



CHICAGO JOURNALS



The Common Cause Principle in Historical Linguistics

Author(s): Christopher Hitchcock

Source: *Philosophy of Science*, Vol. 65, No. 3 (Sep., 1998), pp. 425-447

Published by: [The University of Chicago Press](#) on behalf of the [Philosophy of Science Association](#)

Stable URL: <http://www.jstor.org/stable/188279>

Accessed: 08/04/2014 14:09

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



The University of Chicago Press and Philosophy of Science Association are collaborating with JSTOR to digitize, preserve and extend access to Philosophy of Science.

<http://www.jstor.org>

The Common Cause Principle in Historical Linguistics*

Christopher Hitchcock^{†‡}

Division of Humanities and Social Sciences, California Institute of Technology

Despite the platitude that analytic philosophy is deeply concerned with language, philosophers of science have paid little attention to methodological issues that arise within historical linguistics. I broach this topic by arguing that many inferences in historical linguistics conform to Reichenbach's common cause principle (CCP). Although the scope of CCP is narrower than many have thought, inferences about the genealogies of languages are particularly apt for reconstruction using CCP. Quantitative approaches to language comparison are readily understood as methods for detecting the correlations that serve as premises for common cause inferences, and potential sources of error in historical linguistics correspond to well-known limitations of CCP.

1. Introduction. In a high school philosophy class, one of my fellow students asked the teacher why all of the philosophers we were studying were long dead. Why no twentieth-century philosophers? The teacher, always an opinionated man, replied that twentieth-century philosophy was obsessed with language, rather than with the world that it represents. To him, this was like looking out at the world through a window, and examining in minute detail the cracks and scratches on the window pane. Like all caricatures, this is an exaggeration that contains a grain of truth. Philosophy of language is an area of study that has blossomed in the twentieth century, and occupies a central place in contemporary analytic philosophy. Many philosophers themselves know a good deal

*Received July 1997, revised December 1997.

[†]Send reprint requests to the author, Division of Humanities and Social Sciences, MC 101-40, California Institute of Technology, Pasadena CA 91125; e-mail: cricky@caltech.edu.

[‡]I would like to thank audience members at the Society for Exact Philosophy annual meeting in Montréal, where an earlier version of this paper was presented. For comments upon earlier drafts, I would like to thank David Hull, Merrilee Salmon, Wes Salmon, Elliott Sober, and especially Alexis Manaster Ramer.

Philosophy of Science, 65 (September 1998) pp. 425–447. 0031-8248/98/6503-0005\$2.00
Copyright 1998 by the Philosophy of Science Association. All rights reserved.

about linguistic theory, particularly in the areas of syntax, semantics, and pragmatics. Indeed, semantics and pragmatics would rightly be said to belong to the intersection of philosophy and linguistics.

In recent philosophy of science, there has been a growing interest in the methodology of genealogical inference. This has been sparked, in no small measure, by the work of Elliott Sober. His book *Reconstructing the Past* (Sober 1988), for which he won the prestigious Lakatos prize, brought to philosophers' attention the methodological debates within *systematics*, that part of biology that deals with the classification of species into higher taxa.

Given these two interests, it is surprising that there has been little philosophical interest in the problems attending the reconstruction of linguistic family trees—the provenance of historical linguistics. This is a field that has seen its share of controversy, often over methodological issues that should attract the interest of philosophers. It is a primary goal of this paper to draw philosophers' attention to some of these issues. I will argue that an important type of inference in historical linguistics conforms to the Common Cause Principle (hereafter CCP) familiar to philosophers. In Section 2 I introduce CCP, and make some observations upon a standard illustration. This is followed by an account of inferences in historical linguistics, which are contrasted with seemingly similar inferences in evolutionary biology. I will argue that the former conform to CCP in a manner that the latter do not. In Sections 4 and 5 I discuss a proposal for implementing common cause inferences in historical linguistics, and argue that limitations to the CCP that have been discussed in the abstract by philosophers correspond to textbook caveats regarding inference in historical linguistics. I conclude by pointing to some outstanding problems.

2. The Common Cause Principle. Since its original presentation in Reichenbach's *Direction of Time* (Reichenbach 1956), CCP has been given many different formulations, to various aims. Sometimes it is presented primarily as a thesis about statistical relationships, sometimes as a thesis about causation, or about explanation, or about temporal asymmetry. I will not try to disentangle from this mess the 'true' common cause principle: I will simply provide a formulation that is germane to the current project.

Suppose that A and B are event types that are positively correlated, i.e., such that if AB represents the joint occurrence of events of type A and B , then $P(AB) > P(A)P(B)$. According to CCP, if A is not a cause of B nor vice versa, it is reasonable to postulate a common cause C of A and B , which explains the correlation. Reichenbach added the requirement that the common cause C screen off A and B , i.e., that

$P(AB|C) = P(A|C)P(B|C)$. This is a standard requirement (although see Salmon 1984, Ch. 6, for a dissenting view). This further requirement will have little bearing on what follows. I will not assume that CCP is a universally valid form of inference, nor that it is a universally correct generalization about the relationship between causation and probabilistic correlation. There is strong reason to think that CCP fails in the realm of microphysics, at least. I assume only that there is a range of cases in which inferences conforming to the CCP schema are reasonable ones. (For further discussion of some of the difficulties with CCP, see Sober 1988, Ch. 3, and Arntzenius 1993. We will return briefly to this issue in Section 6 below.)

Here is an (embellished version of an) illustration of CCP provided by Reichenbach (1956, 157). Suppose that a traveling theater troupe is on tour for one thousand days. On ten of those days, one of the male actors, call him the 'leading man', suffered severe gastric distress. Likewise, on ten days, one of the actresses—the leading lady—was similarly indisposed. On nine days, *both* the leading man and the leading lady had stomach illnesses. Assume that these frequencies reliably indicate underlying probabilities, and let A represent the leading man's being sick on a particular night, B the leading lady's being sick, and AB both being sick. Then $P(A) = P(B) = .01$, and $P(AB) = .009$, so A and B are probabilistically correlated. It would be reasonable, according to the common cause principle, to postulate a common cause, such as shared meals containing tainted food, to explain this correlation. This common cause occurs on some nights, but not on others; on those nights it does occur, it dramatically increases the probability that each actor will become sick. The occasional occurrence of this common cause explains the excess of nights on which both are sick, compared to what would be expected by chance.

I wish to make a number of observations upon this example.

(i) Although we assumed that the frequencies reliably indicate probabilities in this example, it is correlations among the probabilities themselves, not the frequencies, that bespeak common causes. In practice, correlations among probabilities are not simply observed, but must be inferred, typically on the basis of observed frequencies. Sometimes these inferences are warranted, sometimes not. Mere correlations among frequencies—call these statistical correlations, in contrast to probabilistic correlations—do not necessarily warrant the postulation of common causes. Indeed, we would expect any sufficiently rich set of data to contain some statistical correlations, no matter how the data were generated. Obviously not all such statistical correlations demand common cause explanations.

(ii) In this example, CCP warrants us in postulating a certain expla-

nation of the probabilistic correlation between two types of event. This correlation is explained by a type of event that occurs with some non-extreme probability. If the leading man and leading lady shared tainted food *every* night, and this common cause screened off the leading man's being sick from the leading lady's being sick, then there would have been no probabilistic correlation between these two event types. In this case, the event type that is postulated as the common cause—shared meals with tainted food—is instantiated within a sequence of instantiations of a broader event type: shared meals generally. That is, it is because the leading man and leading lady share meals in general that they occasionally share meals containing tainted food.

CCP is sometimes taken as a principle for postulating explanations of *particular* joint occurrences. For example, if both the leading lady and the leading man became ill on a particular night, it might be reasonable to postulate a particular common cause of both illnesses. In general, however, this type inference to a particular common cause is only partially warranted by CCP as illustrated in this example. Suppose, to change the example somewhat, that each actor had been sick on 500 nights, and that both had been sick on 290 nights. Assuming that these frequencies are representative of the underlying probabilities, this constitutes a probabilistic correlation. This correlation is consistent with the following hypothesis: the leading man and leading lady share meals, and on 500 of their 1000 nights on the road, their meals were tainted. Each has a .7 probability of being sick when they eat tainted food, and a .3 probability when they do not. Assuming independence conditional upon eating tainted food and not, we would expect them both to be sick on 245 (i.e., $500(.7)(.7)$) of the nights when they eat tainted food, and on 45 (i.e., $500(.3)(.3)$) of the nights when they do not. Thus even if the hypothesis postulated in explanation of the correlation is true, and even if the leading man and leading lady are both sick on a given night, there is still a reasonable chance (about .16) that they did not eat tainted food on that night.

(iii) It is not essential to the example that the tainted food have the same effect in both the leading man and the leading lady. Suppose that instead of suffering from gastric distress, the leading man had passed out on the days in question. The correlation between the leading lady's stomach illness and the leading man's passing out would still warrant a common cause explanation (assuming that it is a genuine probabilistic correlation and not merely a statistical correlation resulting from sampling error). It is thus important to distinguish between *correlation* and *resemblance*. There is an intuitive sense in which the two effects—the two actors' being sick—resemble one another, but this is not what the common cause inference is based upon. Two resembling, but un-

correlated events do not warrant a common cause explanation; two correlated, dissimilar events do. Indeed, Thagard (1988, 162–168) argues that inferring causal connections on the basis of resemblance, rather than correlation, is a symptom of pseudoscience.

(iv) Because of the frequencies cited in the example, the probabilistic inequality required by CCP is rendered meaningful. This is not a deep point about the metaphysics and epistemology of probability. I mean only that the frequencies provide some grounds for assigning the probability values in question, and give us some conceptual grip on what those probabilities are. This is especially true for the joint probability $P(AB)$. We are able to assign a value to this joint probability, independently of the probabilities of A and B taken individually, because it is possible to pair off instances of event types A and B and thus take frequencies of joint occurrences. That is, because we can pair individual episodes of illness in the leading man with episodes of illness in the leading lady (we pair them if they occur on the same day) we can use frequencies to estimate the probability of the event type AB .

By contrast, suppose that the director of a rival theater builds two robots to take on the leading roles in her productions (perhaps hoping to avoid the need for cancellations due to stomach illness). On the first night, the robots break down and cannot perform. Abandoning her idea, the director destroys the robots and never builds another. Does the fact that the robots broke down on the same night constitute a correlation calling for a common cause explanation? That depends upon what the probabilities were. Here, we cannot reliably infer from frequency data the probability that each robot would break down on any given night. We may be able to come up with an a priori estimate, say on the basis of engineering considerations. This will not enable us to estimate a probability for their both breaking down, however, unless we build assumptions about dependence or independence into our estimate (and thus beg the question).

Philosophical errors have resulted from failure to heed these observations. Consider, for example, Salmon's reconstruction of Perrin's argument for the reality of molecules (Salmon 1984, 213–227). Perrin recounts numerous experiments to determine Avogadro's number, the number of molecules in one mole of a substance. The experiments involved very different phenomena, such as Brownian motion, X-ray diffraction, and black body radiation. These experiments all yielded results on the order of 6×10^{23} . According to Salmon, this agreement between methods of computing Avogadro's number constitutes a striking correlation, to be explained in terms of a common cause: the actual presence of a fixed number of molecules in a mole. This, however, is *not* a correlation in the probabilistic sense. In particular, the description

just given is consistent with each experiment being performed just once, in which case we have no basis for inferring any probabilistic correlations whatsoever. Consider for example, just two types of experiment, one involving alpha decay, the other involving black body radiation. Let A represent an experiment of the first type's yielding a result between 5 and 7×10^{23} , and B represent an experiment of the second type's yielding a result in the same range. If each type of experiment is performed just once, each yielding results in the specified range, that is entirely consistent with $P(A) = P(B) = P(AB) = 1$, in which case there is no correlation between A and B . If each type of experiment is performed several times, the frequencies may give us an estimate of $P(A)$ and $P(B)$. In order to get an estimate of $P(AB)$, we need to pair off particular experiments of each type. For instance, suppose we performed one experiment of each type each day for a number of days. Then $P(AB)$ could be estimated by the proportion of days on which *both* experiments yielded results in the specified range. But if we were to conduct this sort of experiment, it would be very *surprising* if $P(AB) > P(A)P(B)$; such a correlation would suggest that Avogadro's number was varying from day to day. Therefore, the resemblance among experimental results reported by Perrin do not furnish us with the sort of probabilistic correlation that permits us to apply the common cause principle. This does not mean that Perrin did not provide impressive evidence for the existence of a constant number of molecules in one mole of a substance—heterogeneity of evidence in support of a hypothesis has been recognized as a virtue by writers on methodology since Whewell—only that his inference cannot be reconstructed using Reichenbach's common cause principle.

3. Genealogical Inference. In this section, I will argue that a certain type of inference in historical linguistics assimilates readily to CCP as illustrated in the previous section. First, I will distinguish between three different types of question that might be asked about the genealogy of a species or language; then I will argue that there is an interesting difference between the ways in which evolutionary biology and historical linguistics attempt to answer one of these types of question.

3.1. Three Questions. Let A , B , and C be languages or biological species. Then we might ask any of the following questions:

Q1: Do A and B share a common ancestor at all?

Q2: Given that A , B , and C share some common ancestor, do A and B share an ancestor that is not shared with C ?

Q3: Given that *A* and *B* share a common ancestor, what does the most recent common ancestor 'look' like?

The recent literature in the philosophy of biology has focused on questions of type Q2, and rightly so, since this type of question has been the focus of the most intense debate within the field of systematics itself (see Hull 1988 for an engaging presentation of this debate). In this context, CCP is subject to well-known limitations. In particular, this principle ignores the crucial distinction between correlations among *derived* features of organisms, called *apomorphies*, and correlations among *ancestral* features, called *plesiomorphies*. According to the cladistic school, it is only correlations among apomorphies that have probative value in addressing questions of type Q2. Note that in order to characterize the traits of organisms *A*, *B*, and *C*, as apomorphies or plesiomorphies, we must know which of these traits were possessed by the common ancestor of *A*, *B*, and *C*; that is, we must have an answer to a question of type Q3. As Sober (1988, §6.5) argues, however, even very rudimentary and unreliable methods of answering Q3 suffice for providing reliable answers to Q2.

Note that questions of both types Q2 and Q3 presuppose an answer to a question of type Q1. In evolutionary biology, however, questions of type Q1 are no longer live questions: all life forms are believed to share some common ancestor. One does see Q1 broached, but only in responses to spurious creationist challenges, and in the introductory chapters of biology textbooks (chapters with titles like "Evidence for Evolution").

In historical linguistics, all three types of question are actively addressed. For example, English belongs to the *Indo-European* family of languages, which includes the Germanic, Romance (Italic), Celtic, Slavic, and Indo-Iranian language groups, among others. The presumed common ancestor of the Indo-European Languages is referred to as *Proto-Indo-European*. The bulk of research in historical linguistics is concerned with the reconstruction of ancestral languages, that is, with questions of type Q3. Of course, question Q3 presupposes an affirmative answer to Q1, but it is possible that this presupposition may be vindicated or undermined by the success or failure of the reconstruction project, rather than established independently. Indeed, many linguists would claim that it is not possible to establish the common origin of two languages without making some headway in attempting to reconstruct their common ancestor. This is an issue that I will finesse in subsequent sections: where linguists talk of languages supplying raw material for reconstruction, which in turn provides evidence for common ancestry, I will talk directly of languages providing evidence of

common ancestry. So long as the features of languages that make reconstruction possible are not distorted in the process, this simplification will have no ill effects. Some linguists (e.g., Ringe 1992) take a similar approach.

Historical linguists also address questions of type Q2. For example, among the Indo-European languages, there are standardly recognized subfamilies, such as Germanic and Romance, but many open questions about the relationships among these subfamilies and among the languages within them. In addressing these questions, linguists use similar techniques to those employed by biological systematists; in particular, shared innovations among languages have a particularly strong evidential bearing. Again, in order to distinguish shared innovations from shared retentions, one must first be in possession of answers to questions of type Q3. At the same time, answers to questions like Q2 are often useful in carrying out reconstructions. For example, to reconstruct the common ancestor of English and Hindi, i.e., Proto-Indo-European, it is helpful to first reconstruct the common ancestor of English and German, *Proto-Germanic*, and the common ancestor of Hindi and Persian, *Proto-Indo-Iranian*, and then attempt a reconstruction of the common ancestor of these reconstructed proto-languages.

Finally, historical linguists actively address issues of type Q1. For example, some linguists have claimed that the Indo-European languages belong to a larger family that includes, among others, the *Uralic* languages such as Finnish and Hungarian. Such claims are highly controversial. Thus linguists are actively concerned with questions such as whether English and Finnish are related at all. Note that in addressing this sort of question, it is impossible to avail oneself of knowledge of shared innovations and shared retentions without begging the question; many of the techniques that allow one to answer questions of type Q2 are not available. The limitations of CCP as a tool for addressing questions of type Q2 become limitations for *all* methods for addressing questions of type Q1. Thus it is an open question whether CCP might be suitable for addressing these questions.

In this paper I will be concerned only with questions of type Q1. These questions have relatively little intrinsic interest in the biological case, where monogenesis is taken for granted; not so in the case of historical linguistics. While questions of type Q2 have rightly dominated the literature in the philosophy of biology, this has obscured some of the important distinctions between the methods of systematic biology and historical linguistics by focusing attention on the area where the methods of the two fields are most closely analogous. By focusing attention on Q1 I hope to illustrate some interesting differ-

ences between the two fields that may serve as a springboard for further research into the methodology of historical linguistics.

3.2. Resemblance and Correlation. Human hands and bat wings have very similar bone structure. Although the corresponding bones have different sizes relative to one another, both structures contain the same number of bones, in similar orientations to one another. Such similarities are commonly taken to be indications of common descent. Analogously, we note the English word ‘feather’ and the German word ‘Feder’ are very similar in structure. Both have two syllables, and contain the same initial and final consonant sounds. (It is the pronunciation, rather than the spelling, of words that is of primary interest to linguists. Where spelling is a reasonable guide to pronunciation, I will not worry about this distinction here.) Correspondences of this sort (including correspondences in morphology and syntax as well as phonology) provide the primary evidence upon which inferences about linguistic genealogy are made.

While these inferences look, *prima facie*, like common cause inferences, to reconstruct them as such would be premature. While both cases clearly involve resemblances in some intuitive sense, we are at risk of repeating Salmon’s error: we have not shown that these resemblances constitute genuine probabilistic correlations. Consider first the case of the structure of the human hand and the bat’s wing. Let A represent the human hand having the bone structure that it in fact has, and B the bat’s wing having a similar structure. Can we get a purchase on $P(A)$, $P(B)$, and $P(AB)$? Most humans have hands with the relevant structure, but a few—due to accidents, congenital conditions, or what have you—do not; likewise for bats. So one possibility would be to use the frequency of the structure within the human and bat populations as an estimate of $P(A)$ and $P(B)$ respectively. If there were a natural way of pairing individual humans with individual bats, we could use frequency within the population of bat-human pairs to estimate $P(AB)$. But of course there is no such pairing: our inference of common descent for bats and humans is not based on these kinds of frequencies.

A second possibility would be to reflect that from an engineering perspective, given the function of bat wings and human hands, there is no necessity in either having the structure it does—indeed, many possible structures can be imagined. Therefore it may seem reasonable to assign some fairly low probability to $P(A)$ and $P(B)$. These probabilities may be thought of intuitively as frequencies over possible worlds. I do think that the plethora of conceivable structures for wings and hands contributes to our being impressed with this similarity, but

it does not help us to construe this resemblance as a probabilistic correlation. For how are we to gain a purchase on $P(AB)$ to determine if it is greater than $P(A)P(B)$? It is no good to reason that the probability $P(AB)$ cannot be all that low, after all, since A and B did both occur: some low probability outcome had to occur. Moreover, such a priori reflections about alternate possibilities raise difficult questions about whether some non-actual creature could have a very different wing structure and still be a *bat*. The problem here is that bats and humans have evolved only once, with only the types of wings and hands that they in fact have, so we do not have any nontrivial frequencies upon which to erect probabilities.

Here is a further possibility: Bats and humans have enough basic morphology in common that it is possible to pair off various of their characteristics—they both have eyes, ears, nose, spine, and so on. Although bats' wings and human hands are quite different, they bear sufficiently similar relationships to other structures that they can be paired off without appealing to their similar internal structures. Let's suppose that we identify 100 human body parts and their bat analogs. Of these, four of the bat's parts have the structure in question, and four of the human's parts do (assuming we specify the structure loosely enough to include feet as well as hands and wings). However, it is the *corresponding* four parts in both bats and humans that have the structure in question. That is, $P(A) = P(B) = P(AB) = .04$. The probabilities here are based not on frequencies over individuals within taxa or over possible worlds, but over body parts. This approach generates the needed probabilistic correlation, but there are at least two difficulties. The first is that it seems to misdescribe what is so striking about the example: the correlation in question is that of all the body parts that might exhibit the structure in question, it is the parts at the ends of limbs in both humans and bats that do. It is surely not this correlation that catches our eye. More importantly, this correlation would exist regardless of the structure of extremities in bats, so long as there was some structure such that all and only the extremities of bats had that structure. (Recall the variant on Reichenbach's example where the leading man passes out.) Yet what is striking about the example is the *similarity* of structure in human hands and bat wings.

Three observations are in order. First, although the inference that humans and bats share a common ancestry cannot be reconstructed as the inference from a detected probabilistic correlation to a common cause, it in no way follows that the inference is invalid. Second, it may be possible to subsume this inference, together with common cause inferences such as that described in the previous section, under some more general principle. Sober (1988) argues that *likelihood* may play

such a role. Third, systematists do sometimes make inferences based upon frequencies not unlike those described in the previous paragraph; but they do so in addressing questions of type Q2. Here is a slightly modified schema: Suppose species *A* and *B* are both known to have descended from some common species ancestor *D*. The question under investigation is whether they share a common ancestor *C* more recent than *D*. The characteristics of individuals within the two species may be paired off as above. Moreover, for each species, it is determined for each character whether that character is inherited directly from *D*, or whether it has been modified through natural selection. If there is a correlation among derived characters that is taken as evidence for a more recent common ancestor. By taking a plethora of characters and reducing them to two basic types, derived and ancestral, it is possible to point to frequencies that may bespeak true correlations. Note, however, that it is only correlations among derived traits that bespeak recent common ancestry, so even here the unrestricted CCP does not apply.

In the case of historical linguistics, by contrast, it is much easier to point to frequencies that could, in principle, ground claims of probabilistic correlation. Consider again the case of ‘feather’ and ‘Feder’. Among words of English and German, some begin with ‘f’, while others do not. It is possible, at least in principle, to determine the relative frequency with which words of English and German start with this consonant sound. These frequencies provide ground for talking of the probability that a word of English will start with ‘f’, and likewise for German. (Note that word-initial ‘v’ in German is usually pronounced like the English ‘f’, so words beginning with this consonant would also be included.) Moreover, there is a natural way of pairing up words of English and German—by synonymy or inter-translatability, or more realistically, similarity of meaning or semantic affinity. Such a system of pairing allows us to determine a relative frequency with which corresponding words in the two languages both begin with ‘f’. (It is relatively common: ‘father’/‘Vater’, ‘fire’/‘Feuer’, ‘four’/‘vier’, etc.) Thus we have a solid grounding for talk about the probability of the joint occurrence of word-initial ‘f’ among English-German synonyms. If this probability is higher than the product of the probabilities of word-initial ‘f’ in the two languages taken individually, then we have a genuine probabilistic correlation that can serve as a premise for a common cause inference. As we shall see in the next section, linguists do in fact rely on these kinds of frequencies in drawing inferences about the histories of languages.

Let us spell out the nature of the explanation of the correlation between word-initial ‘f’ in English and German a little more carefully.

The excess of English-German word pairs in which both begin with ‘f’, compared to what would be expected by chance, is explained by postulating a *type* of common cause: a word (presumably also beginning with ‘f’) from which given English and German words have both descended. Some synonymous (or semantically related) word pairs of English and German have such a common origin, some do not (either because they do not share an ancestor, or they do not share an ancestor of the appropriate sort). Moreover, there is a broader framework—an ancestral language from which English and German have both descended—within which this common cause type is sometimes instantiated. That is, it is because English and German descend from a common ancestral language that some synonymous English-German word pairs are derived from a common word beginning with ‘f’.

Note that the correlated sounds that indicate common ancestry need not be the same. The English consonant ‘f’ is strongly correlated with the consonant ‘p’ in Romance languages: witness English ‘father’, ‘fish’, ‘foot’ and Spanish ‘padre’, ‘pescado’/‘pez’, ‘pie’. A case may still be made that the sounds ‘f’ and ‘p’ *resemble* one another: both sounds result from a stoppage of air flow at the front of the mouth, by the upper and lower lips in the case of ‘p’, and by the upper teeth and lower lip in the case of ‘f’. But even this level of resemblance is not necessary. For example, there is a recognized correlation between the Latin ‘du-’ and Armenian ‘erk-’, as in ‘duo’/‘erku’ (‘two’). For an even more dramatic example, consider Vietnamese, which is a tone language. In Vietnamese, whether a syllable is pronounced with rising pitch, descending pitch, and so on affects the meaning of that syllable. Haudricourt (1953, 1954) established that there is a correlation between the tones of Vietnamese, and the consonants of (other) Mon-Khmer languages, thus providing evidence for the inclusion of Vietnamese in that language family. In this case the correlations are among sounds that are not even in the same phonetic category.

By the same token, similarities per se are not usually taken to constitute evidence of common ancestry. This is a point that has been unappreciated in the philosophical literature. For example, Sober (1993, 42), in talking about the similarity of the French, Italian and Spanish names for numbers, writes: “the fact that these languages assign similar names to numbers is striking evidence [that they are related to each other].” Surprisingly, most linguists do not put much stock in this sort of evidence. A standard textbook example illustrating the dangers of relying on similarity involves the Latin and Greek words ‘deus’ and ‘theos’ (both meaning ‘god’). These words are strikingly similar in sound and meaning, but do not have a common etymology. This observation by itself is not especially telling, since any mode of inference

may yield incorrect conclusions on particular occasions. What is important about this example is that linguists have a systematic reason for rejecting shared etymology in this case: the Latin 'd' and Greek 'th' are not correlated sounds. (By contrast, Latin 'd' is strongly correlated with Greek 'd', and Latin 'f' with Greek 'th'.) Indeed, those linguists who have attempted to base genealogical inferences upon phonetic similarities of words in different languages, most notably Joseph Greenberg, have met with strong resistance. In a recent text, Fox (1995) writes:

. . . instead of presenting correspondence sets for phonological correspondences, . . . and using these correspondences as a basis for linguistic groupings, as discussed above, Greenberg is content to look for some evidence of phonetic similarity between forms of similar meaning in some of the languages compared

It will be evident that [Greenberg's] method runs counter to the principles that we have already examined in earlier chapters of this book, and as a result lays itself open to the various erroneous conclusions that those principles were designed to avoid. (238–240.)

In the case of historical linguistics, then, it is true correlations rather than similarities that underwrite inferences of common ancestry. Whether linguists are *correct* in their rejection of inferences based upon similarity is another matter, alas one whose examination must await another occasion.

It was noted in the previous section that the inference to a general type of common cause from a probabilistic correlation is often more reliable than an inference to a particular common cause from one instance of a correlation: this is certainly true of inferences in historical linguistics. For example, there is a correlation between word initial 'b' in English and in German ('blood'-'Blut', 'bite'-'bissen', and so on) and this is due to *some* English and German words beginning with 'b' evolving from 'b'-initial words in a parent language. It does not follow that whenever an English word and its German counterpart both begin with 'b' that they share a common etymology: 'belly' and 'Bauch' do not, for example. Historical linguists, who are often concerned not only with the discovery of relationships between languages, but also in tracing the etymology of particular words and reconstructing ancestral languages, sometimes fail to recognize that a system of correspondences can serve as evidence of linguistic relatedness even if it includes some false cognates. (This failure is decried by Pinker (1994, 255) for example).

Inferences about the common origins of languages based upon sound correspondences are closely analogous to the common cause inference illustrated by Reichenbach's classic example. In order to de-

tect a probabilistic correlation of the sort that serves as a premise for a common cause inference, it is typically necessary to have a type (or many types) that are instantiated by many different individuals in two populations, where individuals from the two populations may be readily paired with one another. In the case of languages, the types are sounds, the individuals are words, and the words of different languages are paired according to meaning. Biological species do not have this type of structure. If the populations consist of individual organisms, there is no natural way of pairing organisms across species; if the populations consist of body parts, the relevant types of structures are not instantiated by sufficiently many individuals. Thus, in answering questions of type Q1, genealogical inferences in historical linguistics conform to CCP in a manner that seemingly analogous inferences in evolutionary biology do not.

4. Ringe's Proposal. We commonly make inferences that conform to the common cause principle without explicitly estimating the relevant probabilities. Even if the actors in the traveling theater troupe did not maintain careful records of their health history, they no doubt would still have had a qualitative sense that their illnesses were correlated, and have been able to infer the existence of a common cause.

Similarly, in historical linguistics, inferences about the common ancestry of languages are often made without any explicit attempt to estimate probabilities. No one familiar with English and German, for example, can fail to notice the strong correlations between the sounds of the two languages. In a recent textbook, Roger Lass offers the following summary of the method of historical linguistics:

1. If two (or more) languages show regular correspondence . . . ;
2. and if these correspondences cannot be due to chance because of their pervasiveness and apparent systematicity;
3. and if historical factors and/or the systematicity of the similarities rules out diffusion;
4. then the correspondences are due to common origin; or one language is the descendant of the other.
5. If (on independent historical grounds) direct lineal descent is not in question, the correspondences are due to common origin. (Lass 1997, 124)

The possibilities described in 3, the second half of 4, and the antecedent of 5 correspond to the caveat attached to CCP which allows us to infer a common cause of *A* and *B* only after we have ruled out the possibility of direct causal influence from *A* to *B*. Of this, more in Section 5 below. For now, note that 2 specifically requires that a given correspondence

(such as word-initial 'f' in English and German) occur more often than would be expected by chance, exactly as required CCP. Yet in what follows, Lass does not once talk about estimating the frequency of chance occurrences, or about conducting statistical significance tests. Judgments about whether 2 has been satisfied are to be entirely qualitative. This qualitative approach has certainly enjoyed its share of success: for example members of the Indo-European language family are almost universally regarded as having descended from a common source on the basis of primarily qualitative evidence.

In recent decades, however, there has been a trend toward the postulation of ever-broader linguistic families. The basic idea that Indo-European may be part of an even larger language family (usually referred to as *Nostratic*) has been around since the turn of the century, but has received its most ambitious formulation in the work of the Russian linguist Vladislav Illich-Svitych (1971–1984). Illich-Svitych proposed a Nostratic family comprising Indo-European along with Uralic, Afro-Asiatic (Arabic, Hebrew, and various Northern African languages), Altaic (Turkish and Central Asian languages; perhaps also Japanese and Korean), Dravidian (Southern Indian languages such as Tamil and Telugu), and Kartvelian (languages of the Caucasus Mountain region). Joseph Greenberg (1987) has proposed that most of the diverse native languages of North and South America belong to one linguistic family: *Amerind* (Eskimo and Aleut languages comprise a distinct family, as do the *Na-Dené* languages of Alaska, Northwest Canada, and parts of the American Southwest).

These new super-families of languages have met tremendous resistance from historical linguists. The debate has often been acerbic, fueled in part by the perception that the attention these proposals have received in the more popular media (such as *The New York Times*, *Scientific American*, and *Nova*) is disproportionate to the amount of acceptance they have received among experts in historical linguistics.

As we have already seen, one objection to Greenberg's proposal is that it is based upon similarities, as opposed to correlations. Putting this issue aside, a principle difficulty is that the evidence adduced in support of these families consists of sound correspondences found among hundreds of languages. If these correspondences really bespeak common origins, they must exceed in number the correspondences that would be expected on the basis of chance alone. When dealing with such large numbers of languages, scepticism about our ability to make qualitative judgments about probabilistic correlations among them is hardly unreasonable. It is a well-known psychological phenomenon that humans tend to recognize similarities and ignore differences, to the effect that they lose sight of the possibility that the similarities might

be coincidental. It seems, then, that there is a need to get at least a loose quantitative grip on the correspondences between languages to determine whether they constitute genuine correlations. Donald Ringe (1995), an influential critic of the new super-families, writes:

It is not always clear whether the similarities observed between the lexica of different languages could easily be the result of random chance or must reflect some historical relationship. Particularly difficult are cases in which the relationship posited is remote at best; such cases must be evaluated by comparison with mathematically valid models which realistically simulate chance resemblances between languages. (1995, 55)

Ringe (1992) offers a proposal toward that end (although, unfortunately, not one that he adheres to in Ringe 1995; see Hitchcock and Manaster Ramer, in preparation, for more details).

Consider any two languages to be compared, say English and German. We compile a canonical one hundred word list for each language (called a 'Swadesh' list, after a pioneer of the method). The words are chosen so that they are likely to reflect the ancestry of a language. (Obviously, if we are searching for signs of a possible common ancestor millennia ago, comparing words for 'telephone' will not be of much use.) Examples include words for numbers, body parts, relatively obvious categories of plant and animal, colors, logical operators, and so on. Frequencies of, say, initial consonants within this list serve as estimates of probabilities. Thus, for example, five of the English words in Ringe's hundred-word list begin with 't'—'two', 'tree', 'tail', 'tongue', and 'tooth'; so the probability that an English word commences with 't' is estimated to be .05. Likewise, three of the hundred German words begin with 'z' (pronounced 'ts')—'zwei' (two), 'Zunge' (tongue), 'Zahn' (tooth); so the probability of initial 'z' in German is estimated to be .03. Finally, three words begin with 't' in English, *and* with 'z' in German, for a probability of .03. This is much greater than the product of the two probabilities, .0015, so we have a probabilistic correlation between word initial 't' in English and 'z' in German. Since these probabilities are merely estimates based upon frequencies in one hundred word lists, Ringe requires that any such correlations be significant at the .01 level; e.g., if the probability for an English word to start with 't' and its German counterpart to start with 'z' were .0015, there would have to be less than a .01 probability of finding 3 such pairs among one hundred word-pairs (as indeed there is) for the correlation to count. The number of consonant-pairs exhibiting significant correlations are then an indication of the relatedness of languages. For English and German, it turns out that there are seventeen such pairs,

for English and Latin, seven. Likewise, a similar comparison can be made for other positions, such as second consonant or final consonant. To test the legitimacy of a proposed super-family such as Nostratic or Amerind, pairs of languages (or reconstructed proto-languages) comprising that proposed family must be independently tested. Ringe's proposed test is clearly aimed at detecting the sorts of probabilistic correlations that call for common cause explanations.

There are a number of difficulties with this method. Some of these involve garden variety concerns over the probability of false negatives and false positives. On the one hand, Ringe's requirement that correlations be significant at the .01 level seems stringent—correlations between sounds in distantly related languages might well pass through the sieve. On the other hand, in the case of both English and German, there were seventeen initial consonants (sixteen consonants plus the vowels, which are lumped together as the null-consonant) appearing in the one hundred word list, making 289 possible consonant pairings. Assuming this number to be typical for pairs of languages, we should not be surprised to find a few accidental statistical correlations among initial consonants of any two languages. This suggests that Ringe's method may not be sensitive enough to adjudicate borderline cases. Indeed, Baxter and Manaster Ramer (1996) apply Ringe's test to pairs drawn from the following languages: English, Dutch, Welsh, Albanian (Indo-European); Hebrew, Hausa (Afro-Asiatic); Turkish (Altaic). What they find is that Ringe's test does not discriminate between pairs of languages that are currently accepted as distantly related members of the same recognized language group, and pairs drawn from different families that are unrelated, or at best more distantly related.

A further problem is that Ringe's test might yield a false negative, not merely because of insufficient statistical power, but because the related languages do not exhibit the type of correlation being tested for. Languages evolve in systematic ways, and may do so in ways that erase correlations between consonants. For example, some Australian Aboriginal languages, such as Olgolo, have lost *all* of their initial consonants. Nonetheless, traces of the consonant remain in the pronunciation of initial vowels or of consonants occurring later in the word (much as the nasalized vowel in the French 'non' is a vestige of the final 'n' that was once pronounced). (See Dixon 1980 for a summary of some of these findings.) These vestiges allow linguists to determine relationships between Olgolo and other Aboriginal languages. Any test that sought only correlations between word-initial consonants, however, would find none. Recall also the case of Vietnamese, whose tones are correlated with the consonants of other Mon-Khmer languages. Presumably the consonants of some ancestral language have system-

atically evolved into the tones of modern Vietnamese: if this evolution is complete, even the most sensitive statistical test would discover no correlation between the consonants of Vietnamese and the other Mon-Khmer languages.

None of these difficulties shed doubt on the legitimacy of common cause inferences in historical linguistics. We have already noted that CCP is formulated in terms of probabilities, not observed statistical frequencies. Garden variety problems concerning the inference from frequency data to underlying probabilities do not undermine CCP, although they may cause serious problems for any attempted application of CCP. Moreover, CCP does not guarantee that a common cause will give rise to a correlation of any particular sort—it tells us nothing about what sorts of correlations to look for—it tells us only that if we do find a correlation, it is reasonable to postulate a common cause.

5. Common Cause Caveats. Philosophers discussing CCP, without attention to its application in historical linguistics, have uncovered a number of caveats pertaining to the making of common cause inferences. In this section I will mention two, and show how these caveats correspond to textbook caveats issued to historical linguists.

First, recall that if *A* and *B* are correlated, we are only permitted to infer a common cause if we have ruled out the possibility of a direct causal connection between *A* and *B*. In the case of languages, there are two possible ways in which languages can have a direct causal effect upon one another. The most obvious possibility is direct descent. We have substantial knowledge of Latin, Ancient Greek, and Sanskrit, due in part to extensive written records. The sound systems of these languages are strongly correlated with the sound systems of French, Modern Greek, and Hindi, respectively. In this case, the explanation of the correlations would not be descent from a common ancestor, but direct descent of French from Latin, Modern from Ancient Greek, and Hindi from Sanskrit.

A second possibility for the direct causal influence of one language on another is *borrowing* (or diffusion). The English word ‘karaoke’ is very similar in pronunciation to a Japanese word with the same meaning. But this is no indication of a common origin of the English and Japanese languages: rather the English word is taken directly from the Japanese. Moreover, while the English word ‘karaoke’ and the Spanish word ‘karaoke’ have a common cause, that common cause does not lie in the common origin of the two languages, but only in the Japanese language from which both English and Spanish borrowed the word. Isolated loan words will not give rise to systematic correlations between the sounds of two languages, but in some cases borrowing from one

language into another is extensive. English has borrowed extensively from Latin and French, for example. We noted above a correlation between English 'f' and Romance 'p', illustrated by 'father'-'padre', 'fish'-'pez' and 'foot'-'pie'. But we might also find a 'p'/'p' correlation, illustrated by English words such as 'paternal', 'piscine', and 'pedestrian'. If taken at face value, this correlation might lead us to postulate a more recent common ancestry for English and the Romance languages than actually exists. Vietnamese has borrowed so many words from Chinese that the former's relationship with the Mon-Khmer languages (which do not include Chinese) had been obscured until relatively recently.

At this point, we must be careful not to take the analogy between linguistic relationships and family trees too far. I do not mean to suggest that descent with modification is the only true source of linguistic relatedness, and that borrowing merely serves to obscure such relationships. Rather, borrowing is a distinct mechanism wherein one language can be related to another. The moral of the present caveat is that the two mechanisms must be clearly distinguished if genealogical relationships between languages are to be understood. Particularly problematic in this regard are *creole* languages, such as those spoken in Jamaica, Haiti, and Sierra Leone. Creole languages are true hybrids, and thus their lexica are not obviously classifiable as having been borrowed rather than inherited or vice versa. (Similar problems involving migration and hybridization can arise in evolutionary biology as well, although here the barriers to such 'horizontal' transmission of information are stronger.)

The second general caveat concerning common cause inferences is that we must distinguish between two different kinds of common cause. This point is made clearly by Salmon (1984, Ch. 6), who uses the terms 'conjunctive' and 'interactive' to distinguish them. Salmon claims that these two types of common cause are characterized by different probability relations. This claim has been contentious, but I will not address it here. Interactive common causes involve direct physical interactions between systems; conjunctive common causes involve systems that are exposed to similar environments. Reichenbach's example of the traveling theater troupe is of the second sort: the leading man and leading lady do not become ill because they directly interact with one another, but because they are exposed to the same environmental hazard, tainted food. As an example of the first sort of common cause, Salmon asks us to imagine a billiard table containing only the cue ball and the eight ball in a particular configuration. This configuration is such that if the cue ball strikes the eight ball in such a way that the latter will go into the corner pocket, then the cue ball will also go into a corner

pocket. If we imagine an imperfect pool player making many attempts to sink the eight ball in this configuration, the eight ball will sometimes go in, sometimes not; however, whenever the eight ball goes in, the cue ball does as well, so there is a strong correlation between the sinking of the two balls, explained in terms of the occasional direct physical interaction between them. Note that the observations that were made upon Reichenbach's example of the traveling theater troupe in Section 2 carry over readily to the billiard ball example, so those observations are not specific to one type of common cause or the other.

This dichotomy does not map perfectly onto the domain of historical inference in linguistics, but it does correspond at least roughly to a linguistically significant distinction. In the case of historical linguistics, we are interested in whether the genealogies of two languages actually overlap. This suggests that we are looking for *interactive* common causes. We must be wary, then, of the possibility that correlations between languages are brought about by conjunctive common causes; that is, we must be wary of correspondences between languages that arise when the sounds of words are shaped by similar types of processes. A word similar in pronunciation and meaning to the English 'mama' appears in Mandarin Chinese. A plausible explanation is that in both languages, sounds made by infants have become associated with the family members (such as mothers) that are in their presence when the sounds are made. Since 'mama' is among the simplest sounds for infants to make, this is the sound that has become associated with mothers in both languages. Such words are called *nursery words*. If this explanation is correct, the similarity between the Mandarin and English words would be due not to the common origin of the languages in question, but rather to the similar circumstances in which the words enter the language.

Much the same may be true of onomatopoeic words. Several languages contain words similar in both sound and meaning to the English word 'cock-a-doodle-do' (such as the Russian 'kukuriku'); that need not be due to the common origins of the languages in question, but is likely because those languages also have words that resemble the sounds they denote. For another example, consider the English word 'babble', meaning to speak incomprehensibly; Aztec contains the verb *popolo-ca*, 'to speak a foreign language', but this similarity may well be due to onomatopoeia.

While onomatopoeia and nursery words are standard fare in linguistics textbooks, they pose a much greater risk to the enterprise of constructing etymologies of particular words than they do to the enterprise of uncovering historical relationships between languages. In general, the overall impact on the correlation of sounds between lan-

guages that can be expected due to onomatopoeia and nursery words is minuscule, and unlikely to obscure the true relationship (or lack thereof) between languages.

6. Directions for Future Research. I do not pretend to have done more than scratch the surface: to characterize one sort of inference in historical linguistics as conforming to CCP is to provide only the loosest sort of characterization of one part of the field; many problems of philosophical interest remain. Here are a few:

(i) I have taken the legitimacy of common cause inferences in the linguistic domain more or less for granted. As Sober (1988) has forcefully argued, however, individual methods of non-deductive inference cannot be globally justified, but only justified relative to domain-specific empirical assumptions. In this case, common cause inferences must be justified by assumptions about the processes of language change. The standard assumption underlying most work in historical linguistics is that sound change takes place in a lawlike manner within individual languages; like sounds evolve into like sounds. What does this assumption, suitably formulated, actually entail about the success of common cause inference? Is the common cause principle incompatible with competing accounts of language change, such as the wave model, wherein linguistic innovations have an ‘epicenter’ within a particular language but then disseminate into other languages, largely on the basis of geographical proximity?

(ii) As mentioned above, historical linguists are typically concerned primarily with the *reconstruction* of ancestral languages, that is, with answering questions of type Q3. Is philosophy of science able to shed any light on the methods employed in reconstruction? Since reconstruction is based upon sound correspondence in much the way described above, can the methods of reconstruction be characterized in terms of CCP? One obvious problem is that CCP, as formulated in Section 2 above, may allow us to infer that there *is* a common cause, but it does not allow us to infer what form that cause takes. The latter would seem to be essential for reconstruction—unless we conceive of the reconstructed forms as abstracta, as some linguists think we ought.

(iii) Are the arguments against the cogency of inferences based upon resemblance themselves cogent? This is a central issue in the controversy over Greenberg’s proposed Amerind language family. Arguments offered by linguists often appeal to broadly probabilistic considerations, but are rarely made probabilistically rigorous.

(iv) The sorts of mathematical models described in Section 4 were oversimplified—they measured only correlations among initial consonants of synonymous words in two languages. The preponderance of

evidence for genuine historical relationships, however, comprises correspondences among sounds that occur in slightly different positions in words of related but not identical meaning in several languages. Are realistic probabilistic models of such looser correspondences possible? Ringe (1995) purports to make a step in this direction, but his argument is based upon a fundamental misunderstanding. Ringe argues that since the distribution of alleged cognates in the six subfamilies of Nostratic is binomial, the putative cognate sets might well reflect chance resemblances among languages. But so long as the six subfamilies evolved independently *after* having diverged from a common ancestor, so that vocabulary loss was independent and identically distributed, one would equally expect a binomial distribution of cognates (see Hitchcock and Manaster Ramer, in preparation, for more details). In this case, the lack of a concrete and plausible model for the evolution of several language families led one prominent linguist astray.

(v) The mathematical models proposed by linguists are typically formulated in the idiom of classical statistics. Given that strong evidence for linguistic relationships can be supplied by historical, archeological, cultural, textual and even genetic sources, would it not be more appropriate to examine linguistic inferences from a Bayesian perspective, where prior probabilities can reflect the evidence conferred upon certain hypotheses from non-linguistic sources?

While some of these issues have been addressed by historical linguists themselves, discussions often lack the philosophical sensitivity that such issues deserve. It is my hope that this state of affairs will change.

REFERENCES

- Arntzenius, Frank (1993), "The Common Cause Principle", in David Hull, Mickey Forbes, and Kathleen Okruhlik (eds.), *PSA 1992, vol. 2*. East Lansing, MI: Philosophy of Science Association, 227–237.
- Baxter, William and Alexis Manaster Ramer (1996), "Review of *On Calculating the Factor of Chance in Language Comparison*, by Donald A. Ringe, Jr.", *Diachronica*, 13: 371–389.
- Dixon, Robert M. W. (1980), *The Languages of Australia*. Cambridge: Cambridge University Press.
- Fox, Anthony (1995), *Linguistic Reconstruction: An Introduction to Theory and Method*. Oxford: Oxford University Press.
- Greenberg, Joseph (1987), *Language in the Americas*. Stanford: Stanford University Press.
- Haudricourt, André G. (1953), "La Place du Vietnamiens dans les Langues Afroasiatiques", *Bulletin de la Société de linguistique de Paris* 49 (1): 122–128.
- . (1954), "De l'origine des tons en Vietnamien", *Journal Asiatique* 242: 69–82.
- Hitchcock, Christopher and Alexis Manaster Ramer (in preparation), "Binomial Beware: Ringe on the Nostratic Hypothesis".
- Hull, David (1988), *Science as a Process: An Evolutionary Account of the Social and Conceptual Development of Science*. Chicago: University of Chicago Press.
- Illich-Svytich, Vladislav (1971–1984), *Opyt Sravnenija Nostraticheskix Jazykov*. Moscow: Nauka. (Vols. 1 (1971), 2 (1976), 3 (1984).)

- Lass, Roger (1997), *Historical Linguistics and Language Change*. Cambridge: Cambridge University Press.
- Pinker, Steven (1994), *The Language Instinct*. New York: Morrow.
- Reichenbach, Hans (1956), *The Direction of Time*. Berkeley and Los Angeles: University of California Press.
- Ringe, Donald, Jr. (1992), *On Calculating the Factor of Chance in Language Comparison*. *Transactions of the American Philosophical Society*, vol. 82, no. 1. Philadelphia: American Philosophical Society.
- . (1995), “‘Nostratic’ and the Factor of Chance”, *Diachronica* 12: 55–74.
- Salmon, Wesley C. (1984), *Scientific Explanation and the Causal Structure of the World*. Princeton: Princeton University Press.
- Sober, Elliott (1988), *Reconstructing the Past*. Cambridge, MA: MIT Press.
- . (1993), *Philosophy of Biology*. Boulder: Westview Press.
- Thagard, Paul (1988), *Computational Philosophy of Science*. Cambridge, MA: MIT Press.