Synthese (2011) 181:3–21
DOI 10.1007/s11229-009-9593-x

Reichenbach's cubical universe and the problem of the external world

Elliott Sober

Received: 13 May 2009 / Accepted: 8 June 2009 / Published online: 2 July 2009

Abstract This paper is a sympathetic critique of the argument that Reichenbach develops in Chap. 2 of *Experience and Prediction* for the thesis that sense experience justifies belief in the existence of an external world. After discussing his attack on the positivist theory of meaning, I describe the probability ideas that Reichenbach presents. I argue that Reichenbach begins with an argument grounded in the Law of Likelihood but that he then endorses a different argument that involves prior probabilities. I try to show how this second step in Reichenbach's approach can be strengthened by using ideas that have been developed recently for understanding causation in terms of the idea of intervention.

Keywords Common cause · Correlation · External world · Likelihood · Realism · Reichenbach · Solipsism

1

Chapter 2 of Reichenbach's *Experience and Prediction* has two big themes, one negative and one positive. The negative theme is Reichenbach's attack on the positivist theory of meaning. The positive theme is Reichenbach's attempt to show why your sense experience justifies your believing that material objects exist independent of experience. The two themes are connected, in that the negative is a necessary prole-

All citations, unless otherwise noted, are from Hans Reichenbach. (1938). *Experience and Prediction—An Analysis of the Foundations and Structure of Knowledge*. University of Chicago Press.

E. Sober (✉)
Department of Philosophy, University of Wisconsin-Madison, 5185 Helen C White Hall, Madison, WI 53706, USA
e-mail: ersober@wisc.edu

gomenon to the positive. This is because the positivist theory of meaning entails that solipsism and realism about material objects are equivalent rather than incompatible. If the positivists were right about this, the problem of the external world would be a pseudo-problem—something to dissolve, not solve.

My subject here is Reichenbach's epistemology, but I do have a few comments on what he says about meaning. In Chap. 1, Reichenbach (pp. 30–31) formulates the positivist theory of meaning as having two parts:

(1) A sentence has meaning if and only if it is verifiable as true or false.
(2) Two sentences have the same meaning if and only if they obtain the same determination as true or false by every possible observation.

By "verifiable," Reichenbach means deducible from a finite set of jointly consistent observation statements. He points out that many of the statements that are taken seriously in science and in ordinary life are not verifiable in this sense. This leads him to liberalize the positivist theory:

(1′) A sentence has meaning if and only if it has a degree of probability.
(2′) Two sentences have the same meaning if and only if they obtain the same degree of probability from every possible observation (p. 54).

Reichenbach's (1′) is intended to parallel the positivist's (1), in that Reichenbach thinks it is sense experience that makes sentences meaningful by conferring degrees of probability on them. This is Reichenbach's probability theory of meaning.[1]

Reichenbach's proposal is an improvement over positivism, but problems remain. First, there is the question of what "probability" should be taken to mean. Reichenbach refers the reader to his earlier (Reichenbach 1935) treatise where he defends a frequency interpretation of the concept (p. 304). This raises the question of how scientific theories like Einstein's theory of general relativity or Darwin's theory of evolution can be said to have a probability in this sense. If there were finitely many universes, we could think of the prior probability of a theory's being true in our universe in terms of the frequency of universes in which the theory is true.[2] Unfortunately, this suggestion comes to nothing if there is no such set of universes. If we abandon Reichenbach's commitment to a frequency interpretation and follow subjective Bayesians by interpreting probability as an agent's degree of belief, this problem disappears but another pops up in its stead—a single theory will have as many different probabilities as there are agents who have different degrees of certainty about it. And a sentence can be meaningful for me though it is meaningless to you, because I assign it a degree of belief and you do not. Shifting to a third interpretation is problematic as well. If we

[1] It is interesting that Reichenbach identifies a sentence's being meaningful with its *having a probability*, not with its *confirmability*. Bayesians say that a sentence is confirmable if the sentence's probability can be changed by new observations; Skyrms (1984) suggests that this provides an account of testability, which I discuss in Sober (2008, pp. 148–154).

[2] This is not how Reichenbach, in the second edition of his 1935 book (pp. 434–442), proposes to understand "the probability of hypotheses." He discusses Newton's law of gravitation by considering a data set that describes how frequently various measurements on planets conform to the law; Reichenbach applies the straight rule of induction to the data to infer a probability for the law. Notice that this is an exercise in the epistemology of probability, not in its semantics. For an argument that the straight rule is nonBayesian, see Sober (2008, pp. 20–24).

think of probability as normative—as the degree of belief that an agent *ought* to have in a sentence, given some body of observations—it again becomes difficult to see what values those probabilities actually have (since their posterior probabilities will depend on choices of values for prior probabilities).

Another problem for Reichenbach's theory stems from the fact that a sentence can be meaningful even if we have no clue as to how probable it is. Consider the sentence "undetectable angels exist." If this really were meaningless gibberish, we could not discuss its epistemic status or its logical relationships to other sentences. But we can.[3] And surely "undetectable angels exist" and "undetectable unicorns exist" differ in meaning, though they seem to have the same epistemic status. Reichenbach's theory, like the positivism he sought to supersede, draws too close a connection between semantics and epistemology, as Putnam (1975) argued so cogently.

Reichenbach says that it is a virtue of his theory that it explains why the following two sentences differ in meaning:

(3) The material world will continue to exist after I am dead.
(4) The material world will cease to exist when I do (p. 133).

In contrast, the positivist theory says they are synonymous—there is no experience I can ever have that will discriminate between them. Reichenbach asks why positivists bother to buy life insurance policies if they believe that the two sentences are equivalent. He says that he does "not doubt the seriousness of the positivists" when they do so; however, "they cannot justify this carefulness" (p. 135). These good points encourage the reader to expect Reichenbach to conclude that the positivist theory of meaning is false while his own theory is true, but this is not what he says. Rather, his view is that it is a matter of choice, not fact, which theory of meaning we adopt—"we cannot forbid anyone to choose the definition of meaning he prefers (p. 149)." To a post-positivist ear, this sounds like backsliding. If the theory of meaning is a theory about the languages and language users we see around us, then it is answerable to empirical considerations. A theory of meaning is false if it says that (3) and (4) are synonymous sentences of English (Carnap 1953; Sober 2000). Of course, the word "synonymous" can be redefined, but that isn't relevant, since any term can be redefined, trivially.

2

It is one thing to say that (3) and (4) differ in meaning, another to show that (3) is more probable than (4), given the evidence I have. In just the same way, the epistemological task of comparing the probabilities of the following two sentences remains even after we grant that they are incompatible:

(5) The only things that exist are my mind and my experiences.
(6) In addition to my mind and my experiences, material objects also exist.

[3] As Glymour (1977, p. 228) says, "it is only because we can understand something of hypotheses and make rough judgments of their implications that we can test them. Not the other way round."

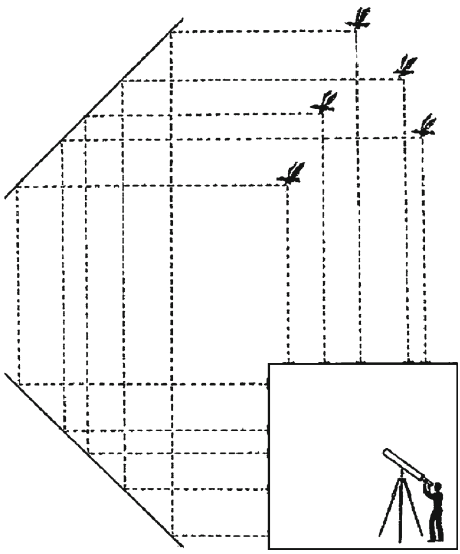

Fig. 1 The observer in the cubical universe sees shadows on the walls that are produced, with the help of mirrors, by the birds outside; from Reichenbach (1938, p. 117)

Proposition (5) expresses the thesis of solipsism, while (6) expresses the thesis of realism about material objects.[4] Reichenbach's question is how my sense experiences justify my postulating entities that are distinct from my mind and the experiences I have. He addresses the difference between (5) and (6) by describing a thought experiment; the picture he uses to explain the example is reproduced here in Fig. 1.

We are asked to imagine

> a world in which the whole of mankind is imprisoned in a huge cube, the walls of which are made of sheets of white cloth, translucent as the screen of a cinema but not permeable by direct light rays. Outside this cube there live birds, the shadows of which are projected on the ceiling of the cube by the sun rays; on account of the translucent character of this screen, the shadow-figures of the birds can be seen by the men within the cube. The birds themselves cannot be seen, and their singing cannot be heard. To introduce the second set of shadow-figures on the vertical plane, we image a system of mirrors outside the cube which a friendly ghost has constructed in such a way that a second system of light rays running horizontally projects shadow-figures of the birds on one of the vertical walls of the cube … As a genuine ghost this invisible friend of mankind does not betray anything of his construction, or of the world outside the cube, to the people within; he leaves them entirely to their own observations and waits to see whether they will discover the birds outside. He even constructs a system of repulsive forces so that any near approach toward the walls of the cube is impossible for the men; any penetration through the walls, therefore, is excluded, and

[4] I take it that solipsism leaves open whether my mind reduces to the experiences I have, or is a distinct entity, and that both (5) and (6) leave open whether phenomenalism is right in its claim that the truth conditions for sentences like "there is a table here now" can be given in a sense datum language.

> men are dependent on the observation of the shadows for all statements they make about the "external" world, the world outside the cube (pp. 115–116).

As is standard, Reichenbach (p. 154) takes the epistemological relationship of objects observed (shadows on the sides of the cube) to objects merely inferred (the birds outside) to be the same as the relationship of your sensations to the physical objects that cause them. He relates these problems to each other by a striking image:

> Our situation with regard to external things is not essentially different from that of the inhabitants of the cubical world with respect to the birds outside; imagine the surface surrounding that world to contract until it surrounds only our own body, until it finally, with some geometrical deformations, becomes identical with the surface of our body—we arrive, then, at the actual conditions for the construction of human knowledge, all our information about the world being bound to the traces which causal processes project from external things to the surface of our body.

Reichenbach (163 ff.) is right that the point of this analogy is not subverted by the fact that we observe physical objects though we rarely, if ever, observe our own experiences. He also is right that the problem of the external world does not require the assumption that our knowledge of our own sensations is absolutely certain; this is inessential, just as the problem of the cubical world does not require that we be absolutely certain about the properties that the shadow-figures have.

Reichenbach concludes his description of the cubical world and its inhabitants with a question: "Will these men discover that there are things outside their cube different from the shadow-figures?" He tells us that, initially, they do not. But after a while, a gifted individual ("a Copernicus")

> …will direct telescopes to the walls and will discover that the dark spots have the shape of animals; and what is more important still, that there are corresponding pairs of black dots, consisting of one dot on the ceiling and one dot on the side wall, which show a very similar shape. If a_1, a dot on the ceiling, is small and shows a short neck, there is a corresponding dot a_2 on the side wall which is also small and shows a short neck; if b_1 on the ceiling shows long legs (like a stork), then b_2 on the side wall shows on most occasions long legs also. It cannot be maintained that there is always a corresponding dot on the other screen but this is generally the case (pp. 117–118).

What interests Reichenbach about these figures on the walls of the cube is the "correspondence" that obtains between the "internal motions" of pairs:

> If the shade a_1 wags its tail, then the shade a_2 also wags its tail at the same moment. Sometimes there are fights among the shades; then, if a_1 is in a fight with b_1, a_2 is always simultaneously in a fight with b_2 (p. 118).

Reichenbach says that the hero of his story, Copernicus,

> will surprise mankind by the exposition of a very suggestive theory. He will maintain that the strange correspondence between the two shades of one pair

> cannot be a matter of chance but that these two shades are nothing but effects caused by one individual thing situated outside the cube within free space. He calls these things "birds" and says that these are animals flying outside the cube, different from the shadow-figures, having an existence of their own, and that the black spots are nothing but shadows (p. 118).

3

Here I will interrupt Reichenbach's story to characterize the two hypotheses he has so far introduced and the bearing of the observations on them. Though my interpretation is close to some of what Reichenbach says, it does not coincide with his in all particulars. The observations are sample frequencies. What we observe, for a given pair of shadow figures i and j, is that

(Obs) $\text{freq}(i \text{ has } F \ \& \ j \text{ has } F) > \text{freq}(i \text{ has } F)\text{freq}(j \text{ has } F)$,

for some range of properties F. I will describe this inequality by saying that we observe that the traits of i and j are *associated* with each other. There are two hypotheses that Reichenbach has so far considered:

(Common Cause) i's having F and j's having F are joint effects of a common cause.
(Coincidence) i's having F and j's having F are causally unconnected.

The point of interest is that,

(7) $\Pr(\text{Obs}|\text{Common Cause}) > \Pr(\text{Obs}|\text{Coincidence})$.

The reason (7) is true is that the common cause model that Reichenbach is considering *entails* that there will be a correlation between i's having F and j's having F, where correlation is a fact about probabilities, not sample frequencies:

(Correlation) $\Pr(i \text{ has } F \ \& \ j \text{ has } F) > \Pr(i \text{ has } F)\Pr(j \text{ has } F)$.

In contrast, the Coincidence model he is considering entails that the two events are probabilistically independent:

(Independence) $\Pr(i \text{ has } F \ \& \ j \text{ has } F) = \Pr(i \text{ has } F)\Pr(j \text{ has } F)$.

In formulating the argument in this way, I am distinguishing observed association from probabilistic correlation. The former involves frequencies in a sample; the latter involves probabilities whose application is not limited to that sample. The hypotheses are, by construction, related deductively to propositions about correlations; they do not have deductive connections to propositions about sample frequencies.

If the Common Cause hypothesis makes the data more probable than the Coincidence hypothesis does, what is the epistemological significance of this fact? The Law of Likelihood (Hacking 1965; Edwards 1972; Royall 1997; Sober 2008) provides an answer:

Observation O favors hypothesis H_1 over hypothesis H_2 precisely when $\Pr(O|H_1) > \Pr(O|H_2)$.

The observed association of the shapes of two shadow-figures favors Common Cause over Coincidence in just the way that the matching of two student essays favors the hypothesis that the students plagiarized from a common source over the hypothesis that they produced their essays independently (Salmon 1984).

One virtue of this likelihood formulation of Reichenbach's argument is that it undercuts the following objection. It is alleged that inferring the existence of external objects is problematic because external objects are objects of a different "kind" from the sensory experiences that provide your evidence. This point about kinds might be relevant if the argument concluded that we are justified in believing that external objects exist. But the Law of Likelihood is not a rule for acceptance. The Law cares nothing for "kinds." The only thing that matters is how different competing hypotheses probabilify the data.

The likelihood reconstruction of Reichenbach's reasoning captures some of what he is getting at when he says that

> … it seems highly improbable that the strange coincidences observed for one pair of dots are an effect of pure chance. It is, of course, not impossible that, when one shade has its shade-tail plucked off, at the same moment the same thing happens to another shade on another plane; it is not even impossible that the same coincidence is sometimes repeated. But it is improbable; and any physicist who sees this will not believe in a matter of chance but will look for a causal connection (pp. 120–121).

However, this passage also suggests that Reichenbach may be venturing beyond what the Law of Likelihood sanctions. Inequality (7) does not show that Common Cause is more probable than Coincidence. It does not follow from (7) that

(8) $\Pr(\text{Common Cause}|\text{Obs}) > \Pr(\text{Coincidence}|\text{Obs})$

unless assumptions are made about the prior probabilities of the two hypotheses.

4

The plot then thickens because Reichenbach sees that his adversary, the positivist, has a reply. The positivist grants that Common Cause and Coincidence "furnish different consequences within the domain of our observable facts" (p. 122) and that the two hypotheses thereby obtain different probabilities. But the positivist then claims that there is a third hypothesis that is predictively equivalent to the hypothesis of Common Cause. This is the hypothesis that the elements within a pair are related as cause to effect:

(Cause/Effect) i's being F causes j's being F, or *vice versa*.

This hypothesis does not postulate the existence of entities outside. Yet, it predicts that there will be a positive association between i's and j's properties, just as the Common Cause hypothesis does. Reichenbach then argues that the Common Cause hypothesis nonetheless is superior to the hypothesis of Cause/Effect. He puts his point in the mouth of "the physicist", who

> … simply states that, wherever he observed simultaneous changes in dark spots like these, there was a third body different from the spots; the changes happened, then, in the third body and were projected by light rays to the dark spots which he used to call shadow-figures … Whenever there were corresponding shadow-figures like the spots on the screen, there was in addition, a third body with independent existence; it is therefore highly probable that there is also such a third body in the case in question (p. 123).

Reichenbach is arguing that there are associated events that occur *inside* the cube that resemble the shadow-images on the cube's surfaces and which have the additional feature that the cube's inhabitants can *observe* what makes the events in a pair similar. The cubists observe that such pairs are *usually* associated because they trace back to common causes and only *rarely* are they associated because they are related to each other as cause to effect.[5]

I see no way to defend this part of Reichenbach's argument. Event types are associated in our world because they are produced by common causes and also because they are related to each other as cause to effect. Both arrangements occur in abundance. For example, consider pairs of birds *inside* our universe; sometimes we see two flying birds change direction in unison because a predator approaches them both and sometimes we see them do so because one bird is chasing the other. Each may look like a "dark spot" that we see against the sky. In general, when we observe that event types e_1 and e_2 are associated because they have a common cause c, we also observe that c and e_1 are associated with each other because they are related to each other as cause to effect, and the same is true of the association of c to e_2. For each Common Cause pattern that we observe, there are two Cause/Effect patterns that we also observe. How, then, could the former arrangement be more frequent? I take it that the same is true of what goes on in the cubical universe. Furthermore, even if Reichenbach were right about what is common and what is rare in the observations that the cubists make inside their universe, it is hard to see how this point would carry over to the problem of the external world. Cubists who see the shadow-images on the cube's surfaces also can observe what is happening in the cube's interior. But when we notice associations between our sensations, we cannot drop back to another sort of association whose causal origins are plain to us.

Reichenbach has a footnote on the next page (p. 124) that makes it clear that we are not here mistaking what his argument is. He says that the Common Cause hypothesis has a higher posterior probability than the Cause/Effect hypothesis because the former has the higher prior and in spite of the fact that the two hypotheses confer on the observed association the same probability (they have identical likelihoods). It isn't the observed association of the shadows that makes it more probable that birds exist outside the cube than that the shadows cause each other; rather, what does the work, according to Reichenbach, is the fact that the existence of birds outside has the higher prior probability.

[5] It is interesting that Reichenbach does not argue that the shadow-images in a pair cannot be related as cause to effect on the grounds that they occur simultaneously (and via appeal to the principle that cause must precede effect). He does say a few times that the events are simultaneous, but his diagram of the cubical world suggests that this would not be exactly true.

5

Reichenbach's epistemology of the cubical universe calls to mind his now-famous Principle of the Common Cause, which he states as follows in his posthumously published book, *The Direction of Time*:

> If an improbable coincidence has occurred, there must exist a common cause … Chance coincidences, of course, are not impossible … The existence of a common cause is therefore … not absolutely certain, but only probable (Reichenbach 1956, pp. 157–158).

Reichenbach does not cite this general principle in *Experience and Prediction*. If it were correct, it would provide a simple rationale for concluding that the correlated shadow-images have a common cause (Salmon 1999). However, a further argument would be needed to show that this common cause is outside the cube, let alone a bird outside. In any event, the common cause principle has problems, so it is good that Reichenbach does not appeal to it in his 1938 book. First, notice that Reichenbach gives two formulations of the principle in the passage quoted above; the first says that a common cause *must* exist while the second says that it *probably* does. The first is an overstatement, since, as Reichenbach notes, improbable coincidences without a common cause are not impossible.[6] The probabilistic formulation, on the other hand, requires that you have prior probabilities for common cause hypotheses, and, as we have seen, these are sometimes difficult to justify. In addition, the principle takes you from correlations to (probable) common causes without pausing to consider, as Reichenbach does in his discussion of the cubical universe, that a correlation might be due to the correlates being related as cause to effect. It therefore is better to formulate Reichenbach's principle as saying that correlated events must be (or probably are) *causally connected*, where this means that the events either have a common cause *or* are related as cause to effect. But it is even better to formulate the principle as an instance of the Law of Likelihood.[7] Unfortunately, this fails to settle why the hypothesis of Common Cause is better than the hypothesis of Cause-Effect.

6

In 1936, two years before the publication of *Experience and Prediction*, Einstein noted that the general theory of relativity predicts that an astronomical object can act as a "gravitational lens" and produce a "gravitational mirage," as shown in Fig. 2. If a massive object lies between the observer and a distant object, light from the distant object will bend around the intervening object, leading the observer to see two (or more) images, not one. This possibility was described qualitatively by the St. Petersburg

[6] Abandoning the "must" in stating the principle of the common cause suffices to ensure that the Bell inequality results are not a counterexample.

[7] Even when modified in this way, the principle still encounters a problem first described by Yule (1926); the mechanisms behind two causally unconnected time series sometimes entail cross-process correlations. See Sober (2001, 2008) for discussion.

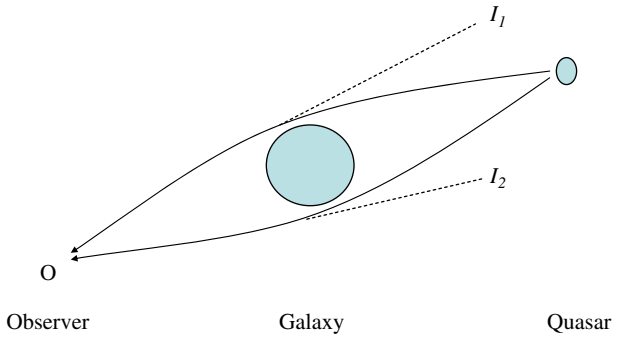

Fig. 2 When light from a quasar bends around a lensing galaxy, an observer will see two images (I_1 and I_2) of the quasar

physicist Orest Chwolson in 1924. In 1936 Einstein considered the example of a single star serving as a lens and concluded that the effect would be too small for current instrumentation to detect. In 1937, Fritz Zwicky realized that a whole galaxy could act as a lens and that the more pronounced image doubling should be easy to detect. The first real example was discovered in 1979 by Dennis Walsh, Bob Carswell, and Ray Weymann using the Kitt Peak National Observatory 2.1 meter telescope. It was called the "Twin Quasar" because it initially looked like two identical quasars; its official name is Q0957 + 561.[8]

Did Reichenbach know about Chwolson's, Einstein's, and Zwicky's ideas? If so, did he consider using the idea of gravitational lensing in *Experience and Prediction* as a real-world example that makes the same point as the shadow images on the walls of his fanciful cubical universe? And even if Reichenbach did not consider this way of telling his story, would real science have served his purpose better than science fiction? I will not try to answer the biographical questions, but I do want to say a little about how gravitational lensing is related to the problem that Reichenbach describes.

We see what appear to be two astronomical objects in different parts of the sky. Our measurements show that they are identical in numerous respects. We consider and reject the possibility that this is just an elaborate coincidence. There must be a causal connection. Either one object causes the other, or the two are effects of a common cause. Both these hypotheses predict the similarities we observe. A difference between the hypotheses may be found in the fact that we have an independently confirmed theory (the General Theory of Relativity) that says that the Common Cause pattern should be exemplified; on the other hand, no current theory describes a process in which one huge astronomical object can cause a carbon copy of itself to exist someplace else. This provides a justification for assigning Common Cause a higher prior probability than Cause/Effect. The prior does not come from *a priori* reasoning, but is grounded by an empirical theory. This appeal to prior probabilities to deal with astronomical twins seems much less artificial than Reichenbach's appeal to prior probabilities when he discusses twin shadows in his cubical universe. Even so, there is something missing

[8] See the Wikipedia article at http://en.wikipedia.org/wiki/Gravitational_lensing.

in this example from physics. Granted, a background theory can furnish us with defensible priors, but what if we have no such theory? This is where the example of the cubical universe has a utility that gravitational lensing does not.

7

Let us return to Reichenbach's claim that Common Cause and Cause/Effect confer the same probability on the observations that Copernicus makes in his cubical universe. Do you really need to appeal to prior probabilities to find an epistemic difference between the two hypotheses? Present-day causal modelers have thought about this problem (Hausman 1998; Woodward 2003). We can take our cue from them.

Here is a simple and shopworn example. Why do we think that barometer readings and storms are joint effects of a common cause (the barometric pressure) rather than being related to each other as cause to effect? The readings and the rain are associated in our data, in that

freq(storm at t_2| low barometer reading at t_1) $>$ freq(storm at t_2| high barometer reading at t_1).

This association is evidence that there is an underlying probabilistic correlation:

(9) Pr(storm at t_2| low barometer reading at t_1) $>$ Pr(storm at t_2| high barometer reading at t_1).

If the two events were joint effects of a common cause, that would explain the observed association, but the association also would be explained if barometer readings caused storms. To discriminate between these two hypotheses, there is no need to appeal to prior probabilities; rather, a different type of fact can be brought to bear. *Manipulating* the barometer reading does not change the probability of a storm:

(10) Pr(storm at t_2| low barometer reading at t_0 & I make the barometer read low at t_1)
$=$ Pr(storm at t_2| low barometer reading at t_0 & I make the barometer read high at t_1).

The manipulation at t_1 takes place shortly after t_0, but before t_2. Manipulating the barometer reading just after the barometer reads high exhibits the same pattern:

Pr(storm at t_2| high barometer reading at t_0 & I make the barometer read low at t_1)
$=$ Pr(storm at t_2| high barometer reading at t_0 & I make the barometer read high at t_1).

Although Common Cause and Cause/Effect both predict that low barometer readings and storms will be associated, the two hypotheses make different predictions about what will happen when we perform manipulations.

The concept of manipulation (or intervention) needs to be specified carefully. It is a causal process (not necessarily one initiated by a conscious agent), but not just any way of causing the barometer to read low counts as an intervention. Interventions must be surgical, not ham-fisted. The concept is defined relationally; we need to define what it takes for I to be an intervention on the barometer reading (R) relative to the

storm (S). The key idea is that any effect that I has on S must pass through R; there can be no second pathway by which this influence is transmitted (nor can I be correlated with a factor Z that influences S by a separate pathway). For example, if you perform an action that simultaneously makes the barometer read low and also raises the barometric pressure, your action is not an intervention.[9] In this circumstance, your act of ham-fistedly changing the barometer reading is associated with an increase in the frequency of storms, but that isn't evidence that barometer readings cause storms.

If barometer readings cause storms, then interventions on the readings can be expected to be associated with storms. If the two are joint effects of a common cause, then interventions on the barometer readings can be expected to not be associated with storms. If no such association is observed, this is evidence favoring Common Cause over Cause/Effect. I say that interventions provide data that favor one hypothesis over another, not that they provide data that definitively refutes either. As noted earlier, the hypotheses under discussion do not have deductively entailments as to whether there will be *associations* in the data; rather, they have entailments about whether there will be probabilistic *correlations*. These correlations are not what we observe; what we observe are associations.

Reichenbach's thought experiment about the cubical world includes the detail that the people in the cube cannot manipulate the shadow-images that appear on the cube's surfaces. As quoted above, he stipulates that powerful force fields prevent them from doing so. This makes the puzzle harder to solve. Had Copernicus or his successors sent rockets to alter or remove one of the shadow-images in a pair (perhaps by smearing whitewash on it), these manipulations would have yielded evidence that favors Common Cause over Cause/Effect.[10]

How does this point about manipulations bear on the problem of the external world? Each of us can manipulate sensations that we find correlated. In my experience, the visual impression (V) that there are waves breaking on a beach is associated with the auditory impression (A) of the crashing of waves. Are these two types of experience related as cause to effect or do they trace back to a common cause? The answer can be found by manipulating each. When I have both experiences, I can easily shut down the one but not the other, and shut down the other but not the one. When I open and close my eyes, my visual impressions flip in perfect synchrony though I continue to have the same auditory impression. When I hold my ears and release them, my auditory impressions stop and start though I continue to have the same visual impressions.[11] The visual (V) and auditory (A) impressions are probabilistically correlated, but manipulating the one does not change the probability of the other. That is, it is true both that

[9] See Woodward (2003, pp. 98–103) for further details on how the concept of intervention should be defined.

[10] Reichenbach (1925, pp. 101–102) anticipates this point about manipulations; he says that "if $A$is the cause of C and a slight alteration is induced in A, a slight alteration will also appear in C. But if C is altered slightly, no change will arise in A." Perhaps he does not use this idea in *Experience and Prediction* because he thought it was not needed to solve the problem at hand.

[11] I am not advancing the strong thesis that there is no manipulation of visual experience that affects audition. My claim is more modest. See McGurk and McDonald (1976) for discussion of a case in which vision influences what we hear.

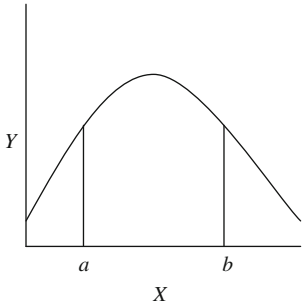

Fig. 3 Even though X causes Y, manipulating X by changing its value from $X = a$ to $X = b$ is not associated with an expected change in the state of Y

(11) Pr(I have A at t_2| I have V at t_1) > Pr(I have A at t_2| I don't have V at t_1)

and

(12) Pr(I have A at t_2| I have V at t_0 and I cause V to occur at t_1 by keeping my eyes open)
= Pr(I have A at t_2| I have V at t_0 and I cause V to not occur at t_1 by closing my eyes).

The Common Cause and Cause/Effect hypotheses agree that (11) should be true, but they disagree about whether (12) is.[12]

The operative principle, both with respect to the barometer and with respect to the beach,
is:

> X causes Y if and only if there is a possible intervention on X that entails an expected change in the state of Y.

This is what distinguishes X's causing Y from Y's causing X and from X and Y's being joint effects of a common cause. There is an existential quantifier in this criterion, not a universal quantifier, because some interventions on a cause can fail to be associated with an expected change in its effect. Consider, for example, the relationship of a cause X to its effect Y that is depicted in Fig. 3. If X is manipulated by changing its value from $X = a$ to $X = b$, the expected value of Y remains the same. But notice that intervening on X by changing its value from $X = a$ to $X = c$ (where $c \neq b$) entails that the expected value of Y will be different.

In general, when X is manipulated and no change in the state of Y occurs, this is evidence, not proof, that favors the hypothesis that X and Y have a common cause over the hypothesis that they are related to each other as cause to effect. It is only evidence for two reasons. First, as already noted, what you observe is a fact about frequencies, not a probabilistic equality. Second, even if a probabilistic equality is true for the particular manipulation you have performed, it may fail to be true for others. When

[12] Consider Locke's comment in Book 4, Chap. 11 of his *Essay Concerning Human Understanding*: "our senses assist one another's testimony of the existence of outward things, and enable us to predict. Our senses in many cases bear witness to the truth of each other's report, concerning the existence of sensible things without us. He that sees a fire, may, if he doubts whether it be anything more than a bare fancy, feel it too; and be convinced, by putting his hand in it."

a manipulation fails to produce a difference in observed frequencies, this is evidence that favors Common Cause over Cause/Effect in the sense described by the Law of Likelihood. The observations are less probable if X is a cause of Y than they would be if X were not. Fiddling with the barometer fails to be associated with a change in the frequency of storms. And my fiddling with my visual impressions of waves by fluttering my eyes fails to be associated with a change in the frequency of my auditory impressions of waves. The observations discriminate between Common Cause and Cause/Effect; you do not need prior probabilities to find a difference between the two hypotheses.

8

Does this successfully complete Reichenbach's argument for the external world? I hope it is clear that the pattern pertaining to my auditory and visual impressions of waves has wide applicability. I do not claim that the argument I have described establishes that there *probably* is a common cause. The more modest thesis to consider is that the associated sensations and the results of my manipulation experiments are *evidence* for this conclusion in the sense described by the Law of Likelihood. But even with this caveat, this evidential argument for the external world[13] is not yet complete. How do I know that my shutting my eyes and my holding my ears are manipulations in the required sense? And why must the common cause for which I claim to have evidence be something outside my own mind?

The worry represented by the first question is depicted in Fig. 4. Perhaps my opening and closing my eyes (E) is not a manipulation of my visual impressions V with respect to my auditory impressions A in the required technical sense. My eye fluttering (E) will be a ham-fisted nonintervention if it affects my seeming to hear the waves (A) by some pathway other than the one that goes through my visual impressions (V), or if E is associated with a causal factor (Z) that affects my auditory impressions by some other pathway. Let's focus on the first of these options, as the issues raised are the same. The problem is that it is possible for V to cause A even though E and A are uncorrelated. This will happen if E raises the probability of A along one pathway and lowers it along the other, where the magnitudes of these two components exactly cancel each other, so that the net effect of E on A is zero. Some causal modelers (Spirtes et al. 2001) adopt an assumption (the "faithfulness" condition) that rules this out. But even without this assumption, the observations still favor Common Cause over Cause/Effect if there is some chance that my fluttering my eyes is an intervention on V with respect to A.

The second question does not contest the thesis that E is an intervention on V with respect to A nor does it deny that the result of this experiment favors the hypothesis that V and A have a common cause. Rather, the question is why I should think that this common cause of my visual and auditory impressions of waves is something outside

[13] I call this the evidential argument for the external world to parallel the evidential argument from evil, which claims, not that the evil we observe *proves* that there is no God, but that it is *evidence against* the existence of God; see the articles collected in Howard-Snyder (1996) and the brief discussion in Chap. 2 of Sober (2008, pp. 164–167).

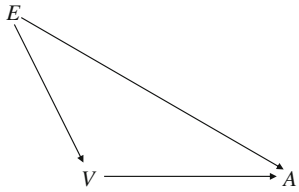

Fig. 4 Suppose E (my closing my eyes) is not an intervention on V (my visual impressions) with respect to A (my auditory impressions) because I not only affects V, but also affects A via a distinct pathway. I is therefore hamfisted

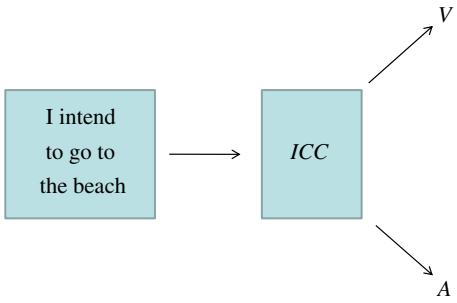

Fig. 5 My intending to go to the beach is a common cause of the visual (V) and auditory (A) experiences of waves that I have. If the intending does not screen off V from A , this is evidence for there being an intervening common cause (an ICC)

my own mind.[14] For example, why cannot the common cause just be the intention I form to go to the beach?

The answer is that the intention *is* a common cause, but that this does not mark a victory for solipsism. The first thing to note is that, in fact, the intention does not screen-off the two types of experience from each other (i.e., it does not render them conditionally independent of each other), since

$$\Pr(V \text{ at } t_2 \text{ and } A \text{ at } t_2 | \text{ I intend to go to the beach at } t_1) > \Pr(V \text{ at } t_2 | \text{ I intend to go to the beach at } t_1)\Pr(A \text{ at } t_2 | \text{ I intend to go to the beach at } t_1).$$

This probabilistic inequality is sanctioned by frequency data. For example, suppose that I do not always have wavy visual and auditory experiences after I form the intention to go to the beach, but that when I have one of these experiences I always have the other. If V and A are correlated even after I take account of my intending to go to the beach, how is this failure of screening-off to be explained? One possibility is depicted in Fig. 5; perhaps there is an *intervening* common cause, one that comes between my intending to go to the beach at $t_1$and my having wavy auditory and visual sensations at t_2.[15] But what is the nature of this intervening common cause? The realist will say that it is a state of the world external to my own mind—perhaps it is the state of my actually being at the beach. Solipsists will reply that there is a state of my own mind that screens off V from A.

[14] This question recapitulates the one I asked earlier about the cubical universe: if the shadow images in a pair have a common cause, why must that common cause be outside the cube, let alone a bird?

[15] In a causal chain from the more distal common cause C_d to the more proximal common cause C_p to the two effects V and A, C_d will fail to screen off V from A if (i) C_p screens off V from A, (ii) C_p screens off V from C_d, (iii) C_p screens of A from C_d, and (iv) all the probabilities involved are between 0 and 1 noninclusive (Sober 1988).

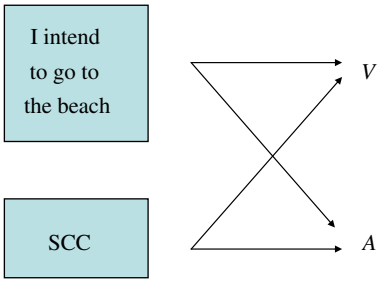

Fig. 6 If my intending to go to the beach fails to screen off my wavy visual (V) and auditory (A) sensations from each other, one possible explanation is that there is a second common cause at work

A second explanation for why my intending to go to the beach does not screen off my wavy auditory and visual experiences from each other is that there is a second common cause (not caused by the first), as depicted in Fig. 6.[16] Here again, the question arises of what kind of state this second common cause is. The solipsist will say that it must be a state of my own mind; the realist will deny this.

I take it that the solipsist must concede that there is evidence that *something* is going on in the production of my wavy visual and auditory sensations besides my intending to go to the beach. The challenge for the solipsist is to point to some other mental state that I occupy that plays the role of this something else. I suggest that often there is nothing of the sort to which the solipsist can point. Suppose, for example, that right after I form the intention to go to the beach that I am rendered unconscious; the next thing I know, I am either experiencing wavy visual and auditory experiences, or I am experiencing neither. When I introspect, I find no further experiences that I can cite to explain this uncanny correlation of V and A. Realists will maintain that I not only have evidence for a common cause, but also have evidence that this common cause is not one of my mental states. "It's the external world that is doing the work, stupid," they will impatiently insist.

Can the solipsist reply that there is a mental state that screens off V and A from each other, but that this is something to which I have no introspective access? I concede that this reply is possible. However, it undermines the epistemological idea that is supposed to motivate solipsism in the first place. The solipsist thinks it is clear that his own mind and his experiences exist, but then he questions whether there is anything else. This clarity and query are supposed to reflect the fact that I am immediately aware of my own mind and mental states—I can have no doubt about their existence—but the physical world is something more iffy. If the solipsist is driven to countenance introspectively inaccessible mental states, what is the basis for his doubting the existence of an external physical world?

[16] If there are two common causes of V and A where each makes a difference in the probability of the two effects, and their joint states screen off V from A, then neither common cause screens-off V from A by itself (Sober 1988).

9

Does the inference from associated sensations to an extra-mental common cause beg the question against solipsism? Empiricists have objected to defenses of scientific realism that appeal to "inference to the best explanation" on this ground (Fine 1984). One difference is that the argument for the external world presented here is evidential and rests just on the Law of Likelihood, whereas the realist's inference to the best explanation involves acceptance. Even so, it is open to the solipsist to reject the Law of Likelihood and therefore to reject Bayesianism. The question then arises of which probabilistic principles for interpreting evidence solipsists are prepared to embrace. If solipsists say "none at all"—that the only inference principles they countenance are deductive—that is reason enough to reject their epistemology, since it prevents them from holding that present and past experiences provide evidence concerning which experiences will occur in the future.[17]

10

The positivists thought that the problem of the external world is a pseudo-problem mired in metaphysical muck. Believing as he did that solipsism and realism about material objects are predictively equivalent, Carnap (1956) concluded that the choice between them is a matter of convenience, not fact; this is what he meant by classifying this philosophical issue as an *external* question, not an *internal* question. Although Reichenbach's probability theory of meaning allows the problem of the external world to be resuscitated, we now have the option of regarding that theory as a ladder that we can kick away once we have climbed it. Solipsism and realism about material objects are *not* synonymous; we can recognize this without needing to have a fully adequate theory of meaning at our fingertips. Reichenbach was right that probabilistic tools can be brought to bear on the epistemological problem of evaluating solipsism and realism as competing hypotheses.[18] We are no more cut off from solving this problem than scientists are cut off from reasoning about the existence of unobservable entities.[19] However, the hypothesis that there is something outside need not be assessed by invoking prior probabilities. These priors are difficult to justify, and placing the entire burden of the argument on the assignment of prior probabilities has the implication that our experiences provide no evidence, one way or the other, for an external world.

[17] I argue that the realism/empiricism debate should be understood as a debate about evidence, rather than as a debate about acceptance, in Sober (2007).

[18] It is curious that Reichenbach's animus against positiv*ism* occasionally manifests itself as *ad hominem* slams against positiv*ists*. For example, in Chap. 2, he says that "the preachers of positivism" remind him of "the fanaticism of a religious sect" (p. 103); he also comments that positivists "usually become offended when they are told that they do not believe that the physical world will continue to exist after their death" (p. 134). Who are these dogmatic and prickly positivists whom Reichenbach found so irritating? Presumably not the gentle Carnap; maybe the testy Neurath or the messianic Ayer or the recently murdered Schlick?

[19] An analog for the problem of the external world that includes a role for manipulations can be extracted from Hacking (1985) discussion of how microscopists distinguish properties of a specimen from artifacts of the apparatus.

Reichenbach's cube is a beautiful analogy for the problem of the external world. And his focus on the correlations that exist between shadow images on the walls of the cube directs our attention to a key issue—the correlations that exist between the different experiences that we have. But there is a disanalogy that needs to be corrected. Because of powerful force fields, the cubists are passive observers of the images on the walls of their cube, just as the chained prisoners in Plato's cave are passive observers of the shadows cast there. In contrast, each of us is related to his or her experiences actively, not passively. We influence which experiences we have by acts of will. This provides us with experimental opportunities that the cubists cannot exploit. Perhaps this is enough to show that I have evidence, not just that some of my experiences have common causes, but that those common causes are not states of my consciousness.

Acknowledgements I thank Martin Barrett, David Chavez, Juan Comesaña, Richard Creath, Joshua Filler, Clark Glymour, Guven Guzeldere, Daniel Hausman, Gurol Irzik, John Koolage, Flavia Padovani, Maria Reichenbach, Alan Richardson, Carolina Sartorio, Friedrich Stadler, Michael Stoelzner, Michael Strevens, Steve Swartzer, and James Woodward for useful comments. I also have benefitted from discussing this paper at the 2008 Reichenbach conference at Bogazici University in Istanbul and at the University of Nebraska, Lincoln.

References

Carnap, R. (1953). Meaning and synonymy in natural languages. *Philosophical Studies, 6*, 33–47.

Carnap, R. (1956). Empiricism, semantics, and ontology. In *Meaning and necessity*. Chicago: University of Chicago Press.

Edwards, A. (1972). *Likelihood*. Cambridge: Cambridge University Press.

Fine, A. (1984). The natural ontological attitude. In J. Leplin (Ed.), *Scientific realism* (pp. 83–107). Berkeley: University of California Press. (Reprinted in *The shaky game*, pp. 112–135. Chicago: University of Chicago Press).

Glymour, C. (1977). Reichenbach's entanglements. *Synthese, 34*, 219–235.

Hacking, I. (1965). *The logic of statistical inference*. Cambridge: Cambridge University Press.

Hacking, I. (1985). Do we see through the microscope? In P. Churchland & C. Hooker (Eds.), *Images of science: Essays on realism and empiricism*. Chicago: University of Chicago Press.

Hausman, D. (1998). , Churchland, P., Hooker, C. (Eds.), *Causal asymmetries*. Cambridge: Cambridge University Press.

Howard-Snyder, D. (Ed.). (1996). *The evidential argument from evil*. Bloomington: Indiana University Press.

McGurk, H., & McDonald, J. (1976). Hearing lips and seeing voices. *Nature, 264*, 746–748.

Putnam, H. (1975). *Mind, language, and reality*. Cambridge: Cambridge University Press.

Reichenbach, H. (1925). Die Kausalstruktur der Welt und der Unterschied von Vergangenheit und Zukunft. In *Sitzungsberichte der Bayerische Akademie der Wissenschaft*, November, pp. 133–175. (Translated as The causal structure of the world and the difference between past and future. In R. Cohen & M. Reichenbach (Eds.), *Hans Reichenbach—Selected Writings, 1909–1953* (1978, pp. 81–119)). Dordrecht: Reidel.

Reichenbach, H. (1935). *The theory of probability—an inquiry into the logical and mathematical foundations of the calculus of probability*, 2nd edn. Berkeley: University of California Press, 1948.

Reichenbach, H. (1938). *Experience and prediction—an analysis of the foundations and structure of knowledge*. Chicago: University of Chicago Press. (Reissued 2006 by University of Notre Dame Press with an introduction by A. Richardson).

Reichenbach, H. (1956). *The direction of time*. Berkeley: University of California Press.

Royall, R. (1997). *Statistical evidence—a likelihood paradigm*. Boca Raton, FL: Chapman and Hall.

Salmon, W. (1984). *Scientific explanation and the causal structure of the world*. Princeton: Princeton University Press.

Salmon, W. (1999). Ornithology in a cubical world: Reichenbach on scientific realism. In D. Greenberger, W. Reiter, & A. Zeilinger (Eds.), *Epistemological and experimental perspectives on quantum physics* (pp. 303–315). Dordrecht: Kluwer.

Skyrms, B. (1984). *Pragmatics and empiricism*. New Haven: Yale University Press.

Sober, E. (1988). The principle of the common cause. In J. Fetzer (Ed.) *Probability and causation: Essays in honor of Wesley Salmon* (pp. 211–228). Dordrecht: Reidel. (Reprinted in E. Sober, *From a biological point of view*. Cambridge: Cambridge University Press, 1994).

Sober, E. (2000). Quine's two dogmas. In *Proceedings of the Aristotelian Society*, Vol. 74, pp. 237–280.

Sober, E. (2001). Venetian sea levels, British bread prices, and the principle of the common cause. *British Journal for the Philosophy of Science, 52*, 1–16.

Sober, E. (2007). Empiricism. In S. Psillos & M. Curd (Eds.), *The Routledge companion to the philosophy of science* (pp. 129–138). London: Routledge.

Sober, E. (2008). *Evidence and evolution—the logic behind the science*. Cambridge: Cambridge University Press.

Spirtes, P., Glymour, C., & Scheines, R. (2001). *Causality, prediction, and search*. Cambridge: MIT Press.

Woodward, J. (2003). *Making things happen: A theory of causal explanation*. Oxford: Oxford University Press.

Yule, G. U. (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.