argument. This can then be seen to suggest that the two probability spaces need to be different. But is it right to claim that the 'Venice-Britain' scenario cannot really be described in a whole single probability space? Put it differently, why could it not be the case that data in Table 1 give rise to genuine correlations? I think this is indeed an option. In particular, while I share Hoover's views as regards mere associations and genuine correlations, I do not see why may not the data in Table 1 be embedded, or modelled if we like, in a single probability space.

---

[15]Cf. (Hoover, 2003).

An argument along these lines is provided by Steel, who makes use of a well known mathematical result, namely the 'mixing theorem'.[16] Very briefly, the 'mixing theorem' provides us with information about the behaviour of a probability distribution resulting from the mixing of the distributions from two populations, each of which with probabilistically independent variables. The theorem tells us, in particular, under what conditions the variables of such a 'mixed' probability distribution are probabilistically independent. The theorem then shows that a probability distribution may display dependencies just because it is the result of the mixing of two other probability distributions. Steel claims this is the case in Sober's 'Venice-Britain' example, and constructs a model from two initial sets of data (of both VSL and BBP), each corresponding to different (distant) time spans. If the probability distributions from these two populations are mixed, the resulting distribution displays probabilistic dependencies in just the manner suggested by Sober.[17] Summing up, I see no convincing reason why probabilistic dependencies such as those in the 'Venice-Britain' example would not be *genuine* correlations.

Once the question as to whether VSL & BBP are (genuinely) correlated has been positively answered, we shall turn to our second question. Is it really the case that no common cause explanation can be given of the correlation between sea levels in Venice and bread prices in Britain?

In relation to this question, we seem to face three possible scenarios. First, we may find out that it is certainly the case that no common cause whatsoever can be provided that explains the correlation. In that case the example would indeed constitute a genuine counterexample to RPCC. Alternatively, and contrary to what Sober thinks, we might find out, with the help of common cause completability for instance, that it is indeed possible to provide a common cause explanation of the correlation. Finally, Sober's criticism could also be avoided if we could show the question to be non-applicable. For example, even if we take the VSL & BBP correlation to be genuine we might want to argue that it needs no causal explanation after all, perhaps due to the fact that it does not reflect any feature of the system itself. It seems to me that the first option —that is, conceding Sober's point— only makes sense once the other two have been ruled out. Let us then consider the two last options in turn.

6 Screening-off Events Exist

Recall that, by common cause completability, an appropriate extension of a common cause incomplete probability space guarantees that there exist a screening-off event. For what it has been said up to now, there is no reason

---

[16]Cf. (Steel, 2003).

[17]See (Steel, 2003) and references therein for further details.

why common cause completability should not also work in our case at hand.

Suppose for instance a new model —an appropriately extended probability space— contains events of the type '$Y_i > \langle \text{Year} \rangle$', which we may call, following Sober's terminology, 'higher than average time values', 'higher than average values of years', or 'absolute values of years', never mind how strange this may sound. We may then assign probabilities to such events in exactly the same way as we did for 'higher than average' values of sea levels and bread prices, that is by referring to their relative frequencies. (Only, we need to make sure that the probability measure is the same for all three values so as to make sure that resulting correlations are indeed genuine correlations.) Thus, we may write, again from the data on Table 1,

$$p\,(Y_i > \langle \text{Year} \rangle) = 1/2.$$

If we now take conditional probabilities we obtain, also looking at the data in Table 1,

$$p\,(\text{VSL}_i > \langle \text{VSL} \rangle \mid Y_i > \langle \text{Year} \rangle) = 1, \qquad (8)$$
$$p\,(\text{BBP}_i > \langle \text{BBP} \rangle \mid Y_i > \langle \text{Year} \rangle) = 1. \qquad (9)$$

It is also easy to check that

$$p\,[(\text{VSL}_i > \langle \text{VSL} \rangle) \wedge (\text{BBP}_i > \langle \text{BBP} \rangle) \mid Y_i > \langle \text{Year} \rangle] = 1. \qquad (10)$$

It is then clear that, as soon as we consider the event '$Y_i > \langle \text{Year} \rangle$', the correlation will vanish. This is because the dependence of the original series washes out conditional on '$Y_i > \langle \text{Year} \rangle$'. In particular if we define a new probability measure $p^Y = p(\cdot \mid Y_i > \langle \text{Year} \rangle)$, the above equations yield

$$p^Y\,[(\text{VSL}_i > \langle \text{VSL} \rangle) \wedge (\text{BBP}_i > \langle \text{BBP} \rangle)]$$
$$- p^Y(\text{VSL}_i > \langle \text{VSL} \rangle) \cdot p^Y(\text{BBP}_i > \langle \text{BBP} \rangle) = 0. \quad (11)$$

This example is of course specific for the case at hand, and the 'trick' has been quite the obvious one, since I transformed both the original data series by describing them in the probability space in which all probabilities are one. In this case the probability space with such 'nice' properties is particularly easy (and obvious) to find since the time dependence of both sea levels and bread prices' higher than average values is exactly the same. However obvious and specific this example might be, I hope it illustrates sufficiently how screening-off events can be provided for such kind of correlations.

This is not quite a causal explanation yet, since the event '$Y_i > \langle \text{Year} \rangle$' does not seem capable of a causal interpretation in an obvious or straightforward way. The question is, more specifically, how can we make sense of events '$Y_i > \langle \text{Year} \rangle$' as (common) causes. I must admit that I do not have an answer to this question. For, in what relevant sense is time a causal

factor in the VSL & BBP example? Indeed, I find it hard to understand '$Y_i > \langle \text{Year} \rangle$' as a cause event.

Of course, problematic issues of this kind will be present whenever a time dependent event needs to be interpreted as a cause, but not only then. In fact, the problem is more general and has to do with the actual significance and applicability of the (formal) *extendability* and *common cause completability* theorems. I haven mentioned already that it is in the spirit of the whole program that the various claims resulting from the formal treatment of the problem can be translated back into claims about the particular physical systems involved, and more specifically into causal claims. In this sense, common cause completability will not achieve its purpose unless we are able to provide the formally obtained screening-off events with a physical interpretation which can be made to fit into a causal story. I see two ways one can try to solve, or at least minimise, the difficulties with the interpretation of such screening-off events. An obvious case for *common cause completability* can be made if we recall that the extension of a common cause incomplete probability space is not unique. We may thus hope (I think quite optimistically but nevertheless realistically) that among all possible extensions of the original probability space, there must exist one that contains a screening-off event which can be adequately interpreted as a common cause. That is, a Reichenbachian common cause. In the particular case above of time dependent events we may further hope, perhaps less realistically, that further conceptual innovations may at some stage provide an adequate framework so as to be able to interpret them as (common) causes.

In sum, the example and the discussion above teach us something both about Sober's example and the actual significance of *common cause completability*. On the one hand, it seems that there are actually good methodological grounds to expect screening-off *common causes* to exist for any given correlation, included those in the 'Venice-Britain' example. On the other hand, we may quite possibly face problems with the interpretation of the screening-off events obtained formally for such correlations. This, although it is not enough, I think, to discard *common cause completability* as a good methodological tool it gives an idea of its limitations and scope of applicability.

7   Purely Formal Correlations

An alternative line of thought would be to deny that the physical significance of the correlations we are trying to explain is at least dubious. This idea would find support in the fact that the screening-off events we have obtained by means of *common cause completability* are difficult, if not impossible, to make sense physically. (Of course 'optimistic' defenders of *common cause completability* would not accept this is so, as I suggested in the previous

section).

Going back to our case, a more thorough analysis of the 'absolute year values' events suggests that the VSL & BBP correlation arises solely due to the time dependence shared by the two systems evolutions. Put it the other way around, we may ask, more specifically, what does the correlation between sea levels and bread prices really say, if anything at all, about the level of the sea in Venice and the price of the bread in Britain?

A closer look to the very events we are dealing with, i.e. 'higher than average values', reveals that these are defined relative to the average of the corresponding quantity over a certain period of time. That is, the correlated events in the 'Venice-Britain' example have some sort of time dependence.

Time dependent data is also known as non-stationary data. It is also well known that non-stationary data display dependencies that do not always reflect a system's inner structure. For instance, it is a consequence of the 'mixing theorem' that two sets of non-stationary data display probabilistic dependencies even if each of them refers to a completely different historical period. Also as a consequence of the 'mixing theorem', if the probabilistic dependence of two data series is due to them being non-stationary the correlation will vanish as soon as we describe the data in a probabilistic model in which one of them is no longer non-stationary. This is in fact what happens with the VSL & BBP correlation, as we have seen in the model-example I outlined in the previous section. We can then conclude that sea levels and bread prices in the VSL & BBL case are only correlated in virtue of telling us something about time, i.e. the correlation does not reflect any information whatsoever about the underlying (physical) structure of the system (if there is any system we can speak of).

The 'Venice-Britain' correlations then seem to be a case of what we can call *purely formal* correlations, i.e. correlations that arise solely as a product of formally modelling experimental data. This is very much along the lines of the argument in Steel (2003). Steel concludes however that, even though RPCC cannot be applied to non-stationary data series, Sober's example constitutes a genuine counterexample to it. For there is a genuine correlation which, even if it is artificial (or 'non-sense'), is not to be explained in terms of a common cause. But, we may ask, should we really demand a common cause explanation when faced with *purely formal* correlations? We can go back to Reichenbach's original intuitions in order to answer this question.

As I said Reichenbach's notion of 'improbable coincidence' is to be captured by the idea of correlation. But do all correlations stand for 'improbable coincidences'? Well, I don't think so. In fact, *purely formal* correlations such as those in the 'Venice-Britain' example do not seem 'improbable' in any sense. Quite the contrary, the structure of the model *formally entails* such coincidences. In this view then, correlations that arise purely formally from the model structure would not be in need of causal explanation. Sober's

example therefore would not constitute a genuine counterexample to RPCC.

Finally, this conclusion may help demarcating further the range of applicability and significance of the *extendability* and *common cause completability* theorems. For, if a given correlation is of purely formal origin, and reflect therefore no physical feature of the system in question, it would seem quite unreasonable to expect it to be screened-off by a physically meaningful event. I think that it is even quite legitimate to suggest that the physically meaningless 'time' events in the previous section explain —although not causally— the correlation, in the sense that it is precisely their lack of physical interpretation that indicates the (formal) character of the correlation. In other words, the physical significance of the *extendability* and *common cause completability* theorems depends crucially on the very nature of the correlations considered as well.

This is also related to a possible objection to the whole argument. Namely, the fact that although it is initially assumed that the events involved in the 'Venice-Britain' scenario, i.e. 'higher than average' events, are real physical events, we have concluded that their correlations are just *purely formal* and do not reflect in any way physical features of the system. Again, what this objection points to is the limitations in the applicability of *common cause completability*. Recall that Hofer-Szabó et al. conceive their formal theorems as tool to effectively discover causal relations in the world. In this sense, the assumption that the events and correlations in the original probability space are real physical events is taken for granted, as it is assumed as well that the result of applying *common cause completability* provides real screening-off common causes. But we have seen this is not always the case. In fact this later assumption may introduce the apparent contradiction raised by in the objection above. An alternative answer to the objection would be to note that Sober himself conceives the 'Venice-Britain' correlations as a conceptual possibility, rather than as a real correlation in the physical world. One could then follow Sober in this and apply *common cause completability* bearing in mind that the screening-off events obtained in doing that will be common causes, also taken as a conceptual possibility.

To conclude, the analysis above suggests that both solutions above may be seen to conflate somehow. For, if *common cause completability* is applied to the 'Venice-Britain' scenario we obtain physically meaningless screening-off events. And this suggests that the 'Venice-Britain' correlation is just an artificial one. Thus, whether 'Venetian sea levels and British bread prices' does constitute indeed a genuine counterexample to RPCC is just a matter of whether or not *purely formal* correlations are in need of (causal) explanation.

References

Cartwright N (1987) How to tell a common cause: Generalizations of the conjunctive fork criterion. In: Fetzer JH (ed) Probability and Causality, Reidel Pub. Co., pp 181–188

Hofer-Szabó G, Rédei M, Szabó LE (1999) On reichenbach's common cause principle and reichenbach's notion of common cause. The British Journal for the Philsophy of Science 50:377–399

Hofer-Szabó G, Rédei M, Szabó LE (2000a) Common cause completability of classical and quantum probability spaces. International Journal of Theoretical Physics 39:913–919

Hofer-Szabó G, Rédei M, Szabó LE (2000b) Reichenbach's common cause principle: Recent results and open questions. Reports on Philosophy 20:85–107

Hoover KD (2003) Nonstataionary time series, cointegration, and the principle of the common cause. The British Journal for the Philosophy of Science 54:527–551

Rédei M (2002) Reichenbach's common cause principle and quantum correlations. In: Placek T, Butterfield J (eds) Non-locality, Modality and Bell's Theorem, Kluwer, pp 259–270

Reichenbach H (1956) The Direction of Time; Edited by Maria Reichenbach. Unabridged Dover (1999) republication of the original University of California Press (1956) publication

Salmon W (1984) Scientific Explanation and the Causal Structure of the World. Princeton University Press

San Pedro I, Suárez M (2009) Reichenbach's principle of the common cause and indeterminism: A review. In: Recio JG (ed) Philosophical essays in physics and biology, Georg Olms

Sober E (1987) The principle of the common cause. In: Fetzer JH (ed) Probability and Causality: Essays in Honor of Wesley Salmon, Reidel, pp 211–228

Sober E (2001) Venetian sea levels, british bread prices, and the principle of the common cause. The British Journal for the Philosophy of Science 52:331–346

Steel D (2003) Making time stand still: A response to sober's counter-example to the principle of the common cause. The British Journal for the Philosophy of Science 54:309–317

REFERENCES

Yule GU (1926) Why do we sometimes get nonsensical relations between time series? a study of sampling and the nature of time series. Journal of the Royal Statistical Society 89:1–64

end