The Journal of Mind and Behavior
Summer and Autumn 2007, Volume 28, Numbers 3 and 4
Pages 175–188
ISSN 0271–0137

Why History Matters: Associations and Causal Judgment in Hume and Cognitive Science

Mark Collier

University of Minnesota, Morris

It is commonly thought that Hume endorses the claim that causal cognition can be fully explained in terms of nothing but custom and habit. Associative learning does, of course, play a major role in the cognitive psychology of the *Treatise*. But Hume recognizes that associations cannot provide a complete account of causal thought. If human beings lacked the capacity to reflect on rules for judging causes and effects, then we could not (as we do) distinguish between accidental and genuine regularities, and Hume could not (as he does) carry out his science of human nature. One might reply that what appears to be rule-governed behavior might emerge from associative systems that do not literally employ rules. But this response fails: there is a growing consensus in cognitive science that any adequate account of causal learning must invoke active, controlled cognitive processes.

Keywords: David Hume, associative learning, rules for judging causes and effects

Hume's account of causation is his most important legacy in the history of philosophy. His first major contribution to the philosophy of causation involves the *definitional* question: How should we analyze the concept of causation? Hume worries that we often use causal terms without the foggiest notion of what they mean. When philosophers are pressed to define causation, for example, they typically maintain that "A *causes* B if and only if A *produces* B or *brings about* B or *necessitates* B." But these causal terms are all synonymous, and thus we have been taken around in a circle (1739/1978, p. 157). Hume makes significant progress, therefore, when he offers a deflationary analysis of causation in terms of invariable regularities.

Requests for reprints should be sent to Mark Collier, Ph.D., Department of Philosophy, University of Minnesota, Morris, 600 East Fourth Street, Morris, Minnesota 56267. Email: mcollier@morris.umn.edu

Hume recognizes that it is one thing to analyze what the term causation means, however, and another thing to explain how we discover causes and effects. Hume's second major contribution to the philosophy of causation involves the *epistemological* question: How can we make causal inferences? The definitional and epistemological issues are, of course, closely related: if causes are constituted by constant conjunctions, then we can discover causes by searching for invariable sequences. For example, every time the ignition key is turned, the car starts. Since one event always follows another, we can pronounce them cause and effect. By taking the mystery out of causation, Hume has apparently taken the difficulty out of causal inference.

This is precisely how Hume characterizes our everyday causal inferences. We ordinarily make causal attributions whenever we observe a frequent repetition between events (p. 87). Why do we make causal inferences in this way? Hume's answer is that these events have become associated in our imaginations; in short, we have become conditioned to do so.

> When the mind . . . passes from the idea or impression of one object to the idea or belief of another, it is not determin'd by reason, but by certain principles, which associate together the ideas of these objects, and unite them in the imagination. (p. 92)

Let us call this the "custom–habit" account. It has several important features. First, it does not postulate any explicit reasoning. We do not ordinarily make causal inferences through reflection, thought, or judgment; rather, our everyday causal attributions depend solely upon associative propensities of the imagination. Second, ordinary causal reasoning is automatic and implicit; when we approach a river, for example, we do not consciously deliberate about whether it is dangerous; rather, "custom operates before we have time for reflexion" (p. 104). Finally, this account does not distinguish between human and animal learning; there is no difference between the way in which we ordinarily make causal inferences and the way in which a dog comes to fear punishment from the tone of his master's voice (p. 178). In sum, Hume appears to be a behaviorist; that is, he seems to hold that causal inferences can be fully explained in terms of associations.

It is tempting to think that Hume endorses radical associationism. After all, one can hardly turn a page of the *Treatise* without encountering the laws of association. Hume is quite explicit about the prominent role that associations play in his theory of mind; he writes at one point that "if any thing can entitle the author so glorious a name as that of an inventor, 'tis the use he makes of the principle of the association of ideas, which enters into most of his philosophy" (pp. 661–662). Indeed, associative learning plays a leading role in every one of Hume's psychological explanations in Book One: it accounts for the origin of ideas, the nature of general language and thought (Collier,

2005b), how we make probabilistic inferences (Collier, 2005a), and why we believe in object permanence (Collier, 1999).

What about causal cognition? Does Hume really maintain that causal thought can be fully explained in terms of associations? The traditional interpretation is that his associationism is intended to be "complete" (Pears, 1991, p. 70) and "comprehensive" (Jessop, 1966, pp. 46–47). But the truth is that Hume does not regard his custom–habit account as a finished product. He is sensitive to the fact that it must be supplemented with a rule-based system in order to paint a complete picture of human causal reasoning. There are two main lines of support for this revisionary interpretation. The first involves Hume's explicit remarks about experimental methods and the "rules for judging causes and effects." The second appeals to the actual experiments performed by Hume in his science of human nature.

Hume's Methodological Pronouncements

Hume observes that genuine causal conditions are often surrounded by "superfluous circumstances."

> In almost all kinds of causes there is a complication of circumstances, of which some are essential and others superfluous; some are absolutely requisite to the production of the effect, and others are conjoin'd by accident. Now we may observe, that when these superfluous circumstances are numerous, and remarkable, and frequently conjoin'd with the essential, they have such an influence on the imagination, that even in the absence of the latter they carry us on to the conception of the usual effect . . . . (p. 148)

For example, imagine that we repeatedly observe two potential causes followed by an effect, where one is genuine, and the other is superfluous. Let us borrow a contemporary example in order to flesh this out (Glymour, 1998, p. 40). Suppose that we are searching for the causes of lung cancer. We examine a population of smokers and observe that everyone who smokes filterless cigarettes (A) develops lung cancer (E). Imagine that it is also the case that everyone in this population has yellow fingers (X). Hume maintains that, in these circumstances, human beings will have a tendency to confuse causation and correlation. After all, X has been repeatedly paired with E, and as a result, these events will become closely associated in our minds. Thus, custom and habit would lead us to infer that people with yellow fingers will get lung cancer, whether or not they smoke.

If Hume maintained that causal thought involves nothing but associations, then this would be all there is to the story. He would have to concede that we cannot discover causes in situations where there are confounding variables. But this is not what he says. Hume proceeds to argue that we can *correct* this propensity by employing what he calls "rules for judging causes and effects."

> We shall afterwards take notice of some general rules by which we ought to regulate our judgment concerning causes and effects . . . . By them we learn to distinguish the accidental circumstances from the efficacious causes; and when we find that an effect can be produc'd without the concurrence of any particular circumstance, we conclude that that circumstance makes not a part of the efficacious cause, however frequently conjoin'd with it. (p. 149)

These rules enable us to screen off accidental regularities. They direct us to search for instances in which an effect occurs without one of its associated conditions; if we discover that a particular circumstance is not always conjoined with an effect, we can eliminate it from the list of efficacious conditions. To return to the smoking example, if we observe someone who develops lung cancer (E) without yellow fingers (X) — suppose she wore gloves — then we will have eliminated this superfluous circumstance from the genuine regularity between smoking (A) and lung cancer (E).

Sometimes nature is kind and provides us with this information. In other cases, we must set up "new experiments" in order to discover the genuine causal conditions.

> There is no phenomenon in nature, but what is compounded and modify'd by so many different circumstances, that in order to arrive at the decisive point, we must carefully separate what is superfluous, and enquire by new experiments, if every particular circumstance of the first experiment was essential to it. (p. 175)

This method requires perseverance. We must vary each of the associated conditions and observe whether the effect still obtains. And we must continue with this procedure, as Hume puts it, until we have "enlarged our sphere of experiments as much as possible" (p. 175). It is only after we have conducted these elaborate investigations that we are entitled to assert that causal relations "really are so" (p. 173). Jonathan Bennett is mistaken, then, when he writes that Hume "exaggerates our intellectual passivity" and "ignores causal judgments which look interrogatively rather than confidently towards the future" (1971, p. 302). Hume recognizes that genuine regularities often lie beneath the surface of things, and that we must use sophisticated experimental procedures in order to discover them.

Hume's Experimental Investigations

There is another important reason why we should not interpret Hume as endorsing radical associationism about causal cognition: he does not rely upon custom and habit when he conducts his own experimental investigations. When Hume conducts his science of human nature, he subjects the regularities that he observes to careful examination. Perhaps the clearest illustration of Hume's sophisticated causal reasoning lies with his account of the indirect pas-

sions in a section of Book Two of the *Treatise* entitled "Experiments to Confirm this System." Hume observes that we typically feel passions such as pride and humility whenever we stand in close relation to agreeable or disagreeable objects. For example, we feel pride when we contemplate our prized possessions, and we feel shame when we reflect upon our embarrassing habits. In order to explain these complex emotions, Hume proposes that they are caused by what he calls the "double association" of impressions and ideas (p. 332). His hypothesis is that we feel indirect passions (E) whenever there are objects that (A) are closely related to us and (X) give us sensations of pleasure or pain.

Let us call this the "double association" hypothesis. In order to test it, Hume devises a series of new experiments that examine what would happen if each of these conditions was varied. For example, what if objects give us pleasure or pain (X), but we are unrelated to them (~A)? And what if we are closely related to these objects (A), but we are indifferent towards them (~X)? In these situations, Hume maintains, we would not feel any of the indirect passions (~E) [pp. 333–335].

> This is the reasoning I form in conformity to my hypothesis; and am pleas'd to find upon trial that every thing answers exactly to my expectation . . . . This exact conformity of experience to our reasoning is convincing proof of the solidity of that hypothesis on which we reason. (p. 338)

Hume isolates the cause of pride and humility, then, by considering situations in which their associated conditions are altered. These test cases provide him with the requisite confidence that he has isolated a genuine regularity in nature.

Hume does not rely upon the custom and habit, then, when he carries out his experimental investigations into the passions. But what is perhaps even more remarkable is that Hume does not rely upon custom and habit when he came up with his custom–habit account! Hume's sophisticated approach to causal reasoning is on clear display when he examines the psychological processes that underlie our everyday causal attributions. Let us look carefully, then, at his methodological approach to the psychology of causal inference. The first step involves observation. Hume notices the following fact about human behavior: we make causal attributions whenever we see constant conjunctions. This observation can itself be viewed as a *second-order* constant conjunction between (A) cases where subjects witness constant conjunctions, and (E) cases where subjects make causal attributions. What Hume observes, in other words, is that A is usually followed by E.

If Hume's own causal reasoning depended upon custom and habit, then he would infer that *A causes E*. But this is not what happens. Hume never asserts that the correlation between A and E is sufficient to make causal judgments.

Rather, he begins the second step of his experimental investigation and attempts to understand what goes on in our minds when we make such inferences. That is, Hume seeks to discover the psychological faculty (X) which, together with experience (A), produces these attributions (E). Notice that causal discovery in this case moves in the opposite direction from the smoking example. In the case of accidental regularities, one must *subtract* causally irrelevant factors in order to locate the genuine regularity; in the case at hand, one must *add* causally relevant factors to do so. This additional stipulation is necessary because the psychological processes that underlie these causal attributions are hidden from view; in contemporary terms, they are hidden variables.

Hume proceeds to formulate two hypotheses concerning the nature of these psychological faculties. According to the first, which we might call the "rationalist hypothesis," it is the faculty of reason which, together with the senses, leads us to make the causal attributions that we do. However, Hume argues that this hypothesis cannot possibly be correct. If reason was the faculty that we are looking for, then we must have non-circular justifications for drawing the conclusions that we do. But there are no such justifications (pp. 89–92). Therefore, reason cannot be responsible for these inferences. The only remaining explanation is the custom–habit hypothesis. On this alternative account, our everyday causal attributions are explained in terms of the interaction between sensory experience and custom and habit. We do not ordinarily make causal inferences by considering reasons or arguments; rather, we are associatively primed to do so. The relevant psychological process is not reflection, but conditioned reflex.

In sum, Hume's own causal reasoning has little in common with the custom–habit account:

1. Hume's causal reasoning is active rather than passive. On the custom–habit account, one sees a constant conjunction and, without knowing why, one makes a causal attribution. But Hume never embraces the naïve view that correlations are sufficient for causal attributions. Hume *qua* scientist of human nature does not just sit back and observe case after case of A followed by E, each time becoming more and more confident until he attributes a causal connection to them. Rather, Hume regards the conjunction between A and E as the starting point for further inquiry.

2. Hume's causal reasoning involves "reflection" (p. 148). Our everyday causal attributions take place implicitly and automatically; as he puts it, "custom operates before we have time for reflexion" (p. 104). But this is not the case with Hume's experimental reasoning. He consciously formulates hypotheses, devises new experiments, and determines which conditions to vary. Sometimes these experiments screen off accidental factors; other times they postulate genuine factors. This method requires careful and tireless investigation. As Hume

puts it, the rules for judging causes and effects are "extremely difficult" to apply and require "the utmost stretch of human judgment" (p. 175).

3. Hume's causal reasoning cannot be explained solely in terms of associations. Indeed, Hume's sophisticated experimental reasoning requires that we *regulate* the associative propensities of the imagination. As we have seen, these associations are what lead us to confuse accidental and genuine circumstances, and it is only when we "correct" this "bias" that we have any hope of making genuine causal discoveries (p. 148).

Hume's Official Position Regarding Causal Reasoning

As we have seen, Hume does not endorse radical associationism either in theory or practice. He is quite sober about the fact that associations cannot tell us all there is to know about how human beings think about causation. If that were the case, then we could not (as we do) distinguish between accidental and genuine regularities, and Hume could not (as he does) carry out his science of human nature. In short, Hume acknowledges that human beings are quite capable of sophisticated causal reasoning.

We can now see that there is a significant tension in Hume's attempt to become the "Newton of the mind." There are two distinct senses in which it might be said that Hume's science of human nature is Newtonian. The first is methodological: Hume maintains that the experimental methods of modern physics can be imported to the human sciences. The second is substantive: Hume hopes to show that the dynamics of the "mental world" can be completely understood in terms of laws of association, just as changes in the material world have been fully explained by laws of motion.

> Here is a kind of ATTRACTION, which in the mental world will be found to have as extraordinary effects as in the natural, and to shew itself in as many and as various forms. (pp. 12–13)

At first glance, these two projects seem to fit together quite well: Hume would use experimental methods in order to discover the universal laws of the mind. But when one probes deeper, it becomes apparent that these two commitments are not easily reconciled.

The problem is that, if the substantive claim is true, then Hume could never have carried out his experimental investigations. Consider his examination of our everyday causal inferences. Hume would have learned to associate A events (i.e., cases where we observe constant conjunctions) and E events (i.e., cases where we make causal attributions). But he could never have moved to the second step of the investigation and discovered the deeper regularity (A+X→E) that explains why we make the attributions that we do. Indeed, we might pose this as a general paradox about radical associationism: *if human*

cognition was governed solely by the laws of association, then nobody could discover that this was the case.

The trouble arises from the self-reflexive nature of Hume's science of human nature. Hume must rely upon his own faculties in order to study human cognition. But we can now see that this imposes constraints on what he can say about the operations of the human mind; in particular, he cannot make substantive claims that would undermine his capacity to carry out his investigations. If the laws of association were presented as a complete account of human cognition, however, it would violate this condition. Hume cannot therefore combine these two aspects of his Newtonian project — something has to give. In the end, what gives is the thesis that associations provide a comprehensive account of causal cognition.

One of Hume's fundamental insights is that human beings have more in common with the rest of the animal kingdom than we are willing to admit. In order to support this deflationary picture of human nature, he sets out to show that much of our cognitive lives can be explained in terms of elementary principles of association. But Hume recognizes, by the end of Book One, that there are essential limits to this reductive project. Causal judgment is one place where reason cannot completely yield to custom and habit. The problem is that human beings not only learn associations, but we also reflect upon them. We are not entirely creatures of habit.

Contemporary Associative Learning Theory to the Rescue?

In cognitive science, the question of how we ordinarily think about causation has been turned into an empirical research program. Whereas Hume relies upon informal observations about human behavior, contemporary psychologists have developed a quantitative experimental paradigm with which to investigate causal reasoning. In these experiments, subjects are presented with trials in which potential causes are followed by an effect. Subjects are then asked to evaluate the causal relationship between these various events. The results of these experiments are unequivocal: subjects are capable of making sophisticated causal judgments.

Consider the phenomenon of blocking. Anthony Dickinson and his colleagues demonstrated that subjects are quite adept at selecting among potential causal factors when making causal judgments (Dickinson, Shanks, and Evenden, 1984). In their "forward blocking" experiments, subjects are initially presented with trials in which A events are followed by E events, and they are subsequently presented with trials in which A+X are followed by E. Subjects are then asked to make a judgment about the relationship between X and E. The results were consistent: subjects deny that X causes E, even though these events were repeatedly paired together. David Shanks (1985)

showed that subjects can also make these judgments when the order of trials is reversed (i.e., "backwards blocking"). If subjects are shown trials in which A+X are followed by E, and subsequently presented with trials in which A alone is followed by E, they inevitably judge that there is no causal relation between X and E.

These blocking experiments demonstrate that subjects disregard potential causes if they are statistically redundant. Indeed, these studies suggest that our untutored causal attributions conform quite closely to the prescriptions of the rules for judging causes and effects. In the forward blocking experiments, for example, subjects observe instances where A alone is followed by E. The rules for judging causes and effects would direct them to conclude that A causes E. And this is precisely what subjects do in subsequent trials: they judge that X is an accidental factor. This is true of the backwards blocking experiment as well. Subjects observe that E can occur without X. The rules for judging causes and effects would direct them to conclude that X is an unnecessary part of the genuine causal regularity. Once again, subjects conform to this normative recommendation and judge that A is the cause of E.

The contemporary controversy over causal learning concerns the question of how subjects manage to solve these problems (Shanks, Holyoak, and Medin, 1996). What are the computational processes that underlie behavior in these tasks? On one side of the debate are those who defend an "associative approach." These researchers maintain that the causal learning experiments can be fully explained in terms of excitatory and inhibitory connections in associative memory; some of them appeal to the Rescorla–Wagner model of conditioning (Dickinson, Shanks, and Evenden, 1984; Shanks, 1985), whereas others invoke connectionist models (Read and Marcus–Newhall, 1993; Van Overwalle and Van Rooy, 1998). On the other side of the debate are those who defend a "rule-based approach." Some of these researchers maintain that subjects use statistical rules in order to compute causation (Cheng and Holyoak, 1995); others postulate that subjects employ deductive rules in order to do so (De Houwer and Beckers, 2002).

Ockham's razor dictates that, all else being equal, we should choose the simplest hypothesis. If the causal learning experiments can be accounted for in terms of nothing but associations, then, we ought to prefer this explanation over ones that posit more elaborate computational processes. But is it actually the case that the associative hypothesis can save the phenomena? As Hume pointed out, it is rather difficult to see how this could be so. Associations simply do not seem up to the task of explaining how we manage to distinguish between accidental and genuine causal regularities; quite the contrary, it appears that they are directly responsible for our propensity to collapse correlation and causation. In order to correct these mistakes, Hume maintains, we must invoke rules for judging causes and effects.

Hume's official position is that sophisticated causal thought cannot be explained in terms of associations. But is this merely a lack of imagination on his part? One might object that Hume drastically underestimates the power and complexity of associative learning theory. After all, associative psychology has made significant advances since the time that Hume wrote (Gormezano and Kehoe, 1981). Moreover, it has been shown that simple associative systems can perform extraordinary complex computations (Siegelmann and Sontag, 1995). Indeed, it is a truism in cognitive science that what *seems like* rule-governed behavior can emerge from systems that do not literally employ rules (Bereiter, 1991; Rumelhart and McClelland, 1986, p. 32). There is reasonable hope, then, that our capacity to make sophisticated causal judgments can be reduced to the level of associations.

This is not idle speculation. Contemporary psychologists have labored hard over the last twenty years to show that the causal learning experiments can be accommodated by the associative approach. For example, Dickinson, Shanks, and Evenden (1984) showed that the forward blocking studies can be explained in terms of the principles of classical animal conditioning. The crucial point is that classical conditioning involves cue competition: when multiple potential causes are followed by an effect, they will compete to predict its next occurrence (Rescorla and Wagner, 1972). It is also significant that associative learning in this model is driven by error-correction. This means that when a potential cause reliably predicts an effect, there will be no change in associative weights. Given these principles, it becomes relatively easy to account for the forward blocking data: the association formed between A and E in the first trial will block any association between X and E in subsequent trials. After all, A has zero competition during the initial trials, and as a result, the connection between A and E will achieve maximal strength; since A reliably predicts E in the later trials, the association between X and E will remain unchanged.

The backwards blocking experiments are more difficult to explain in terms of classical conditioning. The problem is that the standard model does not allow *absent cues* to enter into cue competition. The subsequent trials in which A alone is followed by E, therefore, cannot weaken the initial associations that were formed between X and E. This problem led Van Hamme and Wasserman (1994), among others, to modify the classical model of conditioning so that the strengths of absent cues can be lowered, as long as they were previously paired with a cue that is present. This slight revision enables the associative hypothesis to account for the backwards blocking studies. Since A and X were paired during the initial trials, X would be expected to accompany A on subsequent trials; as a result, its absence would weaken the connection between X and E.

Higher-Order Causal Reasoning

Contemporary psychologists appear to have proven Hume wrong. Hume maintains that the laws of association do not allow us to distinguish between accidental and genuine regularities. In retrospect, we can see that he was pushed towards this negative conclusion simply because his own version of associationism did not allow for cue competition. It seems that a slight modification in associative learning theory is all that is required in order to explain sophisticated causal reasoning.

But associationism is not out of the woods yet. Recent causal learning experiments have presented a serious challenge to the associative approach (see De Houwer, Beckers, and Vandorpe, 2005, for an overview). For example, De Houwer and Beckers (2002) performed an experiment in which subjects were required to solve *recursive* backwards blocking tasks. Subjects were initially presented with trials in which A+B were followed by E and subsequently shown trials in which B+C were succeeded by E. In the final stage of the experiment, subjects were split into two groups: one group was presented with trials in which A was followed by E, and the other was shown trials in which A was not followed by E. The crucial result was that subjects in the former group were much more likely to judge that C, but not B, is a genuine cause of E. De Houwer and Beckers maintain that this result cannot be explained in terms of conditioning models, including those that incorporate absent cues (p. 149). These models might explain why these subjects deny that B is a genuine cause, since A had been paired with B during the initial stage of the experiment, and thus B would have been an absent but expected cue. But they cannot account for the fact that these subjects would draw any conclusions at all about C, since it was never paired with A at any point during the experiment.

This study suggests that subjects use higher-order reasoning in order to make causal judgments. Indeed, De Houwer and his colleagues conclude that subjects could not have solved the recursive blocking task unless they consciously reflected upon rules for judging causes and effects (De Houwer, Vandorpe, and Beckers, 2005, p. 46). The thought is that subjects must have reasoned as follows: A+B is followed by E; but A alone is followed by E; thus B is a superfluous factor; yet B+C is followed by E; thus C must be an efficacious cause of E. Even though A and C were never paired together, in other words, subjects can use deductive rules to draw conclusions about their relationship. De Houwer and Beckers (2003) provide further evidence for the claim that subjects use conscious reasoning strategies in order to make causal judgments. Their experiments demonstrate that the ability to solve recursive backwards blocking problems is severely attenuated when subjects are asked

to perform difficult secondary tasks. This implies that causal learning is effortful and involves controlled cognitive processes.

Even if these recent blocking experiments could be accommodated by the associative approach (which seems unlikely at the moment), there are independent reasons to think that this approach cannot provide an exhaustive account of causal reasoning. One problem is that associative learning theory can only explain our ability to select among potential causal factors when we are provided with the appropriate negative instances. But as Hume pointed out, nature is not always so kind. It might very well be the case that the rooster always crows before the sun rises. In cases such as these, we must *intervene* in the order of things to make causal discoveries. Indeed, recent evidence suggests that both adults (Lagnado and Sloman, 2004) and children (Sommerville, 2007) carry out such interventions in order to distinguish between accidental and genuine regularities.

Another problem is that associative learning theories fail to explain how subjects can discover unobserved causal factors. The blocking experiments involve situations where subjects must subtract causally irrelevant factors; but what about cases where they must add causally relevant features? As Hume pointed out, if we are to make causal judgments on the basis of partial information, we must use background knowledge and generate hypotheses about the complete set of factors at work. In fact, recent experiments have shown that people often rely upon sophisticated causal models when asked to make causal judgments in such circumstances (Waldmann, Hagmayer, and Blaisdell, 2006). As these researchers put it, "people . . . do not stick to the superficial level of event covariations, but reason and learn on the basis of deeper causal regularities" (p. 307).

The problem is that the associative approach makes causal learning an entirely passive affair. This lesson is particularly damaging for those like Paul Churchland (1989) who defend an associative account of scientific reasoning. Churchland maintains that recurrent connectionist networks can explain how scientists recognize events "qua instance of a type of causal or law-governed process" (1995, p. 105). As he puts it, these regularities would be represented as "lines" rather than "points" in the network's state-space. But it is difficult to see how this could be so. These networks simply learn to associate whatever events follow one another in time. This makes them excellent models of how we ordinarily make probabilistic inferences (Collier, 2005a). But it makes them poor models of sophisticated causal reasoning: the networks lack the capacity, as it were, to interrogate nature and discover her secrets. The networks extract statistical regularities from the data, but they cannot determine which ones are genuinely causal.

Conclusion

There is a growing consensus among contemporary psychologists that associations cannot fully explain causal learning (De Houwer, Vandorpe, and Beckers, 2005; Dickinson, 2001; Pineno and Miller, 2007; Shanks, 2007). Researchers have begun to direct their attention to hybrid models that incorporate associative and rule-based components. This enables researchers to ask important new questions, such as: In which contexts do subjects rely upon associations or rules? How are the rules for judging causes and effects acquired? How are they represented in the mind? In the end, the interesting issue is no longer whether causal cognition should be explained in terms of *either* associations *or* rules (Allan, 1993). As Hume taught us, this is a false dilemma: *both* associations *and* rules are required for a complete and comprehensive account of how we think about causation.

References

Allan, L.G. (1993). Human contingency judgments: Rule based or associative? *Psychological Bulletin, 114*, 435–448.

Bennett, J. (1971). *Locke, Berkeley, Hume: Central themes*. Oxford: Clarendon Press.

Bereiter, C. (1991). Implications of connectionism for thinking about rules. *Educational Researcher, 20*, 10–16.

Cheng, P.W., and Holyoak, K.J. (1995). Complex adaptive systems as intuitive statisticians: Causality, contingency, and prediction. In H.L. Roitblat and J.A. Meyer (Eds.), *Comparative approaches to cognitive science* (pp. 271–302). Cambridge, Massachusetts: MIT Press.

Churchland, P.M. (1989). *A neurocomputational perspective: The nature of mind and the structure of science*. Cambridge, Massachusetts: MIT Press.

Churchland, P.M. (1995). *Engine of reason, seat of the soul*. Cambridge, Massachusetts: MIT Press.

Collier, M. (1999). Filling the gaps: Hume and connectionism on the continued existence of unperceived objects. *Hume Studies, 25*, 155–170.

Collier, M. (2005a). A new look at Hume's theory of probabilistic inference. *Hume Studies, 31*, 21–36.

Collier, M. (2005b). Hume and cognitive science: The current status of the controversy over abstract ideas. *Phenomenology and the Cognitive Sciences, 4*, 197–207.

De Houwer, J., and Beckers, T. (2002). Higher-order retrospective revaluation in human causal learning. *Quarterly Journal of Experimental Psychology, 55*, 137–151.

De Houwer, J., and Beckers, T. (2003). Secondary task difficulty modulates forward blocking in human contingency learning. *Quarterly Journal of Experimental Psychology, 56B*, 345–357.

De Houwer, J., Beckers. T., and Vandorpe, S. (2005). Evidence for the role of higher order reasoning processes in cue competition and other learning phenomena. *Learning and Behavior, 33*, 239–249.

De Houwer, J., Vandorpe, S., and Beckers, T. (2005). On the role of controlled cognitive processes in human associative learning. In A.J. Wills (Ed.), *New directions in human associative learning* (pp. 41–63). Mahwah, New Jersey: Lawrence Erlbaum Associates.

Dickinson, A. (2001). Causal learning: An associative analysis. *Quarterly Journal of Experimental Psychology, 54B*, 3–25.

Dickinson, A., Shanks, D, and Evenden, J. (1984). Judgment of act-outcome contingency: The role of selective attention. *Quarterly Journal of Experimental Psychology, 54*, 29–50.

Glymour, C. (1998). Learning causes: Psychological explanations of causal explanation. *Minds and Machines, 8*, 39–60.

Gormezano, I., and Kehoe, J.E. (1981). Classical conditioning and the law of contiguity. In P. Harzem and M.D. Zeiler (Eds.), *Predictability, correlation, and contiguity* (pp. 1–45). New York: John Wiley and Sons.

Hume, D. (1978). *A treatise of human nature*. Oxford: Oxford University Press. (Original work published 1739)

Jessop, T.E. (1966). Some misunderstandings of Hume. In V.C. Chappell (Ed.), *Hume: A collection of critical essays* (pp. 35–52). New York: Anchor Books.

Lagnado, D., and Sloman, S. (2004). The advantage of timely interevention. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 30*, 856–876.

Pears, D. (1991). *Hume's system: An examination of the first book of his Treatise*. Oxford: Oxford University Press.

Pineno, O., and Miller, R. (2007). Comparing associative, statistical, and inferential reasoning accounts of human contingency learning. *The Quarterly Journal of Experimental Psychology, 60*, 310–329.

Read, S.J., and Marcus–Newhall, A. (1993). Explanatory coherence in social explanations: A parallel distributed processing account. *Journal of Personality and Social Psychology, 65*, 429–447.

Rescorla, R.A., and Wagner, A.R. (1972). A theory of Pavlonian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In A.H. Black and W.F. Prokasy (Eds.), *Classical conditioning II: Current research and theory* (pp. 64–99). New York: Appleton Books.

Rumelhart, D.E., McClelland, J.L., and the PDP research group. (1986). *Parallel distributed processing: Explorations in the microstructure of cognition*. Volume I. Cambridge, Massachusetts: MIT Press.

Shanks, D.R. (1985). Forward and backward blocking in human contingency judgment. *Quarterly Journal of Experimental Psychology, 37*, 1–21.

Shanks, D.R. (2007). Associationism and cognition: Human contingency learning at 25. *Quarterly Journal of Experimental Psychology, 60*, 291–309.

Shanks, D.R., Holyoak, K.J., and Medin, D.L. (Eds.). (1996). *The psychology of learning and motivation (Volume 34): Causal learning*. San Diego, California: Academic Press.

Siegelmann H.T., and Sontag E.D. (1995). On the computational power of neural networks. *Journal of Computer and System Sciences, 50*, 132–150.

Sommerville, J.A. (2007). Detecting causal structure: The role of interventions in infants' understanding of psychological and physical causal relations. In A. Gopnik and L. Shulz (Eds.), *Causal learning: Psychology, philosophy, and computation* (pp. 48–57). Oxford: Oxford University Press.

Van Hamme, L.J., and Wasserman, E.A. (1994). Cue competition in causality judgments: The role of nonpresentation and compound stimulus elements. *Learning and Motivation, 25*, 127–151.

Van Overwalle, F., and Van Rooy, D. (1998). A connectionist approach to causal attribution. In S.J. Reed and L.C. Miller (Eds.), *Connectionist models of social reasoning and behavior* (pp. 143–171). Mahwah, New Jersey: Lawrence Erlbaum Associates.

Waldmann, M.R., Hagmayer, Y., and Blaisdell, A.P. (2006). Beyond the information given: Causal models in learning and reasoning. *Current Directions in Psychological Science, 15*, 307–311.