THE QUARTERLY JOURNAL OF EXPERIMENTAL PSYCHOLOGY, 2002, *55B* (4), 289–310

A review of recent developments in research and theories on human contingency learning

Jan De Houwer
University of Ghent, Ghent, Belgium
Tom Beckers
University of Leuven, Leuven, Belgium

Over the past 20 years, human contingency learning has resurfaced as an important topic within experimental psychology. This renewed interest was sparked mainly by the proposal that associative models of Pavlovian conditioning might also apply to human contingency learning—a proposal that has led to many new empirical findings and theoretical developments. We provide a brief review of these recent developments and try to point to issues that need to be addressed in future research.

Our ability to detect contingencies between events is crucial for our functioning and survival because it allows us to predict and control events in the environment. But how do we learn that one event predicts or causes the presence or absence of another event? This question has occupied the minds of philosophers and psychologists for many centuries (e.g., Cheng, 1997; Hume, 1739/1987; Kant, 1781/1965; Michotte, 1954; Shanks, Holyoak, & Medin, 1996). In most modern studies on human contingency learning (HCL), participants receive information about a number of situations in which certain cues and certain outcomes are either present or absent, and they are asked to judge the extent to which the presence of a cue is related to the presence of the outcome. For instance, on each trial, participants might see a list of foods that a fictitious patient has eaten (i.e., the cues) together with information about whether the patient showed an allergic reaction after eating those foods (i.e., the outcome). On the basis of this information, participants judge how likely it is that the patient will have an allergic reaction after eating a particular food or combination of foods.

At a procedural level, studies of HCL are very similar to those of Pavlovian conditioning: Stimuli (cues and outcomes) are paired in a certain way, and the resulting changes in the responses to the stimuli (contingency judgements) are assessed. In 1984, Dickinson, Shanks, and Evenden proposed that HCL and Pavlovian conditioning might be similar not only at the

Requests for reprints should be sent to Jan De Houwer, Department of Psychology, University of Ghent, Henri Dunantlaan 2, B-9000 Ghent, Belgium. Email: Jan.DeHouwer@rug.ac.be

Tom Beckers is a Research Assistant for the Fund for Scientific Research (FWO, Flanders, Belgium). We thank Steven Glautier and Peter White for their helpful comments on previous drafts of this paper.

© 2002 The Experimental Psychology Society
http://www.tandf.co.uk/journals/pp/02724995.html DOI:10.1080/02724990244000034

procedural level, but also with respect to the effects that can be observed and the processes that underlie these effects. They argued that phenomena that have been observed in Pavlovian conditioning might also occur in HCL, and that models of the former might also be valid for the latter form of learning. This proposal has led to a proliferation of research and theories concerning HCL.

In the present paper, we present a brief review of these recent developments. We focus on associative models of Pavlovian conditioning that have been applied to HCL, on the main non-associative models of HCL that were contrasted with those associative models, and the studies that highlight the strengths and weaknesses of both types of model. Although this encompasses a large part of the recent literature on HCL, our paper provides only a selective review of modern HCL research. For instance, we largely ignore (1) studies in which information about the relation between stimuli is summarized in the form of tables rather than presented sequentially (i.e., on a trial-by-trial basis), (2) studies in which participants learn about the relation between responses that they are free to perform and the outcome of those responses, and (3) models of causal reasoning and causal judgement that are mainly based on the results of such studies and that have not been contrasted directly with associative models (e.g., Anderson & Sheu, 1995; White, 2002). Furthermore, our paper is primarily about human contingency learning (i.e., learning about the relation between the presence of two stimuli) rather than about causal learning (i.e., learning about whether one stimulus produces another stimulus). But because causal learning can be regarded as one particular form of contingency learning, we discuss some models of, and studies about, causal learning, in particular those that inform us about the strengths and weaknesses of general models of contingency learning.

Despite the inevitable limitations in the scope of this review, we bring together a large variety of findings and theories. By doing so, we hope to provide a useful introduction to modern HCL research and to help identify some of the issues that need to be addressed in future research. In the first part of the paper, we focus on the dominant models of HCL and the research that was inspired by these models. The second part is devoted to a number of findings that do not fit with any of the dominant models but that can, in most cases, be explained within alternative theoretical frameworks.

FORWARD BLOCKING AND THE RESCORLA–WAGNER MODEL

At the time that Dickinson and colleagues (1984) published their seminal work, most theories of HCL postulated that judgements about the contingency between a target cue T and an outcome do not merely depend on the number of times that the target cue and the outcome have co-occurred, but also on information about the number of times that only the cue, only the outcome, or neither was present in a certain situation (see Shanks, 1995, for a review). For instance, according to the well-known ΔP model, judgements about the contingency between a target cue T and an outcome O reflect the probability with which the outcome was present when the target cue was present, $P(O \mid T)$, minus the probability of the outcome when the cue was not present, $P(O \mid \sim T)$. Models such as the ΔP model thus postulate that contingency judgements are a fairly accurate reflection of the actual statistical contingencies between events.

This emphasis on the role of statistical contingency was not new. During the 1960s, scholars of Pavlovian conditioning had already put forward the idea that the strength of conditioned responses is a function of the statistical contingency between the conditioned stimulus (CS) and the unconditioned stimulus (US; e.g., Rescorla, 1968). However, they quickly discovered that the statistical contingency between CS and US is not the only thing that matters. Most important, Kamin (1969) demonstrated that conditioned responses toward a target CS, T, also depend on the statistical contingency between other CSs and the US. In his experiments, a target CS (e.g., a light) was presented with an alternative CS (e.g., a tone), and this compound was followed by the US (we refer to such trials as AT+ trials, where T stands for the target CS, A for the alternative CS, and + for the presence of the US). The crucial result was that on test the target CS, T, evoked a weaker conditioned response when the AT+ trials were preceded by trials on which the alternative CS, A, was paired individually with the US (A+) than when A was not trained on its own (no A+ trials).

This effect, which became known as *blocking*, led to the development of many new models of Pavlovian conditioning, the most influential being the Rescorla–Wagner model (Rescorla & Wagner, 1972). According to this model, associations are formed between the representations of CSs and USs that are presented together. The strength of the conditioned response that is evoked by a CS will depend on the strength of the association between the representations of the CS and the US. Importantly, each time a CS is presented, the strength of the CS–US association is updated according to the following learning rule:

$$\Delta V_n = \alpha\beta(\lambda - \Sigma V_{n-1}) \quad (1)$$

According to this rule, the extent to which the associative strength of a CS will change on a given trial (ΔV_n) depends not only on the associative strength that the cue had acquired previously, but also on the existing associative strengths of other cues that are present on the trial (ΣV_{n-1}). It is the latter assumption that allows the Rescorla–Wagner model to explain blocking and related cue competition effects. During A+ trials that precede AT+ trials, A will gain associative strength up to some maximum value (λ) that depends on the intensity of the US. Therefore, when A is afterwards presented together with T, the combined associative strength of A and T (ΣV_{n-1}) will equal λ (λ has a fixed positive value when the US is present but is zero when the US is absent), and thus ΔV_n will be zero. As a result, T will gain little if any associative strength and will not evoke a (strong) conditioned response when presented on its own. Put in psychological rather than mathematical terms, the Rescorla–Wagner model postulates that learning (i.e., the change in associative strength) depends on the extent to which the presence or absence of the US is expected (i.e., λ–ΣV_{n-1}). Because participants learn on the A+ trials that A predicts the US, they are not surprised by the presence of the US on the AT+ trials and as a result learn little about the relation between T and the US.

Dickinson et al. (1984) pointed out that the Rescorla–Wagner model can easily be applied to HCL. One only needs to assume that associations are formed between the representations of cues and outcomes, that the strength of these associations is updated according to the Rescorla–Wagner learning rule, and that judgements about the contingency between a target cue and an outcome are a reflection of the strength of the association that links the representations of that cue and outcome. They then reasoned that if the Rescorla–Wagner model is a valid model of HCL, blocking should also occur in HCL. Dickinson et al. (1984, Experiment 2) indeed found that contingency judgements for a target cue T were lower when AT+ trials

were preceded by A+ trials than when no A+ trials were presented. This result was crucial because it could not be explained on the basis of the existing models of HCL, whereas it was fully anticipated by the Rescorla–Wagner model. It therefore lent substantial support to the claim that the Rescorla–Wagner model is a valid model of HCL.

BACKWARD BLOCKING AND THE PROBABILISTIC CONTRAST MODEL

Soon after the influential studies of Dickinson et al. (1984), Shanks (1985) tested another prediction of the Rescorla–Wagner model. According to the model, blocking should occur if the A+ trials are presented before the AT+ trials, but not if they are presented after the AT+ trials. In the latter case, A does not have associative strength prior to the AT+ trials, and T can, therefore, acquire associative strength. Moreover, the subsequent A+ trials should have no effect on the associative strength of T. This follows from the assumption that the associative strength of a cue will change only on trials where that cue is present. Because T is not present on the A+ trials, the A+ trials should have no effect on the associative strength of T. Therefore, contingency judgements for T should be the same regardless of whether AT+ trials are followed by A+ trials or not. Shanks (1985), however, showed that contingency judgements for T were lower when A+ trials were presented after the AT+ trials than when no A+ were presented. This result, known as *backward blocking*, raised serious questions about the validity of the Rescorla–Wagner model as a model of HCL. Similar "retrospective revaluation" effects have meanwhile been found in many other studies (e.g., Chapman, 1991; De Houwer, Beckers, & Glautier, 2002; Dickinson & Burke, 1996; Larkin, Aitken, & Dickinson, 1998; Wasserman & Berglan, 1998; Williams, Sagness, & McPhee, 1994).

Perhaps encouraged by this failure of the Rescorla–Wagner model, attention was redirected to the kind of probabilistic models that dominated HCL research before the publication of Dickinson et al.'s (1984) work. Most prominently, Cheng and colleagues (e.g., Cheng & Holyoak, 1995; Cheng & Novick, 1990, 1992; also see Waldmann & Holyoak, 1992) developed the probabilistic contrast model. This model is similar to the ΔP model in that it states that contingency judgements will reflect the objective contingency between the presence of a target cue and the presence of the outcome. The main difference is that the objective contingency is calculated on the basis of a focal set of situations in which the presence of other cues is kept constant. For instance, when the target cue T is always presented in compound with an alternative cue A, contingency judgements for T will reflect the difference between the probability of the outcome when both A and T are present, $P(O \mid A.T)$, and the probability of the outcome when only A is present, $P(O \mid A.\sim T)$. The logic underlying the model is straightforward: In order to make valid inferences about the relation between T and the outcome, one needs to compare situations that only differ with regard to the presence of T (e.g., both A and T are present versus A but not T is present) rather than situations that differ with regard to the presence of several cues (e.g., both A and T are present versus no cue is present).

The probabilistic contrast model provides a straightforward explanation for both forward and backward blocking effects. Regardless of whether the A+ trials are presented before or after the AT+ trials, they will result in an increase of the probability that the outcome occurs in the presence of A and the absence of T and thus in a decrease of the difference between this probability and the probability of the outcome in the presence of both A and T. Because

contingency judgements for T will reflect the latter difference, $P(O|A.T) - P(O|A.\sim T)$, judgements for T will be lower when A+ trials precede or follow AT+ trials than when no A+ trials are presented.

REVISIONS OF ASSOCIATIVE MODELS

Demonstrations of backward blocking and related retrospective revaluation effects did not, however, mark the end for associative models of HCL. Just as earlier probabilistic models had been modified in order to account for (forward and backward) blocking, researchers started to examine whether associative models could be adapted to accommodate backward blocking and other forms of retrospective revaluation effects. Van Hamme and Wasserman (1994; also see Markman, 1989) pointed out that the Rescorla–Wagner model had to be revised only slightly in order to render it compatible with the data on retrospective revaluation. In the original model, the associative strength of a cue could change only on trials where that cue was actually present. Van Hamme and Wasserman proposed a revision of the model in which the associative strength of a target cue could also change on trials where the cue was absent but expected. The presence of a target cue can be expected either because of instructions or because another cue is present that was previously associated with the target cue. For instance, if AT+ trials are followed by A+ trials, T is expected on the A+ trials by virtue of an association formed between the representations of T and A during the AT+ trials. Van Hamme and Wasserman proposed that the direction of the change in associative strength of expected but absent cues is opposite to the direction of the change in associative strength of cues that are actually presented. They modelled this by giving the α parameter in Equation 1 a positive value for a cue that is present and a negative value for a cue that is expected but absent.

This minor modification allows the Rescorla–Wagner model to explain backward blocking. As mentioned earlier, T is expected but absent on the A+ trials that follow AT+ trials and will thus have a negative α value on these trials. Because the presence of the outcome is not fully expected on the A+ trials (A shared the available associative strength with T on the AT+ trials), $\lambda - \Sigma V_{n-1}$ in Equation 1 will be positive on the initial A+ trials. Consequently, the change in the associative strength of T (ΔV_n), calculated by multiplying α (which has a negative sign) with $(\lambda - \Sigma V_{n-1})$ (which has a positive sign), will be negative, and the associative strength of T will decrease on the A+ trials. Therefore, contingency judgements for T will be reduced as a result of a backward blocking procedure.

Dickinson and Burke (1996) proposed an alternative associative model that can also account for both forward and backward blocking effects. Their model is a revision of the Standard Operating Procedures (SOP) model that was first proposed by Wagner (1981). In the revised (and original) SOP model, stimuli are represented as nodes that consist of several elements. Each of the elements in a node can be in one of three states: an inactive state (I), a primary active state (A1), or a secondary active state (A2). The elements of the representation of a stimulus, X, can switch from an inactive state to an A1 state only when Stimulus X is presented. Those same elements can switch from the inactive state to an A2 state when another stimulus, Y, is presented that has previously been associated with the Stimulus X. One could thus say that the A1 state models the actual presence of a stimulus whereas the A2 state models the extent to which a stimulus is expected to be present.

The extent to which Stimulus Y can activate the representation of Stimulus X (i.e., can lead to an expectation of Stimulus X) depends on the strength of the excitatory association between the representations of both stimuli relative to the strength of the inhibitory association between those representations. In the revised SOP model of Dickinson and Burke (1996), the excitatory association between two representations will increase in strength as a function of the number of elements of both stimulus representations that are in the same active state (A1–A1 or A2–A2). Inhibitory associations will increase in strength as a function of the number of elements of the two stimulus representations that are in different active states (A1–A2 or A2–A1). Contingency judgements are assumed to be a function of the difference between the strength of the excitatory association and the strength of the inhibitory association that link the representation of the cue and the outcome.

Within the revised SOP model, backward blocking is explained in the following way: Because A and T have been presented together on the AT+ trials, there will be a strong excitatory association between the representations of A and T. Because of this association, elements of the representation of T will be activated to the A2 state on the A+ trials (T is expected but absent), whereas the elements of the outcome representation will be in the A1 state (the outcome is present). Therefore, the strength of the inhibitory association between the representations of T and the outcome will increase on A+ trials that follow AT+ trials. Hence, contingency judgements for T will be lower when AT+ trials are followed by A+ trials than when no A+ trials are presented. Some studies suggest that the revised SOP model of Dickinson and Burke (1996) is better able to account for the fine-grained characteristics of retrospective revaluation effects than is the revised version of the Rescorla–Wagner model (Larkin et al., 1998; but see Le Pelley & McLaren, 2001).

Both the revised Rescorla–Wagner model and the revised SOP model make a straightforward prediction about the conditions under which retrospective revaluation effects can be observed. In both models, the associative strength of an absent cue can change only when the absent cue was expected to be present. In the absence of instructions about the presence or absence of a target cue, its presence on a trial is only expected when it was previously associated with a cue that is presented on that trial. Therefore, manipulations that influence the strength of the association between the target cue and alternative cues should influence the magnitude of retrospective revaluation effects. This prediction was verified by Dickinson and Burke (1996; also see Aitken, Larkin, & Dickinson, 2001; Larkin et al., 1998; Wasserman & Berglan, 1998). Dickinson and Burke (1996) conducted a study that consisted of two conditions. In one condition, a target cue T was presented together with three different alternative cues (inconsistent group; e.g., A1T+, A2T+, A3T+), whereas in the other condition, Cue T was consistently paired with only one alternative cue (consistent group; e.g., three A1T+ trials). Subsequent information about the probability of the outcome when only A1, A2, or A3 were presented alone influenced judgements for T only in the consistent group. Importantly, when the A trials were presented before the AT trials (forward cue competition), A trials had an impact on judgements for T regardless of whether T was paired with just one A cue or with all three A cues. Thus, prior information about the contingency between A and the outcome influences judgements for T regardless of whether T is consistently associated with A, but subsequent information about the relation between A and the outcome only influences judgements for T when A and T have repeatedly co-occurred. This is precisely the pattern of results predicted by the revised Rescorla–Wagner and SOP models. Forward cue competition

depends only on the associative strength of the A cue prior to the AT+ trials: Cue T is prevented from gaining associative strength on AT+ trials regardless of whether it is always paired with the same A stimulus or with different A stimuli on different trials. But T loses associative strength on A+ trials that follow AT+ trials only if it was repeatedly paired with that particular A stimulus on earlier trials.

These results cannot be accommodated by any of the existing non-associative models. It is difficult to envisage how, for instance, the probabilistic contrast model would explain the fact that retrospective cue competition depends on the extent to which the target cue and alternative cues have been associated. And even if one could find a way to reconcile the probabilistic contrast model with the observation that cue–cue associations have an effect on retrospective revaluation, the model would still need to explain why these associations do not have a similar impact on forward cue competition effects. The crucial point here is that the probabilistic contrast model provides the same explanation for forward and retrospective cue competition effects, as presenting A+ trials leads to an increase in P(O | A.T) and thus in a decrease of the difference between P(O | A.T) and P(O | A~T) whereas the revised associative models posit different mechanisms for forward and backward cue competition.

Researchers have also identified other characteristics of HCL that can only be explained by associative models. First, several studies have shown that the order in which trials are presented has a marked impact on contingency judgements. The fact that the strength of the association between the alternative and the target cue has a different impact on forward than on retrospective cue competition (e.g., Dickinson & Burke, 1996) already demonstrates the importance of the order in which trials are presented. Other effects of trial order have also been observed. For instance, Lopez, Shanks, Almaraz, and Fernandez (1998; see also Shanks, Lopez, Darby, & Dickinson, 1996) presented AT+ and A– trials followed or preceded by BT– and B+ trials. Participants judged the relation between T and the outcome to be stronger when the AT+ and A– trials came last than when they came first, even though participants could remember the earlier trials as well as the later trials. These results thus suggest that more recent information has a stronger impact on contingency judgements than does information that was presented earlier.

Both the Rescorla–Wagner and the SOP models provide a straightforward explanation for such recency effects. Both models imply that learning is expectancy driven: More will be learned on a trial (i.e., associative strength will change more) when the presence or absence of the outcome is unexpected than when its presence or absence is expected. When, as in the study of Lopez et al. (1998), AT+ and A– trials are followed by BT– and B+ trials, the absence of the outcome on the BT– trials is particularly surprising because T was a good predictor of the outcome during the AT+ and A– trials. (In the Rescorla–Wagner model, this surprise is modelled by a large difference between λ and ΣV_{n-1}; in the SOP model, it is modelled by the fact that many elements of the outcome representation will be in the A2 state). Therefore, more will be learned on the BT– trials (i.e., the associative strengths will decrease more) when these trials are presented after the AT+ and A– trials than when they are presented before the AT+ and A– trials. This illustrates that when participants receive conflicting information about the contingency between a cue and an outcome, the information provided most recently will have the biggest impact. It is markedly more difficult to explain these trial order effects on the basis of non-associative models such as the ΔP model and the more recent probabilistic contrast model (e.g., Cheng & Novick, 1992). According to these models, contingency

judgements reflect the outcome of probabilistic contrasts, for example, P(O | T) – P(O | ~T). Because probabilities are unaffected by the order in which events are presented, these non-associative models predict that trial order should have no effect on contingency judgements.[1]

Not only the order, but also the number of trials can have an impact on HCL. Several studies showed that contingency judgements become more accurate as the number of trials concerning a particular cue–outcome relation increases (e.g., Lopez, Almaraz, Fernandez, & Shanks, 1999; Shanks, 1987; Van Overwalle & Van Rooy, 2001). Associative models are compatible with both the existence and the shape of such learning curves. For instance, when a cue is repeatedly paired with an outcome (i.e., there is a strong positive contingency between the two), associative strength will gradually increase over trials until it reaches an asymptote that reflects the actual contingency between the cue and outcome. Moreover, increases in associative strength will be bigger on the initial cue–outcome trials because on these trials, the presentation of the cue does not elicit much of an expectation of the outcome, so that the subsequent presentation of the outcome is highly surprising. Non-associative models such as the probabilistic contrast model are not able to provide a straightforward account of learning curves because probabilities are not affected by the amount of evidence on which they are based.

However, not all results regarding the effects of cue–cue associations, trial order, and number of trials are consistent with the revised Rescorla–Wagner and SOP models. First retrospective revaluation effects can be observed even when the alternative cue is not directly associated with the target cue. For instance, De Houwer and Beckers (2002a, b) presented AT1+ and T1T2+ trials that were followed by either A+ trials (A and outcome present) or A– trials (only A present). Whether A on its own was paired with the outcome not only influenced judgements for T1 (with which A was associated directly) but also judgements for T2. The latter result is crucial because T2 was associated only indirectly with A (through a common association with T1). Therefore, participants had no reason to expect the presence of T2 on the A+ trials, and the associative strength of T2 should therefore not have changed. Other studies have demonstrated that alternative cues need not even be associated indirectly with the target cue. Matute and colleagues (Matute & Pine o, 1988a, b; Pine o, Ortega, & Matute, 2000) showed that responding to a target cue T is reduced when an alternative cue is paired with the same outcome (A+) after the T+ trials. They pointed out that such results are similar to retrospective interference effects that have been reported in the paired-associative learning literature (i.e., impaired recollection of an association A–B when A–B trials are followed by C–B trials).

Second, whereas the results of Lopez et al. (1998) showed that more recent trials have a bigger impact on judgements than do earlier trials, Denis and Ahn (2001) recently observed the exact opposite pattern of results; that is, a stronger influence of initial than of later trials. They argued that, under certain conditions, participants form a hypothesis about the relation between the cue and the outcome on the basis of the initial trials. Due to a confirmation bias,

[1]One could, however, argue that (1) trial order determines the probability that a certain trial will be taken into account for calculating the relevant probabilistic contrast (but see Shanks, Lopez, et al., 1996), (2) participants notice that contingencies have changed and thus intentionally disregard earlier trials (e.g., Waldmann, 2000), and (3) surprising events are more likely to be taken into account for calculating the relevant probabilistic contrast.

they subsequently pay more attention to trials that confirm this hypothesis than to trials that disconfirm the hypothesis. Hence, initial trials will have more impact on judgements than will subsequent trials. Although the reasons for the discrepancy between the results of Dennis and Ahn (2001) and the results of Lopez et al. (1998) are still unclear, the fact that earlier trials can exert a bigger impact on judgements than can more recent trials is problematic for the (revised) Rescorla–Wagner and SOP models. Finally, although learning curves have been found in some HCL studies, such learning curves are not always observed (e.g., Catena, Maldonado, & Cándido, 1998; Waldmann, 2000).

To summarize, despite the fact that initial findings confirmed certain central predictions of the revised Rescorla–Wagner and SOP models, more recent findings raise doubts about the generality of those initial findings and thus about the validity of the revised models. Moreover, researchers have also created new non-associative models. The studies that were inspired by these non-associative models further call into question the validity of the revised Rescorla–Wagner and SOP models. These developments are summarized in the next section.

THE FURTHER DEVELOPMENT OF PROBABILISTIC MODELS

The Power PC model of causal learning

Probabilistic models such as the ΔP model and the probabilistic contrast model (e.g., Cheng & Novick, 1992) are in essence normative models that were developed on the basis of a normative analysis of what the optimal contingency judgement would be in a given situation (see Shanks, 1995, Chapter 1).[2] Cheng and Novick (1992), for instance, argued that when a target cue is always presented together with another, alternative cue, normatively it would be optimal to give a judgement for T that corresponds to the outcome of a conditional probabilistic contrast, for example, $P(O|A.T) - P(O|A.\sim T)$. As we pointed out earlier, the logic behind this proposal is

[2]Shanks (1995) argued that the probabilistic contrast model provides a different kind of explanation from that provided by associative models such as the Rescorla–Wagner model. Both types of model can be used to map an input (i.e., events) to an output (i.e., contingency judgements), but only the latter make assumptions about the processes and representations that are involved in the transformation of input to output. The probabilistic contrast model says nothing about how events are coded in memory or what processes people use to arrive at judgements. It merely states that the judgements will reflect the outcome of an appropriate probabilistic contrast. In fact, one can mathematically prove that under certain conditions (e.g., assuming certain parameters and after a sufficient number of trials), the (revised) Rescorla–Wagner model actually generates contingency judgements that reflect the outcome of an appropriate probabilistic contrast (e.g., Cheng, 1997). One can therefore argue that the Rescorla-Wagner model is one possible process or algorithmic implementation of the normative theory that contingency judgements reflect the outcome of appropriate probabilistic contrasts (see Marr, 1982). However, within Aristotle's framework of types of causal explanation (see Killeen, 2001), both the probabilistic contrast model and the associative models are formal models that can be used to predict an output (behaviour) on the basis of a certain input. Within this framework, the only thing that matters is whether models make different predictions regarding the input–output relation. Researchers have further developed the normative theory by trying to delineate the conditions under which judgements that are based on the outcome of probabilistic contrasts are normatively optimal (e.g., Cheng, 1997), and some have proposed probabilistic models that do incorporate assumptions about how input is mapped onto output (e.g., Waldmann, 2000). These new models do yield different predictions from those of the Rescorla–Wagner model and should therefore be contrasted with the Rescorla–Wagner model.

that inferences about the relation between T and the outcome should be based on a comparison of the probability of the outcome in situations that differ only with regard to the presence of T. If the probability of the outcome in such situations is the same regardless of whether T is present or absent, one can conclude that T is not related to the presence of the outcome.

Cheng (1997; Cheng & Holyoak, 1995), however, pointed out that this conclusion is warranted only if the outcome does not always occur to a full extent within the focal set of situations. Assume, for instance, that a physician wants to test whether a person is allergic to a certain substance T. He or she therefore puts scratches on the patient's arm, some that contain the substance, others that do not. Unexpectedly, hives break out at each spot, regardless of whether the scratch contained the substance, $P(O|A.T) - P(O|A.\sim T) = 1$. If we regard scratching the skin as an alternative cue, A, the probabilistic contrast for T equals zero, $P(O|A.T) - P(O|A.\sim T) = 0$, which should thus result in a low contingency judgement for T. However, it is also possible that T does cause an allergic reaction but that the effect of T is not observable because the alternative cause (i.e., scratching) on its own always produces an allergic reaction. This example from Cheng (1997) illustrates that the outcome of an appropriate probabilistic contrasts, for example, $P(O|A.T) - P(O|A.\sim T)$, does not always provide a normatively correct basis for contingency judgements.

Whereas Cheng and Holyoak (1995) simply stated that judgements will not reflect the outcome of appropriate contrasts when ceiling effects might apply, Cheng (1997) proposed the Power PC model of causal learning, which incorporates this auxiliary assumption. At the formal level (see Cheng, 1997, for the normative analysis that underlies the formal model), judgements about the extent to which a cue, T, produces an outcome will reflect the outcome of the following equation:

$$P = \frac{[P(O|A.T) - P(O|A.\sim T)]}{[1 - P(O|A.\sim T)]} \qquad 2$$

This equation expresses the idea that causal judgements will be a function not only of the outcome of an appropriate probabilistic contrast (as is assumed in the probabilistic contrast model), that is, $P(O|A.T) - P(O|A.\sim T)$, but also of the base rate of the outcome within the focal set of situations, that is, $P(O|A.\sim T)$. This base rate reflects the probability of the outcome in situations that are identical to situations in which T is present, except for the fact that T is not present. If the outcome always occurs even when T is absent, that is, base rate = $P(O|A.\sim T) = 1$, the denominator of Equation 2 will be zero, and the equation will thus be undefined. As a result, participants will refrain from interpreting the appropriate probabilistic contrast. When the outcome never occurs in the absence of T, that is, $P(O|A.\sim T) = 0$, the Power PC model reduces to the standard probabilistic contrast model (because the denominator in Equation 2 will be equal to 1).[3] As such, the Power PC model can account for the same

[3] At intermediate levels of base rate, causal judgements will be higher than when only based on the outcome of the probabilistic contrast model and more so the higher the base rate of the outcome. This is because in such cases the denominator has a positive value smaller than one, which implies that the nominator will be multiplied by a value larger than one. Stated differently, the Power PC model implies that increases in the base rate of the outcome, $P(O|A.\sim T)$, will result in increases in the weight that is given to the outcome of the probabilistic contrast, $P(O|A.T) - P(O|A.\sim T)$, apart from when the base rate is 1, in which case there is uncertainty about the causal status of the target cue.

phenomena as the probabilistic contrast model, but in addition predicts ceiling and base-rate effects.

If ceiling effects are important, one would expect that blocking will be observed only when the outcome does not always occur to a maximal extent on the A+ and AT+ trials. Recent studies seem to support this hypothesis. For instance, De Houwer et al. (2002; see Lovibond, Been, Mitchell, Bouton, & Frohardt, 2001, for related results) found weaker (forward and backward) blocking when the intensity of the outcome was always maximal on the A+ and AT+ trials (i.e., an explosion with an intensity of 10 on a scale from 0 to 10) than when the intensity of the outcome on those trials was below maximum (i.e., an explosion with an intensity of 10 on a scale from 0 to 20). When the intensity of the outcome is always maximal, participants cannot infer whether T has an additional effect on the outcome, and hence they should be uncertain about their judgement for T. However, when the outcome has an intensity below the maximum on both the A+ trials and the AT+ trials, participants can infer with greater certainty that T does not have an effect on the outcome. This result is important not only because it supports the idea that ceiling effects play a role in HCL but also because such effects cannot be explained on the basis of the (revised) Rescorla–Wagner and SOP models. One should note, however, that the results of De Houwer et al. (2002) do not directly support the Power PC model because the model is not developed to deal with outcomes that vary in intensity. It remains to be tested whether ceiling effects also influence blocking effects when the *probability* rather than the *intensity* of the outcome is manipulated. Moreover, several studies have shown that the Power PC model does not provide an accurate account for effects of the base rate of the outcome (e.g., Lober & Shanks, 2000; Perales & Shanks, 2001). Although the notion that ceiling effects are important in HCL seems to be valid, questions can thus be raised about the implementation of this in the Power PC model.

Causal model theory

Not long after Cheng and Novick (1992) presented the probabilistic contrast model, Waldmann and Holyoak (1992) proposed a theory that also postulates that judgements reflect the outcome of appropriate probabilistic contrasts. However, their theory emphasizes not so much the fact that judgements that reflect the outcome of such contrasts are (under certain conditions) normatively accurate, but the fact that people know when and why probabilistic contrasts provide a good basis for contingency judgements and that they act accordingly.

Waldmann (2000, 2001; Waldmann & Holyoak, 1992) pointed out that the same set of cue–outcome events can be interpreted in many ways. He proposed that contingency judgements depend on the causal model that participants adopt: that is, on a set of assumptions about the nature of the cues and outcomes and the nature of the relation between cues and outcomes. In some situations, participants will adopt a *common-cause model*; that is, they will regard the different cues as different potential causes of the outcome. For instance, it is likely that participants adopt a common-cause model if instructions state that the cues are substances in the blood that can cause a disease (i.e., the outcome). When several potential causes of an outcome are present (as is the case on AT+ trials), it is not clear which of the causes actually produced the presence of the outcome. If, however, additional trials are presented in which A on its own produces the outcome to the same extent as A and T together, that is,

$P(O \mid A.T) = P(O \mid A.\sim T)$, participants can infer that T is not a cause of the outcome (except for when the outcome always occurs to a maximal extent, see earlier) and thus that the outcome will not be present when only T is present.

However, this inference is valid only when participants adopt a common-cause model. They can, for instance, also adopt a *common-effect model*; that is, they can assume that cues A and T are effects and that the outcome is the potential cause of those effects. For instance, a common-effect model is appropriate when A and T are described as substances in the blood that might be caused by the disease (i.e., the outcome). When participants see trials on which both Effect A and Effect T are present when the cause (i.e., outcome) is present (AT+ trials), they can infer that T is an effect of the cause (and thus that the presence of T indicates the presence of the cause) regardless of whether there are other situations in which only Effect A and the cause are present (A+ trials). These latter situations do not change the fact that the presence of T can be used to predict the presence of the outcome (see Waldmann, 2000; Waldmann & Holyoak, 1992, for a more detailed discussion).

For these reasons, causal model theory predicts that blocking will occur when participants adopt a common-cause model but not when they adopt a common-effect model. Although there has been a debate about the conclusiveness of the initial findings of Waldmann and Holyoak (1992; see Waldmann, 2000, for a review), Waldmann (2000) recently provided strong evidence that cue competition is indeed more likely to occur when instructions encourage participants to adopt a common-cause model (e.g., when cues are described as different substances in the blood, and the outcome is said to be a disease that can be produced by certain substances in the blood) than when instructions favour a common-effect model (e.g., when the outcome is described as a substance in the blood and the cues as different diseases that can be produced by that substance). These results are important not only because they demonstrate that people possess and can use abstract knowledge about causal relations (e.g., the knowledge that the effects of different causes should add up) but also because it is very difficult to see how any associative model could explain the fact that instructions about the nature of the cues and outcomes could have such a dramatic impact on contingency judgements (see De Houwer et al., 2002, for additional evidence).

Although causal model theory is not very specific with regard to the exact representations and processes that underlie HCL, it does imply that the beliefs and assumptions of the participants play a crucial role. De Houwer (2002; De Houwer & Beckers, 2002a) recently made a similar proposal. He suggested that contingency judgements might not reflect the objective contingency as determined by the actual events that have been presented, but the contingency as determined by the way participants encode and recode events. This hypothesis is in line with two recent findings. De Houwer presented A+ trials followed by AT+ trials, but prevented participants from seeing whether Cue T was present during the A+ trials (the position where T could appear was hidden behind a screen that was removed at the start of the AT+ trials). Half of the participants were told after the A+ and AT+ trials that T had been present during the A+ trials, whereas the other participants were told that T had not been present during those trials. Only the second group showed a blocking effect. This result can be explained in the following way. When participants are told that cue T was absent on the A+ trials, they can infer that the outcome was as likely and intense on the A+ trials as on the AT+ trials, and thus that T is not a cause of the outcome. But when they are told that T was present on the A+ trials, they can retrospectively recode the A+ trials as AT+ trials. As a result, they have no

information about the probability of the outcome when only A is present and can therefore not make an inference about the causal status of cue T.

The results of De Houwer and Beckers (2002a, b) can be explained in a similar way. As we mentioned earlier, they presented AT1+ and T1T2+ trials that were followed by either A– or A+ trials. Judgements for T2 were higher after A+ trials than after A– trials. When A– trials are presented, participants can infer on the basis of a probabilistic contrast, $P(O|A.T) - P(O|A.\sim T) = 1$, that T1 produced the outcome on the AT1+ trials and thus that these trials can be recoded as T1+ trials. As a result, the appropriate probabilistic contrast for T2 can be used to infer that T2 is not a generative cause of the outcome, that is, $P(O|T.T2) - P(O|T1.\sim T2) = 0$. When A+ trials are presented, however, these trials can be used to infer that TI is not a cause of the outcome, $P(O|A.T1) - P(O|A.\sim T1) = 0$, and hence that T2 was responsible for the outcome on the T1T2+ trials.

To summarize this section, new probabilistic models have been formulated, which have led to new findings concerning the role of ceiling effects, causal models, and active recoding of events. Whereas these findings are compatible with the idea that judgements reflect the outcome of probabilistic contrasts, they raise serious doubts about the validity of the revised Rescorla-Wagner and SOP models. The latter models fail because they are bottom-up, abstractionist models: Associative strengths are updated on the basis of the objective properties of events, and the impact that an event has on associative strength cannot be revised later on. Both the available evidence and common sense suggest that contingency judgements depend not merely on the objective properties of events but also on how these events are interpreted in the light of content-specific and abstract knowledge that the observer possesses (also see White, 2000, 2002). Moreover, people are able to recode prior events on the basis of new information about those events and to adjust their judgements accordingly. One could thus conclude that the Rescorla–Wagner and SOP models fail as models of HCL because they underestimate the active role that observers play when encoding and retrieving knowledge about contingencies.

OTHER FINDINGS AND THEORIES

Until now, we have focused on the Rescorla–Wagner model (and its cousin, the SOP model), the probabilistic contrast model (and its cousins, the Power PC model and causal model theory), and the studies that were inspired by these models. We have seen that neither the associative models nor the probabilistic models are able to account for all the results of these studies. Over the years, other findings have been reported that also raise questions about the validity of the models that we have discussed so far and instead provide support for alternative models. We summarize some of these findings and theories in this section, starting with those that were inspired by Pavlovian conditioning research.

Configural learning

When we see a stimulus such as the letter T, we can either regard it as a compound stimulus that consists of a horizontal and a vertical line, or we can regard it as a unique configural stimulus that is more than the sum of its elements (i.e., we can regard it as a letter T). Likewise, whenever two cues are presented together, participants can either encode the compound

stimulus in an elemental manner (i.e., as consisting of two individual cues) or in a configural manner (i.e., as a unique configural cue that is more than the sum of its elements).

The Rescorla–Wagner and SOP models are essentially elemental models. That is, they postulate that people learn about the elements of a compound. However, several studies have shown that people can adopt either an elemental or a configural approach and that contingency judgements depend heavily on the approach that people adopt. For instance, in a blocking-like study, Williams et al. (1994) found an impact of A trials when participants were encouraged to think of the AT compound as consisting of two individual cues (i.e., A present and T present) but not when they encoded the compound as a unique configural cue (i.e., AT is perceived as a new cue, X). Other research also supports the idea that participants can adopt a configural approach. For instance, Shanks, Charles, Darby, and Azmi (1998) showed that when participants are exposed to A+, AB–, and B+ trials, they correctly infer that the outcome is more likely to occur after A or B than after the AB compound stimulus. This result strongly suggests that participants regard the AB compound as unique configural cue whose relation with the outcome is (to a certain extent) independent from the relation between its elements (i.e., A and B) and the outcome. Although people seem to have a strong tendency to regard compound stimuli as more than the sum of the elements, under certain conditions they will nevertheless adopt an elementary approach (e.g., Shanks et al., 1998; Williams et al., 1994). That is, sometimes they will respond to the compound on the basis of the relation between its elements and the outcome. It appears that people possess a certain flexibility in adopting a configural or elemental approach.

Because the Rescorla–Wagner and SOP models are elemental models, they cannot readily explain that people can respond differently to a compound stimulus than to the elements of the compound. One could assume that in addition to learning about the elements of a compound, participants also learn about some unique features of the compound (e.g., Wagner & Brandon, 2001). Although this extends the range of phenomena that can be explained by the Rescorla–Wagner and SOP models, some findings argue against this solution and suggest that people in some cases do not learn about the elements of a compound at all (e.g., Shanks et al., 1998).

Pearce (1987, 1994) has proposed a model that only encodes the contingency between configurations of stimuli and the outcome. When Cues A and T are presented together, and this compound is followed by the outcome (i.e., AT+), participants will form an association between a representation of the configuration of A and T, on the one hand, and the representation of the outcome, on the other. Such AT+ trials will have no effect on the association between the representations of A and the outcome nor on the link between the representations of T and the outcome (unlike what is assumed in the Rescorla–Wagner and SOP models). A second crucial feature of the configural model is that it explicitly deals with generalization between stimuli. The basic idea is that responding to a certain cue or configuration of cues will depend not only on the associative strength of that cue or configuration, but also on the associative strength of similar cues and configurations. For instance, when A+ trials are followed by a test trial on which the compound AT is presented, participants will to a certain extent expect the outcome to occur even though the compound AT has itself never been paired with the outcome. This results from the fact that A is similar to AT because the latter incorporates the former. As a result of this similarity, the expectancy of the outcome that is generated by A will generalize to AT. Because learning is assumed to be expectancy based (i.e., more will be learned when the presence or absence of the outcome is unexpected), the model correctly

predicts that contingency judgements for T will be lower when AT+ trials are preceded by A+ trials (i.e., forward blocking).

The main strength of the configural model of Pearce (1987, 1994) is, of course, that it can deal with findings that show that people often learn about configurations of stimuli rather than about the individual stimuli (e.g., Shanks et al., 1998). On the negative side, it cannot deal with the observation that participants sometimes do adopt an elemental approach and that participants are to some extent flexible in adopting a configural or elemental approach. Moreover, in its present form, Pearce's configural model cannot account for retrospective cue competition effects. Ironically enough, the model is also at odds with one of the findings that demonstrates the importance of configural processing in HCL. Williams et al. (1994) showed that forward blocking does not occur when participants adopt a configural approach. According to Pearce's model, at least a partial blocking effect should occur under these conditions because part of the excitatory strength of A that was acquired on the A+ trials should generalize to the subsequent AT+ trials, thus resulting in a decrease in the associative strength that can be gained by T on the AT+ trials. Note, however, that there are other configural models that are able to account for retrospective cue competition (e.g., Le Pelley & McLaren, 2001) and that do not necessarily allow for a generalization from A to AT (e.g., Kruschke, 1992). Although these configural models are ingenious and have a lot of potential, it is difficult to see how they could model the effects of beliefs and recoding on contingency judgements (e.g., De Houwer, 2002; Waldmann, 2000).

Learning and performance

The associative models that we have discussed so far can all be regarded as learning-based models. That is, they postulate that cues acquire associative strength during the learning phase and that judgements are a direct reflection of this acquired associative strength. This essential feature of the associative models is perhaps best illustrated by the way in which they account for blocking and other cue competition effects: On the A+ trials that precede AT+ trials, A gains associative strength, which prevents T from gaining associative strength on the AT+ trials. Blocking is thus assumed to be due to a failure to learn.[4]

It is often incorrectly assumed that all associative models emphasize processes that occur during learning, whereas all probabilistic models emphasize processes that occur at the time of testing (see Miller & Escobar, 2001, for a discussion of this issue). Miller (Miller & Matzel, 1988; Miller & Schachtman, 1985), for instance, developed an associative model that emphasizes processes that occur at the time of testing. The central idea behind their *comparator model* is that associatively induced changes in behaviour depend on the comparison of the strength of different associations at the time of testing. For instance, each time that a cue T is paired with an outcome, the strength of the T–outcome association will increase. If, however, T is presented in compound with other cues, those cues will also become associated with both T and

[4]Mackintosh (1975) introduced an attentional theory of learning according to which blocking is due to the fact that participants learn not to attend to redundant cues when they already have a good predictor of the outcome. When an established predictor (A) is presented together with a new cue (T), participants will thus fail to encode T and as a result will not register the contingency between T and the outcome. Recent results indeed suggest that such an attentional mechanism also influences HCL (Glautier, 2002; Kruschke & Blair, 2001).

the outcome. During performance, judgements about the T–outcome contingency will depend on the associative strength of T relative to the associative strength of alternative cues with which T was associated (i.e., the comparator cues). Moreover, the extent to which the associative strength of the alternative cue modulates the expression of the T–outcome association depends on the strength of the association between T and the alternative cues.

The basic idea that judgements depend on the comparison of associative strengths at the time of testing allows the model to explain both forward and retrospective cue competition effects. For instance, regardless of whether AT+ trials are preceded or followed by A+ trials, the associative strength of A will increase on the A+ trials. Because responding to T will depend on the associative strength of T in comparison to the associative strength of A, responding to T will decrease as a result of the A+ trials. Note that this account of blocking is very similar to the account provided by probabilistic contrast models (see earlier). One only needs to replace "the associative strength of A"' with "the probability of the outcome in the presence of A and the absence of T" and "the associative strength of T" with "the probability of the outcome when both A and T are present".

Blaisdell, Bristol, Gunther, and Miller (1998; also see Denniston, Savastano, & Miller, 2001) recently proposed an extension of the comparator model in which even more complex comparisons take place at the time of testing. Assume that AT1+ and T1T2+ trials are followed either by A+ trials or by A– trials. According to the extended comparator model, judgements about T2 will also be influenced by whether A on its own was paired with the outcome. Judgements for T2 depend on the associative strength of T2 but also on the associative strength of T1. However, the extent to which the associative strength of T1 modulates responding to T2 will in turn depend on the associative strength of A. Although we do not go into detail regarding this extended version of the model, its prediction concerning the effect of higher order comparators has received some support both in studies of Pavlovian conditioning with animals (e.g., Blaisdell et al., 1998) and in HCL studies (De Houwer & Beckers, 2002a, b).

Despite the fact that the comparator model is fundamentally different from other associative models, like all other associative models it has difficulties in dealing with the effects that beliefs seem to have on HCL (e.g., Waldmann, 2000). Moreover, because associative strengths are updated on a trial-by-trial basis, the model cannot account for the fact that judgements depend on how participants retrospectively recode trials (e.g., De Houwer, 2002). Therefore, although associative models may differ with regard to the emphasis on processes during learning or testing, the effects of beliefs and active recoding of information seem to lie beyond the scope of all current associative models.

Rule learning

Recent research demonstrates that people can infer abstract rules from the specific events that they witness and can use those rules to make judgements about new (combinations of) cues. Shanks and Darby (1998), for instance, presented A+, B+, AB–, C–, D–, CD+ trials together with I+, J+, M–, and N– trials. During a test phase, participants judged that the outcome was more likely to occur after the (previously unseen) compound MN than after the (also previously unseen) IJ compound. The crucial aspect of this result is that it occurred even though the elements of the MN compound were not associated with the outcome, whereas the elements of the IJ compound were associated with the outcome. Judgements rather reflected the rule that

the likelihood of the outcome after a compound of two stimuli (i.e., AB–, CD+) is the reverse of the likelihood of the outcome after the elements of the compound (i.e., A+, B+, C–, D–). This important ability of people to infer and use abstract rules seems to lie beyond the scope of all current models of HCL (but see Lien & Cheng, 2000).

Trial order and frequency of judgement

Earlier we described some findings that show that the order in which trials are presented can have a profound impact on contingency judgements (e.g., Lopez et al., 1998). Catena et al. (1998) looked in detail at these trial order effects and showed that the last trial before a contingency judgement can have a disproportionate impact on the judgement. This effect, however, was present only when participants made judgements after each trial rather than after a series of trials. Thus, trial order effects can depend on the frequency with which participants are asked to make judgements. Catena et al. further observed that judgements overall became less accurate as the frequency of judgements (after a series of trials versus after each trial) increased.

To account for these findings, Catena et al. (1998) developed the belief revision model, which postulates that each contingency judgement is based on (1) the previous judgement about that same contingency and (2) the new evidence that has been presented since that previous judgement (see Hogarth & Einhorn, 1992, for a related model). Unlike the other probabilistic models but similar to most associative models, the belief revision model is an abstraction model in that all the information presented before a certain judgement contributes to one value (i.e., the numeric value of that judgement), which is updated on the basis of new information at the time of the next judgement. Abstraction thus occurs on a judgement-by-judgement basis. It is similar to the other probabilistic models in that the new evidence is summarized in terms of the probability of the outcome in the presence and absence of the cue during the events between two judgements.

When participants are asked to make a judgement after each trial, the new evidence will consist only of the trial preceding the judgement. However, if there are many trials separating two judgements, the last trial before a judgement will be only one of many trials that provide the new evidence. Therefore, the information presented on the last trial before a judgement will have a bigger impact on the judgement when participants make a judgement after each trial than when they make judgements less frequently. There is also other evidence that supports the belief revision model. Most important, Shanks and colleagues (Lober & Shanks, 2000; Perales & Shanks, 2001) were able to demonstrate, through simulations, that the belief revision model provided the best account for effects of the base rate of the outcome on HCL.

What is perhaps most interesting about the belief revision model is the idea that participants take short cuts when judging the same contingency repeatedly. They will not consider all the presented evidence on each occasion, but rather constantly revise the beliefs or hypotheses that they had already formed at an earlier stage during learning. However, the belief revision model needs further development before it can be regarded as an adequate model of HCL. Most important, it is not able to deal with cue competition effects at all. The model furthermore lacks any assumptions about how causal beliefs or configural processing can influence contingency judgements.

CONCLUDING COMMENTS

Models of HCL

Much has been learned about HCL over the past 20 years. However, none of the models that have been proposed can explain all of the existing data. Most important, all non-associative models fail to account for the effect of cue–cue associations on retrospective revaluation (e.g., Dickinson & Burke, 1996). All associative models, on the other hand, are incompatible with the fact that beliefs and retrospective recoding of events can have such a profound impact on contingency judgements (e.g., De Houwer, 2002; Waldmann, 2000).

Perhaps a hybrid model is needed that incorporates elements of both associative and non-associative models. Several authors have indeed argued that there might be more than one mechanism responsible for HCL (e.g., Dickinson, 2001; Evans & Over, 1996; Lovibond et al., 2001; Mackintosh, 1995; McLaren, Green, & Mackintosh, 1994). Most of them propose that contingency learning might depend both on rational reasoning and on simple associative mechanisms. There can be little doubt about the fact that people can make rational inferences about the contingencies or causal relations between events. It is therefore not surprising that contingency judgements *can* depend on causal beliefs or on the retrospective recoding of trials. Results such as those of Waldmann (2000, 2001) and De Houwer (2002; De Houwer & Beckers, 2002a, b) only demonstrate that participants in HCL studies *can* act in a rational, almost scientific way. It should be clear, however, that people do not always act in such a way. Even the most rational of people will only make rational inferences when she or he has the motivation and opportunity to do so. When motivation or opportunity to engage in rational reasoning are lacking, contingency judgements might well be based on the operation of associative mechanisms.

It is interesting to note that the studies that provided the strongest evidence against non-associative models (e.g., Dickinson & Burke, 1996; Larkin et al., 1998) involved a large number of cues, whereas the studies that provided the strongest evidence against associative models (e.g., De Houwer, 2002; Waldmann, 2000) involved only a few cues. One could argue that the higher the number of cues, the less likely it is that participants can keep track of all the events and use them to come to rational, justifiable judgements (see Dickinson, 2001), and the more likely it is that contingency judgements are based on the operation of associative mechanisms. It would thus be interesting to manipulate the motivation and opportunity for rational reasoning (e.g., by manipulating the number of cues, or by adding secondary tasks or time constraints). Also, one could argue that associative models might provide a particularly good account of implicit (unconscious) forms of learning because these forms of learning by definition do not depend on rational inferences (e.g., Hendrickx & De Houwer, 1997).

The interplay between Pavlovian conditioning and HCL research

Over the past 20 years, Pavlovian conditioning research has provided a rich source of inspiration for scholars of HCL and will probably continue to do so in the years to come. For instance, context-dependent learning (i.e., occasion setting) is currently a prominent topic of research in Pavlovian conditioning. Examining whether HCL is also context dependent could shed important light on the nature of the processes that underlie contingency learning.

On the other hand, knowledge that has been gained from recent HCL research will probably start influencing Pavlovian conditioning research and theory. For instance, the role of causal beliefs has rarely been studied in Pavlovian conditioning. Initial research suggests that they might well have a considerable impact on Pavlovian conditioning. Mitchell and Lovibond (2002), for instance, found that ceiling effects also influence the magnitude of blocking effects in Pavlovian conditioning of human autonomic responses. HCL research could also inspire research on Pavlovian conditioning in non-human animals. For example, although it is difficult to assess the role of beliefs in Pavlovian conditioning with animals, recent findings in HCL regarding the role of causal beliefs might revive the idea that Pavlovian conditioning in animals can be conceived of as a form of causal attribution rather than as a form of merely predictive learning (e.g., Revusky & Garcia, 1970; Testa, 1974).

We are therefore convinced that the interplay between Pavlovian conditioning and HCL research will become even more intense in the years to come. Note that the merits of this interplay do not depend on whether one can demonstrate that the same phenomena occur in Pavlovian conditioning and HCL. If anything, conflicting results might even be more informative about the processes that are involved in both forms of learning. We therefore eagerly look forward to the next 20 years of research.

REFERENCES

Aitken, M.R.F., Larkin, M.J.W., & Dickinson, A. (2001). Re-examination of the role of within-compound associations in the retrospective revaluation of causal judgements. *Quarterly Journal of Experimental Psychology*, *54B*, 27–51.

Anderson, J.R., & Sheu, C-F. (1995). Causal inferences as perceptual judgments. *Memory and Cognition*, *23*, 510–524.

Blaisdell, A.P., Bristol, A.S., Gunther, L.M., & Miller, R.R. (1998). Overshadowing and latent inhibition counteract each other: Support for the comparator hypothesis. *Journal of Experimental Psychology: Animal Behavior Processes*, *24*, 335–351.

Catena, A., Maldonado, A., & Cándido, A. (1998). The effect of the frequency of judgement and the type of trials on covariation learning. *Journal of Experimental Psychology: Human Perception and Performance*, *24*, 481–495.

Chapman, G.B. (1991). Trial order affects cue interaction in contingency judgement. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *17*, 837–854.

Cheng, P.W. (1997). From covariation to causation: A causal power theory. *Psychological Review*, *104*, 367–405.

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

Cheng, P.W., & Novick, L.R. (1990). A probabilistic contrast model of causal induction. *Journal of Personality and Social Psychology*, *58*, 545–567.

Cheng, P.W., & Novick, L.R. (1992). Covariation in natural causal induction. *Psychological Review*, *99*, 365–382.

De Houwer, J. (2002). Forward blocking depends on retrospective inferences about the presence of the blocked cue during the elemental phase. *Memory & Cognition*, *30*, 24–33.

De Houwer, J., & Beckers, T. (2002a). Second-order backward blocking and unovershadowing in human causal learning. *Experimental Psychology*, *49*, 27–33.

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

De Houwer, J., Beckers, T., & Glautier, S. (2002). Outcome and cue properties modulate blocking. *Quarterly Journal of Experimental Psychology*, *55A*, 965–985.

Dennis, M.J., & Ahn, W. (2001). Primacy in causal strength judgements. *Memory & Cognition*, *29*, 152–164.

Denniston, J.C., Savastano, H.I., & Miller, R.R. (2001). The extended comparator hypothesis: Learning by contiguity, responding by relative strength. In R.R. Mowrer & S.B. Klein (Eds.), *Handbook of contemporary learning theories* (pp. 65–117). Mahwah, NJ: Lawrence Erlbaum Associates, Inc.

Dickinson, A. (2001). The 28th Bartlett memorial lecture: Causal learning: An associative analysis. *Quarterly Journal of Experimental Psychology*, *54B*, 3–25.

Dickinson, A., & Burke, J. (1996). Within-compound associations mediate the retrospective revaluation of causality judgements. *Quarterly Journal of Experimental Psychology*, *49B*, 60–80.

Dickinson, A., Shanks, D., & Evenden, J. (1984). Judgement of act–outcome contingency: The role of selective attribution. *Quarterly Journal of Experimental Psychology*, *36A*, 29–50.

Evans, J.St.B.T., & Over, D.E. (1996). *Rationality and reasoning*. Hove, UK: Psychology Press.

Glautier, S. (2002). Spatial separation of target and competitor cues enhances blocking of human causality judgements. *Quarterly Journal of Experimental Psychology*, *55B*, 121–136.

Hendrickx, H., & De Houwer, J. (1997). Implicit covariation detection. *Psychologica Belgica*, *37*, 29–50.

Hogarth, R.M., & Einhorn, J. (1992). Order effects in belief updating: the belief–adjustment model. *Cognitive Psychology*, *24*, 1–55.

Hume, D. (1987). *A treatise of human nature* (2nd ed.). Oxford: Clarendon Press. (Original work published 1739.)

Kamin, L.J. (1969). Predictability, surprise, attention, and conditioning. In R.M. Church (Ed.), *Punishment and aversive behaviour* (pp. 279–296). New York: Appleton-Century-Crofts.

Kant, I. (1965). *Critique of pure reason*. London: Macmillan. (Original work published 1781.)

Killeen, P.R. (2001). The four causes of behavior. *Current Directions in Psychological Science*, *10*, 136–140.

Kruschke, J.K. (1992). ALCOVE: An exemplar-based connectionist model of category learning. *Psychological Review*, *99*, 22–44.

Kruschke, J.K., & Blair, N.J. (2001). Blocking and backward blocking involve learned inattention. *Psychological Bulletin and Review*, *7*, 636–645.

Larkin, M.J.W., Aitken, M.R.F., & Dickinson, A. (1998). Retrospective revaluation of causal judgements under positive and negative contingencies. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *24*, 1331–1352.

Le Pelley, M.E., & McLaren, I.P.L. (2001). Retrospective revaluation in humans: Learning or memory. *Quarterly Journal of Experimental Psychology*, *54B*, 311–352.

Lien, Y., & Cheng, P.W. (2000). Distinguishing genuine from spurious causes: A coherence hypothesis. *Cognitive Psychology*, *40*, 87–137.

Lober, K., & Shanks, D.R. (2000). Is causal induction based on causal power? Critique of Cheng (1997). *Psychological Review*, *107*, 195–212.

Lopez, F.J., Almaraz, J., Fernandez, P., & Shanks, D.R. (1999). Adquisicion progresiva del conocimiento sobre relaciones predictivas: Curvas de aprendizaje en juicios de contingencia [Gradual acquisition of knowledge of predictive relationships: Learning curves in contingency judgments]. *Psicothema*, *11*, 337–349.

Lopez, F.J., Shanks, D.R., Almaraz, J., & Fernandez, P. (1998). Effects of trial order on contingency judgements: A comparison of associative and probabilistic contrast accounts. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *24*, 672–694.

Lovibond, P.F., Been, S.-L., Mitchell, C.J., Bouton, M.E., & Frohardt, R. (2001). *Forward and backward blocking of causal judgement is enhanced by additivity of effect magnitude*. Manuscript submitted for publication.

Mackintosh, N.J. (1975). A theory of attention: Variation in the associability of stimuli with reinforcement. *Psychological Review*, *82*, 276–298.

Mackintosh, N.J. (1995). Categorization by people and pigeons: The twenty-second Bartlett memorial lecture. *Quarterly Journal of Experimental Psychology*, *48B*, 193–214.

Markman, A.B. (1989). LMS rules and the inverse base-rate effect: Comment on Gluck and Bower (1988). *Journal of Experimental Psychology: General*, *118*, 417–421.

Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information*. San Francisco: Freeman.

Matute, H., & Pine o, O. (1998a). Cue competition in the absence of compound training: Its relation to paradigms of interference between outcomes. In D.L. Medin (Ed.), *The psychology of learning and motivation* (Vol. 38, pp. 45–81). San Diego, CA: Academic Press.

Matute, H., & Pine o, O. (1998b). Stimulus competition in the absence of compound conditioning. *Animal Learning & Behavior*, *26*, 3–14.

McLaren, I.P.L., Green, R.E.A., & Mackintosh, N.J. (1994). Animal learning and the implicit/explicit distinction. In N.C. Ellis (Ed.), *Implicit and explicit learning of languages* (pp. 313–332). New York: Academic Press.

Michotte, A. (1954). *La perception de la causalité*. Leuven, Belgium: Publications Universitaires de Louvain.

Miller, R.R., & Escobar, M. (2001). Contrasting acquisition–focused and performance-focused models of acquired behavior. *Current Directions in Psychological Science*, *10*, 141–145.

Miller, R.R., & Matzel, L.D. (1988). The comparator hypothesis: A response rule for the expression of associations. In G.H. Bower (Ed.), *The psychology of learning and motivation* (Vol. 22, pp. 51–92). San Diego, CA: Academic Press.

Miller, R.R., & Schachtman, T.R. (1985). Conditioning context as an associative baseline: Implications for response generation and the nature of conditioned inhibition. In R.R. Miller & N.E. Spear (Eds.), *Information processing in animals: Conditioned inhibition* (pp. 51–88). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

Mitchell, C.J., & Lovibond, P.F. (2002). Backward and forward blocking in human autonomic conditioning requires an assumption of outcome additivity. *Quarterly Journal of Experimental Psychology*, *55B*, 311–329.

Pearce, J.M. (1987). A model for stimulus-generalization in Pavlovian conditioning. *Psychological Review*, *94*, 61–73.

Pearce, J.M. (1994). Similarity and discrimination—A selective review and a connectionist model. *Psychological Review*, *101*, 587–607.

Perales, J.C., & Shanks, D.R. (2001). *Normative accuracy of regularity-based causal estimation in one-cause one-outcome learning situations: Effects of contingency and causal power*. Manuscript submitted for publication.

Pine o, O., Ortega, N., & Matute, H. (2000). The relative activation of the associations modulates interference between elementally-trained cues. *Learning and Motivation*, *31*, 128–152.

Rescorla, R.A. (1968). Probability of shock in the presence and absence of CS in fear conditioning. *Journal of Comparative and Physiological Psychology*, *66*, 1–5.

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

Revusky, S., & Garcia, J. (1970). Learned associations over long delays. In G.H. Bower & J.T. Spence (Eds.), *The psychology of learning and motivation* (Vol. 4, pp. 1–77). New York: Academic Press.

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

Shanks, D.R. (1987). Acquisition functions in causality judgement. *Learning and Motivation*, *18*, 147–166.

Shanks, D.R. (1995). *The psychology of associative learning.* Cambridge: Cambridge University Press.

Shanks, D.R., Charles, D., Darby, R.J., & Azini, A. (1998). Configural processes in human associative learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *24*, 1353–1378.

Shanks, D.R., & Darby, R.J. (1998). Feature- and rule-based generalization in human associative learning. *Journal of Experimental Psychology: Animal Behavior Processes*, *24*, 405–415.

Shanks, D.R., Holyoak, K.L., & Medin, D.L. (Eds.). (1996). *The psychology of learning and motivation; Vol. 43, Causal learning*. San Diego, CA: Academic Press.

Shanks, D.R., Lopez, F.L., Darby, R.J., & Dickinson, A. (1996). Distinguishing associative and probabilistic contrast theories of human contingency judgement. In D.R. Shanks, K.J. Holyoak, & D.L. Medin (Eds.), *The psychology of learning and motivation* (Vol. 34, pp. 265–311). San Diego, CA: Academic Press.

Testa, T.J. (1974). Causal relationship and the acquisition of avoidance responses. *Psychological Review*, *78*, 3–43.

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

Van Overwalle, F., & Van Rooy, D. (2001). When more observations are better than less: A connectionist account of the acquisition of causal strength. *European Journal of Social Psychology*, *31*, 155–175.

Wagner, A.R. (1981). SOP. A model of automatic memory processing in animal behavior. In N.E. Spear & R.R. Miller (Eds.), *Information processing in animals: Memory mechanisms* (pp. 5–47). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

Wagner, A.R., & Brandon, S.E. (2001). A componential theory of Pavlovian conditioning. In R.R. Mowrer & S.B. Klein (Eds.), *Handbook of contemporary learning theories* (pp. 23–64). Mahwah, NJ: Lawrence Erlbaum Associates, Inc.

Waldmann, M.R. (2000). Competition among causes but not effects in predictive and diagnostic learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 26, 53–76.

Waldmann, M.R. (2001). Predictive versus diagnostic causal learning: Evidence from an overshadowing paradigm. *Psychonomic Bulletin & Review*, *8*, 600–608.

Waldmann, M.R., & Holyoak, K.J. (1992). Predictive and diagnostic learning within causal models: Asymmetries in cue competition. *Journal of Experimental Psychology: General*, *121*, 22–236.

Wasserman, E.A., & Berglan, L.R. (1998). Backward blocking and recovery from overshadowing in human causal judgement: The role of within–compound associations. *Quarterly Journal of Experimental Psychology*, *51B*, 121–138.

White, P.A. (2000). Causal judgement from contingency information: Relation between subjective reports and individual tendencies in judgement. *Memory and Cognition*, *28*, 415–426.

White, P.A. (2002). Perceiving a strong causal relation in a weak contingency: Further investigation of the evidential evaluation model of causal judgement. *Quarterly Journal of Experimental Psychology*, *55A*, 97–114.

Williams, D.A., Sagness, K.E., & McPhee, J.E. (1994). Configural and elemental strategies in predictive learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *20*, 694–709.

Original manuscript received 24 January 2002
Accepted revision received 8 March 2002