



A New Look at Hume's Theory of Probabilistic Inference

Mark Collier

Hume Studies Volume 31, Number 1, (2005) 21 - 36.

Your use of the HUME STUDIES archive indicates your acceptance of HUME STUDIES' Terms and Conditions of Use, available at

<http://www.humesociety.org/hs/about/terms.html>.

HUME STUDIES' Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the HUME STUDIES archive only for your personal, non-commercial use.

Each copy of any part of a HUME STUDIES transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

For more information on HUME STUDIES contact

humestudies-info@humesociety.org

<http://www.humesociety.org/hs/>

A New Look at Hume's Theory of Probabilistic Inference

MARK COLLIER

1. Hume's Theory of Probabilistic Inference

Historians of philosophy do not usually take Hume's theory of probabilistic inference seriously. For some scholars, Hume's account is dismissed because of its misguided reliance upon psychological rather than logical methods.¹ Others are more sympathetic to Hume's naturalistic approach, but regard the specific proposals of his positive account as hopelessly naïve.

If his contributions are to be judged as part of the empirical science of man . . . then his 'results' will appear ludicrously inadequate, and there will be no reason to take him seriously.²

Still others are willing to defend many of Hume's positive proposals, but single out his account of probabilistic inference as "unsatisfactory"³ and "dubious."⁴ In this paper, I challenge these disparaging assessments. I argue that Hume's theory of probabilistic inference is neither misguided nor inadequate; quite the contrary, it stands at the leading edge of our contemporary science of the mind.

Hume agrees with Leibniz that previous philosophers have been "too concise when they treat of probabilities" (T Abs.4; SBN 647; cf. EHU 6.4; SBN 59). In section 1.3.12 of the *Treatise*, he attempts to remedy this situation by laying out a theory of "conjectural or probable reasonings" (T 1.3.12.20; SBN 139). It is important to be

Mark Collier is Assistant Professor of Philosophy, University of Minnesota, Morris, MN 56267 USA.
e-mail: mcollier@morris.umn.edu

clear at the outset that Hume is interested primarily in a *psychological* rather than a *metaphysical* approach to probability. Hume does in fact subscribe to a particular metaphysical interpretation of probability, according to which it is nothing but a reflection of our ignorance concerning hidden causes, but this position is clearly of secondary interest and is not one that he defends at any length. His main concern is to explain how we manage to make predictive inferences under conditions of uncertainty, and for this issue, questions about the metaphysical nature of probability are idle; our philosophical interpretations of probability, he maintains, have no influence on how we carry out probabilistic inferences in our everyday lives.

When causes are not followed by their usual effects, the vulgar take this as an indication of “contingency” in the cause, by virtue of which the *same* cause can sometimes produce *different* effects. In contrast, philosophers retain their commitment to the causal principle, and explain away putative counter-examples in terms of the “secret operation of contrary causes.”

The vulgar, who take things according to their first appearance, attribute the uncertainty of events to such an uncertainty in the causes, as makes them often fail of their usual influence, tho' they meet with no obstacle nor impediment in their operation. But philosophers observing, that almost in every part of nature there is contain'd a vast variety of springs and principles, which are hid, by reason of their minuteness and remoteness, find that 'tis at least possible the contrariety of events may not proceed from any contingency in the cause, but from the secret operation of contrary causes. (T 1.3.12.5; SBN 132)

Nevertheless, philosophers who reject causal indeterminacy have no choice but to rely upon probabilities for guidance in their everyday lives.

But however philosophers and the vulgar may differ in their explication of the contrariety of events, their inferences from it are always of the same kind, and founded on the same principles. (T 1.3.12.6; SBN 132; cf. EHU 6.4; SBN 58)

When philosophers make decisions or predictions under conditions of uncertainty, they must make probabilistic calculations just like the vulgar.

Hume's primary concern in T 1.3.12 involves the nature of our commonplace probabilistic inferences. What types of sensory perceptions lead us to make them? What degrees of belief do they generate? Which faculties of the mind enable us to draw such inferences? Hume regards these as empirical questions, and in order to make progress on them, he turns to the resources of his science of human nature. His strategy is to show that probabilistic inferences are a *species* of inductive

inferences, and therefore can be explained in terms of the “same principles” (T 1.3.11.1; SBN 124). In order to properly understand Hume’s theory of probabilistic inference, then, we must briefly review his psychological explanation of induction in T 1.3.6.

Hume begins his psychological explanation of induction by describing the behavior of his fellow men: we have a tendency to make inductive inferences whenever we observe a conjunction between two types of events (T 1.3.6.2; SBN 87). Following Stroud, we can reconstruct Hume’s description in the following terms.

Inference from Experience

Past Experience (PE): Previously observed A-type events have been followed by B-type events.

Present Impression (PI): X is an A-type event.

Future Expectation (FE): X will be followed by a B-type event.

In other words, Hume discovers the following psychological fact about human beings: each time we witness sensory perceptions such as (PE) and (PI) in the above schema, we come to have the type of expectation described in (FE).

In the next step of his investigation, Hume attempts to explain this fact by drawing on his theory of the imagination. His hypothesis is that our capacity to make inferences from experience depends upon the interaction of the sensory information registered in (PE) and (PI) with associative principles of the imagination.

When the mind, therefore, passes from the idea or impression of one object to the idea or belief of another, it is . . . determin’d . . . by certain principles, which associate together the ideas of objects, and unite them in the imagination. (T 1.3.6.12; SBN 92; cf. EHU 5.2–5; SBN 41–3)

The faculty of imagination is governed by three laws of association: contiguity, resemblance, and causation (T 1.1.4.1–4; SBN 11–12). The task of Hume’s psychological explanation is to show that these minimal resources are all that is needed in order to explain why we make the inductive inferences that we do.

It is the principle of resemblance, according to Hume’s hypothesis, that accounts for why we assimilate the present impression (PI) to the previously observed event types (PE).

In reality, all arguments from experience are founded on the similarity, which we discover among natural objects, and by which we are induced to expect effects similar to those which we have found to follow from such objects. (EHU 4.20; SBN 36; cf. T 1.3.6.14; SBN 93)

Hume borrows here from the theory of general ideas that he developed in T 1.1.7. On that account, the imagination automatically categorizes objects and events according to the class of instances towards which they have the highest degrees of resemblance (T 1.1.7.15; SBN 23). We classify the event token (PI) as an A-type event, in other words, because it is more similar to A-type events than any other event types in memory.⁵

The principle of resemblance accounts for why we categorize the event token as an A-type event, but it does not explain why we anticipate that it will be followed by a B-type event. In order to deal with this further fact, Hume appeals to another law of association, the principle of causation. The principle of causation states that whenever we repeatedly observe the relation of contiguity between event types, they will become connected in the imagination. Thus, since A-type events have always been followed by B-type events, the principle of causation entails that we will infer from the A-type event (PI) to the B-type event (FE). Hume's official psychological explanation of induction, then, is that the *principle of resemblance* explains why the event token (PI) is assimilated to the A-type events in memory (PE), and the *principle of causation* explains the expectation of a B-type event (FE).

Hume's project in T 1.3.12 is to demonstrate that these same associative principles can be put to work in order to explain how we make probabilistic inferences.

The probabilities of causes are of several kinds; but are all deriv'd from the same origin, viz. *the association of ideas to a present impression*. (T 1.3.12.2; SBN 130)

Hume begins his examination of probabilistic inferences, as he did with inductive inferences, with an observation concerning the common behavior of human beings. He notices that we typically make probabilistic inferences whenever we perceive *inconstant* conjunctions between events. Moreover, he observes that when we perceive inconstant conjunctions, our future expectations are accompanied by *partial degrees of belief*; we say that it is "likely" or "probable" that they will co-occur in the future.⁶ In the next step of his investigation, Hume attempts to explain *why* inconstant conjunctions give rise to partial degrees of belief. His strategy once again is to appeal to his theory of the imagination; his goal is to show that the various species of probabilistic inference can be explained, without remainder, in terms of the elementary principles of association.

The first species of probabilistic inference, according to Hume, occurs whenever the conjunction between events in (PE) involves a *small sample* (T 1.3.12.2; SBN 130–1). We can characterize this type of probabilistic inference in the following terms.

Inference from Small Sample

PE: A *small sample* of previously observed A-type events have been followed by B-type events.

PI: X is an A-type event.

FE: X has a *minimal likelihood* of being followed by a B-type event.

When we observe one event follow another repeatedly, but not extensively, we are willing to infer one from the other, but we do so with hesitation; as Hume puts it in the *Enquiries*, “it is only after a long course of uniform experiments in any kind, that we attain a firm reliance and security, with regard to a particular event” (EHU 4.20; SBN 36).

Hume maintains that the associative principle of causation explains the fact that our inferences from small samples are characterized by relatively low degrees of belief.

As the habit, which produces the association, arises from the frequent conjunction of objects, it must arrive at its perfection by degrees, and must acquire new force from each instance, that falls under our observation. The first instance has little or no force: The second makes some addition to it: The third becomes still more sensible; and 'tis by these slow steps, that our judgment arrives at a full assurance. (T 1.3.12.2; SBN 130)

According to the principle of causation, the strength of the association between events is a function of the frequency of their co-occurrence. Hume's associationist hypothesis, therefore, predicts that our assurance will gradually increase in proportion to the size of the sample and that small samples would generate relatively low levels of confidence.

The second species of probabilistic inference involves cases where the sample in (PE) is large, but where the present impression in (PI) has a *partial resemblance* to the events in memory.⁷ Hume refers to such cases as “probability deriv'd from analogy” (T 1.3.12.25; SBN 142). He has in mind the following type of inference from experience.

Inference from Analogy

PE: Previously observed A-type events have been followed by B-type events.

PI: X *partially resembles* an A-type event.

FE: X will *likely* be followed by a B-type event.

Hume maintains that inferences from analogy can be explained in terms of the principles of association. The principle that does the explanatory work in this case is the associative law of resemblance.

[I]n the probability deriv'd from analogy, 'tis the resemblance only, which is affected. Without some degree of resemblance, as well as union, 'tis impossible there can be any reasoning: but as this resemblance admits of many different degrees, the reasoning becomes proportionally more or less firm and certain. (T 1.3.12.25; SBN 142)

The principle of resemblance entails that inferences from analogy will be attended with varying levels of uncertainty. The crucial point is that the principle of resemblance “admits of many different degrees.” The stronger the similarity between past and present events, then, the more inductive confidence we will have in our future expectations. Since the degrees of belief in (FE) are proportional to the degree of resemblance between (PI) and (PE), it follows that whenever we observe partial resemblances, our degrees of belief will be partial as well.

Hume has shown that his associationist psychology can account for how we ordinarily make inferences from small samples and partial resemblances. He manages to do so, we have seen, because of the flexibility of the principles of association. There are two ways in which the conjunction between events can be inconstant: the *quantity* of the sample can be small or there can be *qualitative* variation among its instances. The principles of the imagination explain how we make probabilistic inferences in each case, Hume maintains, because the strength of the association between events will vary in proportion to the constancy of the conjunction. As Hume puts it, “[i]f you weaken either the union or resemblance, you weaken the principle of transition, and of consequence that belief, which arises from it.”⁸ Hume’s associationist hypothesis not only explains how we make inferences from experience in such cases, then, but it also explains why we do so with varying levels of confidence.

Hume’s associationist hypothesis faces a more difficult challenge, however, with the third species of probabilistic inference. The third species of probabilistic inference involves cases where the conjunction between events in (PE) is composed of *mixed frequencies*, or what Hume calls “contrariety” (T 1.3.12.4–19; SBN 131–8).

Twou’d be very happy for men in the conduct of their lives and actions, were the same objects always conjoin’d together, and we had nothing to fear but the mistakes of our own judgment, without having any reason to apprehend the uncertainty of nature. But as ’tis frequently found, that one observation is contrary to another, and that causes and effects follow not in the same order, of which we have had experience, we are oblig’d to vary our reasoning on account of this uncertainty, and take into consideration the contrariety of events. (T 1.3.12.4; SBN 131)

We can represent inferences from mixed frequencies in terms of the following schema.

Inference from Mixed Frequency

PE: *Some* previously observed A-type events have been followed by B-type events, and *some* previously observed A-type events have been followed by C-type events.

PI: X is an A-type event.

FE: X will *likely* be followed by a B-type event or a C-type event.

Hume's examples of such "irregular" conjunctions typically involve medical cases. Sometimes rhubarb proves a purge and sometimes it does not; sometimes opium puts one to sleep and other times it does not (EHU 6.4; SBN 57–8). Once again, philosophers do not regard such irregularities as violations of the causal principle, but merely as a reflection of our ignorance concerning the real causes at work; nevertheless, "[o]ur reasonings . . . and conclusions concerning the event are the same as if this principle had no place" (T 1.3.12.25; SBN 142). That is, when philosophers must decide whether or not to ingest rhubarb or opium, they have no choice but to rely upon mixed frequencies in order to calculate the probability that these medicines will prove to be effective cures.

In order to explain how we make inferences from mixed frequencies, Hume once again turns to the resources of his science of human nature. Let us suppose, for simplicity's sake, that A-type events have been followed by B-type events four times, and A-type events have been followed by C-type events three times. According to Hume's exemplar-based theory of general ideas, this frequency information will be represented in memory in terms of separately stored instances.

AB = {a1 b1, a2 b2, a3 b3, a4 b4}

AC = {a1 c1, a2 c2, a3 c3}

What will happen, then, the next time we observe an A-type event?

Hume claims that there are "two hypotheses" concerning the manner in which we "transfer" these event sequences from memory to our future expectations.

First, That the view of the object, occasion'd by the transference of each past experiment, preserves itself entire, and only multiplies the number of views. Or, *secondly*, That it runs into the other similar and correspondent views, and gives them a superior degree of force and vivacity. (T 1.3.12.19; SBN 138)

The first hypothesis, in other words, is that we transfer *all* the particular events that have been associated with A-type events in the past. If this were the case,

then the content of our future expectation would consist of a disjunctive list of event tokens.

$$(FE) = \{b1 \vee b2 \vee b3 \vee b4 \vee c1 \vee c2 \vee c3\}$$

Hume maintains that we need only introspect, however, in order to recognize the implausibility of this hypothesis. Experience informs us that our future expectations consist “in one conclusion, not in a multitude of similar ones” (T 1.3.13.25; SBN 142). Moreover, it is implausible on theoretical grounds to maintain that the mind can represent, at one time, a long list of events; as Hume puts it, the disjunctive list events would usually be “too numerous to be comprehended distinctly by any finite capacity” (T 1.3.13.25; SBN 142).

When we make inferences from mixed frequencies, then, it must be the case that the “similar views run into each other, and unite their forces” (T 1.3.13.25; SBN 142). The only plausible hypothesis, in other words, is that we perform a *summary computation* when we transfer event sequences from memory. The separately stored instances, as Hume puts it, are united by the principles of the imagination into a “general view” (T 1.3.12.17; SBN 137).

When we transfer contrary experiments to the future, we can only repeat the contrary experiments with their particular proportions; which cou’d not produce assurance in any single event, upon which we reason, unless the fancy melted together all those images that concur, and extracted from them one single idea or image, which is intense and lively in proportion to the number of experiments from which it is deriv’d, and their superiority above their antagonists. (T 1.3.12.22; SBN 140; cf. EHU 6.3; SBN 57)

Through this process of amalgamation, similar events combine their strengths and contrary events cancel each other out. As a result, a composite representation of the frequency information will be formed in the imagination. Moreover, when we observe another A-type event, our future expectations will be proportional to the relative frequencies with which B-type and C-type events have followed A-type events in the past.

Hume does not provide any precise account of the processes whereby the separately stored event sequences are brought together into a single representation. The problem is that principles of the imagination do not supply him with the resources to do so. The principle of resemblance accounts for the classification of the event token in (PI), but it does not explain why we develop the future expectations that we do. The principle of causation explains why the A-type event (PI) gives rise to the future expectation of B-type or C-type events, but it does not explain how we manage to form a single probability estimate that is proportional to the mixed frequencies in memory. Indeed, the principle of causation would appear to support the first hypothesis concerning the transference of particular

event sequences from memory; after all, *each* of the particular event tokens has been associated with A-type events in the past.

Barry Gower is correct to point out, then, that “it is hard to see how to account for any probabilities arising from insufficiency of evidence in terms of an associationist psychology.”⁹ Hume recognizes this difficulty, and it leads him to lapse into metaphorical talk of particular event tokens “melting together” into composite representations. Hume glimpses, somewhat darkly, that the solution to this problem must involve summary computations whereby the mixed frequencies stored in memory are combined into a unified representation. But Hume does not offer a sufficient explanation of the combinatorial process through which “we extract a single judgment from a contrariety of past events” (T 1.3.12.8; SBN 134).

This lacuna in Hume’s theory of probabilistic inference has led Hume scholars to turn to non-associationist resources in order to explain how inferences from mixed frequencies are performed. The dominant tendency is to reconstruct these inferences in terms of the Carnap-Reichenbach Straight Rule or the mathematical rules of Bayesianism.¹⁰ Others interpret the combination of mixed frequencies in terms of a theory of mental oscillations; on this account, probabilities are measured by the amount of time it takes to survey the various event sequences.¹¹ But these interpretations unnecessarily leave behind the spirit and letter of Hume’s account. Hume makes it quite clear, after all, that the combination of mixed frequencies involves “an operation of the *fancy*” (T 1.3.12.22; SBN 140).¹² What is needed in order to defend Hume, in a manner consistent with his general approach to human nature, is a precise associationist account of how particular event sequences in memory can be combined into a single representation. And as we shall see in the next section, such an account is now available to us. Contemporary associative theories provide the resources that are needed in order to explain how single probability estimates can be performed on the basis of mixed frequencies.

2. Recent Evidence for Hume’s Theory of Probabilistic Inference

In the Introduction to the *Treatise*, Hume promises to ground his science of human nature on “careful and exact experiments” (T Intro.8; SBN xvii). In practice, however, Hume’s experimental methods appear substandard when compared to those of his contemporaries, such as the physicists in the Royal Society.¹³ Hume was aware of the laboratory experiments being performed in the physical sciences; he merely thought that they could not be applied to the human sciences.

Moral philosophy has, indeed, this peculiar disadvantage, which is not found in natural, that in collecting its experiments, it cannot make them purposely, with premeditation, and after such a manner as to satisfy itself concerning every particular difficulty which may arise. When I am

at a loss to know the effects of one body upon another in any situation, I need only put them in that situation, and observe what results from it. But should I endeavor to clear up after the same manner any doubt in moral philosophy, by placing myself in the same case with that which I consider, 'tis evident this reflection and premeditation would so disturb the operation of my natural principles, as must render it impossible to form any just conclusion from the phenomenon. (T Intro.10; SBN xix)

It simply never occurred to Hume that he need not perform these experiments on himself, and that he could make use of experimental subjects (such as undergraduates) who would carry out the tasks without “premeditation.” Of course, one can easily excuse Hume for this oversight, since as Daniel Robinson points out in *Toward a Science of Human Nature*, it was not until the nineteenth and twentieth centuries that psychologists developed the rigorous methods with which we are familiar today.¹⁴

One must concede the point, then, that “Hume had no way of empirically testing his hypothesis.”¹⁵ We need not speculate about how well his hypothesis would have held up under examination, however, since contemporary psychologists have devised an experimental paradigm with which to test it. In these experiments, known as probability learning tasks, subjects are presented with frequency information about the co-occurrence of events, and are asked to estimate the degree to which these events are related.¹⁶ These subjective ratings are then compared with a normative standard, called “contingency,” which measures the actual co-variation between the events.¹⁷

There are three important results of the probability learning task experiments. First, the experimental findings demonstrate that subjects are extremely sensitive to the degrees of contingency in the data.¹⁸ In experiment after experiment, the subjective ratings of the relation between events correspond quite closely to their actual co-variation. Second, the experimental results show that contingency learning proceeds in gradual fashion; the ratings typically start close to zero, and increase in small steps until they correspond to the objective degree of contingency in the data.¹⁹ Finally, the probability learning tasks reveal that the contingency ratings depend upon the degree of resemblance between the events presented during training; the more similar the events, the more strongly they become associated.²⁰

These experimental studies provide confirmation, then, for Hume’s claim that we adjust our degrees of belief according to the resemblance, contrariety, and sample size of the events that are perceived. But do they also support Hume’s contention that our capacity to make probabilistic inferences can be exhaustively explained in terms of the principles of association? The probability learning experiments demonstrate that subjects proportion their degrees of belief to the evidence,

but they do not tell us *how* they manage to do so. The psychological experiments provide us with precise measures of the sensory input (frequency information and stimulus similarity) and behavioral output (confidence ratings), but they remain silent about the psychological processes that underlie performance in the tasks.

The most influential explanation of contingency learning in contemporary psychology is the Rescorla-Wagner Model.²¹ According to this model, probabilistic learning can be analyzed in terms of a competitive learning rule that modifies associative weights on a trial-by-trial basis.²² The Rescorla-Wagner Model has proven extremely effective in accounting for the results of the probabilistic learning task experiments.²³ First, a learning rate parameter in the model predicts that the learning curves observed in the experiments will be gradual in nature. Second, the model can explain, through stimulus-generalization, why the degree of resemblance between events plays an important role in the subjective ratings.²⁴ Third, the competitive nature of the learning rule entails that subjects will adjust their probability estimates according to the contrariety between events. Indeed, it has been demonstrated that, when there are two variables, the Rescorla-Wagner learning rule is mathematically equivalent to the measure of contingency.²⁵ When the Rescorla-Wagner learning rule modifies associative strengths on a trial-by-trial basis, therefore, it implicitly calculates the co-variation between events.

The probability learning experiments can also be explained, at the mechanistic level, in terms of adaptive neural networks. Gluck and Bower demonstrate that two-layer connectionist networks can simulate the behavior in the experiments.²⁶ In addition, David Shanks shows that these networks exhibit excellent fit to the learning curves described in the experimental results; the networks, like the subjects, begin with no sensitivity to the contingency between events, but improve on a trial by trial basis, until eventually they converge on the actual degree of contingency in the training sample.²⁷ The ability of adaptive neural networks to explain how we make contingency judgments comes as no surprise. These networks rely upon a learning rule, known as the Delta Rule, which is formally equivalent to the Rescorla-Wagner rule.²⁸ As a result, when a connectionist network modifies its associative weights according to the Delta Rule it is, in effect, computing the contingency between events.²⁹

The networks in these simulations incorporate feed-forward architectures, however, and thus are unable to learn about statistical dependencies that span across event sequences.³⁰ In order to model temporal contingency learning, Axel Cleeremans and his colleagues at Carnegie Mellon University have turned to Simple Recurrent Networks. The network is recurrent because information not only flows from the sensory input layer to the hidden layer, but also back down from the hidden layer to the input layer. This recurrent connection provides the network with short-term memory, which is necessary in order for the network to learn about event sequences that unfold over time, such as the following example.

ABCDEABCEDEABCEDEABCEDEACBEDABCEDEABCEDE . . .

Notice that this sequence is composed of *recurring* event types with different relative frequencies; for example, although A-type events reliably predict B-type events, C-type events only sometimes predict D-type events. From a computational point of view, the sequence of events constitutes a *probabilistic function*, and the task the network faces is to learn the mapping between event types.

The Simple Recurrent Network learns to approximate this probabilistic function by changing its weights in such a way as to drive down the errors in its predictions from state to state, and it eventually settles on a single set of weights that associates each of the cues with their respective outcomes. Suppose, for example, that the data set on which the network is trained consists of a non-deterministic sequence in which C-type events have been followed by D-type events forty percent of the time and E-type events sixty percent of the time. With sufficient training, the network's hidden units will be "shaded" in such a way that the next time it observes a C-type event it will estimate the conditional probabilities of the possible successors—D and E—in a manner that is proportional to their past frequencies.³¹

These computational simulations show, then, that our capacity to make complex probabilistic inferences can be accounted for in terms of the operations of a simple associative mechanism. After all, when the networks classify a novel event token as an instance of a type, they do so according to its *degree of resemblance* to event classes in memory. Moreover, the future expectations of the networks depend upon the *degree of constancy* in the conjunction between event types; the networks simply expect the event type that has followed the novel event token with the highest frequency. The SRN model of contingency learning also provides an elegant solution to the problem, recognized by Hume, concerning the processes whereby mixed frequencies are combined into a single probability estimate. In order to solve the probability learning task, the SRN updates the configuration of its hidden unit weights on each trial, which entails that the network will automatically summarize the event frequencies through a process known by connectionist researchers as superposition.³²

These psychological experiments and computational models from cognitive science provide convergent evidence, therefore, for Hume's hypothesis that the various species of probabilistic inferences can be explained in terms of elementary associationist principles. In the Rescorla-Wagner Model and neural networks, probabilistic inferences are understood solely in terms of the automatic, implicit, trial-by-trial adjustment of associative connections. Contemporary researchers on associative learning agree with Hume that our commonplace probabilistic inferences can be understood in this manner; they merely disagree with Hume over whether or not he provided a complete account of the associative learning principles. Hume thought that our probabilistic inferences could be exhaustively explained in terms of the principles of resemblance and causation. The recent

evidence from cognitive science, however, suggests that his account must be supplemented with “constraining principles” such as superposition and cue competition.³³ In any case, this addition would be welcomed by Hume, who admits that his enumeration of the principles of association is revisable and open-ended (EHU 1.3.3; SBN 24).

While historians of philosophy tend to treat Hume’s theory of probabilistic inference as an embarrassment, his account receives much warmer praise from contemporary researchers in cognitive science who have turned the question of how we make probabilistic inferences into an empirical research program.

It is Hume’s model, or refinements of it, which have come to be adopted by many contemporary psychologists, and which seem indeed to be best confirmed by the experimental data on animals and humans.³⁴

It is no overstatement to say that over the last half century “the Humean view” has become the dominant position in psychological research on how we make inferences under conditions of uncertainty.³⁵ We must resist, therefore, the tendency of Hume scholars to treat the particular hypotheses of his science of human nature as a source of disrepute. In the case of Hume’s theory of probabilistic inference, at least, his account has much more going for it than interpreters have previously recognized. It is simply unfair to claim that Hume’s theory of probabilistic inference is “speculative psychology which is not . . . of lasting interest.”³⁶ On the contrary, it stands at the leading edge of our contemporary science of the mind.

NOTES

An earlier version of this paper was presented at the 30th Annual Hume Conference on “Probability, Chance, and Judgment” at the University of Nevada, Las Vegas.

References to Hume’s works will be inserted into the text, using a letter or acronym for the titles followed by the page number, as follows:

EHU = David Hume, *An Enquiry Concerning Human Understanding*, ed. T. L. Beauchamp (Oxford: Oxford University Press, 1999).

T = David Hume, *A Treatise of Human Nature*, ed. D. F. Norton and M. J. Norton (Oxford: Oxford University Press, 2000).

SBN = David Hume, *A Treatise of Human Nature*, ed. L.A. Selby-Bigge, 2nd edition, revised by P. H. Nidditch (Oxford: Clarendon Press, 1978) and David Hume, *Enquiries concerning Human Understanding and concerning the Principles of Morals*, ed. L.A. Selby-Bigge, 3rd edition, revised by P. H. Nidditch (Oxford: Clarendon Press, 1975).

1 D. C. Stove, *Probability and Hume's Inductive Skepticism* (Oxford: Clarendon Press, 1973), 120; cf. A. Flew, *David Hume: Philosopher of Moral Science* (Oxford: Blackwell Publishers, 1986), 124.

2 B. Stroud, *Hume* (London: Routledge, Kegan & Paul, 1977), 223.

3 N. Kemp Smith, *The Philosophy of David Hume* (London: MacMillon and Co., 1941), 430.

4 D. G. C. MacNabb, *David Hume: His Theory of Knowledge and Morality* (Archon Books, 1966): 84.

5 For a more detailed account of Hume's associationist theory of general ideas, see my "Hume and Cognitive Science: The Current Status of the Controversy over Abstract Ideas," in *Phenomenology and the Cognitive Sciences* 4 (2005).

6 This is not the case with inferences from *constant* conjunctions. When we have observed a large and uniform sample, our confidence levels are characterized by maximal levels of insurance; as Hume puts it, they "exceed probability" and are "entirely free from doubt and uncertainty" (T 1.3.11.2; SBN 124).

7 The notion of resemblance is a notoriously difficult notion, and Hume says little in the *Treatise* to clarify his use of the term. His famous note in the Appendix to the *Treatise*, however, sheds a good deal of light on how he conceives of this notion. If X and A are complex ideas, then we can say that X partially resembles A in so far as the two ideas share "common circumstances" (T 1.1.7n; SBN 637). That is, if X contains the simple ideas d, e, and f, and B contains the simple ideas d and e, then they partially resemble one another in respect of properties d and e. Of course, the resemblance of properties d and e, in so far as they are simple ideas, cannot be further explained in this manner. As Quine points out, empiricists must commit themselves to the doctrine that human beings are born with innate quality spaces. See W. V. O. Quine, *Ontological Relativity and Other Essays* (New York: Columbia University Press, 1969), 123. Indeed, such a doctrine would seem to be presupposed by Hume's assertion that the simple ideas BLUE and GREEN are *intrinsically* more similar than the simple ideas BLUE and SCARLET (T 1.1.7n; SBN 637).

8 Ibid.

9 B. Gower, "Hume on Probability," *British Journal for the Philosophy of Science* 42 (1991): 1–19.

10 A. Mura, "Hume's Inductive Logic," *Synthese* 115 (1998): 307; cf. Gower, 15.

11 P. Maher, "Probability in Hume's Science of Man," *Hume Studies* 7 (1981): 149; cf. L. Loeb, *Stability and Justification in Hume's Treatise* (Oxford University Press, 2002), 234.

12 Admittedly, Hume states that many of our inferences from mixed frequencies are reflective and explicit in character. As he puts it, "we commonly take knowingly into consideration the contrariety of past events" and "carefully weight the experiments, which we have on each side" (T 1.3.12.7; SBN 133). But Hume immediately clarifies these statements by adding the proviso that reflection arises from habit in an "*oblique* manner" (ibid.). The rest of T 1.3.12 is dedicated to a psychological explanation of the

indirect manner by which the imagination extracts a single probability estimate from mixed frequencies. Hume's official position, then, is that many of our probabilistic inferences involve conscious awareness, but the operations of the mind that underlie them involve nothing but associative propensities of the mind. As Loeb puts it, "Hume's objective is to show that reflection and deliberation on the probability of causes is itself an associationist process" (Loeb, 230).

13 J. Noxon, *Hume's Philosophical Development: A Study of His Methods* (Oxford: Clarendon Press, 1973), 120.

14 D. Robinson, *Toward a Science of Human Nature* (New York: Columbia University Press, 1982).

15 E. Fales and E. A. Wasserman, "Causal Knowledge: What Can Psychology Teach Philosophers?" *Journal of Mind and Behavior* 13 (1992): 1.

16 G. Chapman and S. Robbins "Cue Interaction in Human Contingency Judgment," *Memory & Cognition* 18 (1990): 537.

17 D. Shanks, "Hume on the Perception of Causality," *Hume Studies* 11 (1985): 105.

18 A. Dickinson and D. Shanks, "Animal Conditioning and Human Causality Judgment," in *Perspectives on Learning and Memory*, ed. L. Nilsson and T. Archer (Hillsdale, N.J.: Erlbaum, 1985); cf. L. G. Allan, "Human Contingency Judgments: Rule Based or Associative?" *Psychological Bulletin* 114 (1993): 440.

19 D. Shanks, *The Psychology of Associative Learning* (Cambridge: Cambridge University Press, 1995): 31–3.

20 D. S. Blough, "Steady State Data and a Quantitative Model of Operant Generalization and Discrimination," *Journal of Experimental Psychology: Animal Behavior Processes* 1 (1975): 3–21.; R. A. Rescorla and D. R. Furrow, "Stimulus Similarity as a Determinant of Pavlovian Conditioning," *Journal of Experimental Psychology: Animal Behavior Processes* 3 (1977): 212. *Technical note:* Blough employs a "common element" notion of similarity in his experiments. He shows that the degree of association between events is a function of their common elements. Rescorla and Furrow demonstrate that when paired events are qualitatively similar the association between them is "substantially superior." In their first two experiments, Rescorla and Furrow show that this phenomenon holds when the stimulus similarity involves a *common* modality; when auditory cues are paired with auditory cues, for example, the association is stronger than when auditory cues are paired with visual cues. In their third experiment, they demonstrate that stimulus similarity *within* a modality (e.g., the dimensions COLOR and SHAPE) also increases the strength of the association.

21 R. A. Rescorla and A. R. Wagner, "A Theory of Pavlovian Conditioning: Variations in the Effectiveness of Reinforcement and Non-Reinforcement," in *Classical Conditioning II: Current Research and Theory*, ed. A. H. Black and W. F. Prokasky (New York: Appleton-Century-Crofts, 1972).

22 Allan, 439.

23 R. Miller et al. "Assessment of the Rescorla-Wagner Model," *Psychological Bulletin* 117 (1995): 381.

24 Ibid., 365.

25 G. Chapman and S. Robbins, "Cue Interaction in Human Contingency Judgment," *Memory & Cognition* 18 (1990): 545.

26 M. Gluck and G. Bower, "From Conditioning to Category Learning: An Adaptive Network Model," *Journal of Experimental Psychology: General* 117 (1988): 241.

27 Shanks, 114–5.

28 R. S. Sutton and A. G. Barto, "Toward a Modern Theory of Adaptive Networks: Expectation and Prediction," *Psychological Review* 88 (1981): 155–6.

29 Shanks, 110.

30 J. L. Elman, "Finding Structure in Time," *Cognitive Science* 14 (1990): 189.

31 D. Servan-Schreiber, A. Cleeremans, and J. L. McClelland, "Graded State Machines: the Representation of Temporal Contingencies in Simple Recurrent Networks," *Machine Learning* 7 (1991): 181.

32 For a discussion of superposition in connectionist networks, see my "Filling the Gaps: Hume and Connectionism on the Continued Existence of Unperceived Objects," *Hume Studies* 25 (1999): 161.

33 I. Gormezano and E. Kehoe, "Classical Conditioning and the Law of Contiguity," in *Predictability, Correlation, and Contiguity*, ed. P. Harzem and M. D. Zeiler (New York: John Wiley & Sons 1981), 38.

34 Fales and Wasserman, 8.

35 P. W. Cheng et al., "A Causal-Power Theory of Focal Sets," in *Causal Learning: Advances in Research and Theory*, ed. D. R. Shanks, D. L. Medin, and K. J. Holyoak (San Diego: Academic Press, 1996), 314.

36 I. Hacking, "Hume's Species of Probability," *Philosophical Studies* 33 (1978): 23.