

Comparing associative, statistical, and inferential reasoning accounts of human contingency learning

Oskar Pineño

University of Seville, Seville, Spain

Ralph R. Miller

State University of New York at Binghamton, Binghamton, NY, USA

For more than two decades, researchers have contrasted the relative merits of associative and statistical theories as accounts of human contingency learning. This debate, still far from resolution, has led to further refinement of models within each family of theories. More recently, a third theoretical view has joined the debate: the inferential reasoning account. The explanations of these three accounts differ critically in many aspects, such as level of analysis and their emphasis on different steps within the information-processing sequence. Also, each account has important advantages (as well as critical flaws) and emphasizes experimental evidence that poses problems to the others. Some hybrid models of human contingency learning have attempted to reconcile certain features of these accounts, thereby benefiting from some of the unique advantages of different families of accounts. A comparison of these families of accounts will help us appreciate the challenges that research on human contingency learning will face over the coming years.

“Then you better start swimmin’, or you’ll sink like a stone, for the times they are a-changin’.” With these words of a classic Bob Dylan song, one could describe the feeling left after attending the Human Contingency Learning (HCL) meeting in the Ardennes, Belgium, in May 2004. Why such a feeling? The reason is that a third account of HCL, based on inferential reasoning, joined the debate between statistical and associative theories that has been in progress for the last two decades. Moreover, it became obvious at the HCL meeting that this third account constitutes a challenge to these prior views. Advocates of the

inferential reasoning approach claim that both the associative and statistical accounts have ignored the role of higher order, complex cognitive processes and have reported several phenomena that seemingly support this assertion. This challenge will require extending existing associative and statistical theories if they are to survive and will probably bring with it scientific advancement.

Theoretical traditions in HCL

The relevance of the debate among statistical, associative, and inferential reasoning accounts of

Correspondence should be addressed to Ralph R. Miller, Department of Psychology, SUNY-Binghamton, Binghamton, NY 13902-6000, USA. E-mail: rmiller@binghamton.edu

The development of this paper was possible due to support from Department of Universities, Research, and Technology of the Andalucía Government (Junta de Andalucía) and NIMH Grant 33881. We would like to thank Jeffrey C. Amundson, Gonzalo Urcelay, and Daniel Wheeler for their comments on an earlier version of this manuscript.

HCL can only be understood by considering the issues that demanded the most attention of researchers in HCL over the last 20 years. In fact, each of these theoretical views was developed to account for certain critical findings. In this section we briefly describe these theoretical accounts in relation to the findings that prompted their appearance. It is not our purpose to provide here a detailed discussion on the wide range of different models within each account, but to explain the general ideas underlying these views (for other recent reviews, see Allan & Tangen, 2005; De Houwer & Beckers, 2002a; Wills, 2005). For the sake of simplicity, the statistical, associative, and inferential reasoning families of models are discussed based on the most prominent member of each family.

Early contingency studies: The hegemony of statistical theories

Statistical theories of learning were initially developed as an attempt to explain how human participants learn the contingencies underlying the occurrence of two or more covarying events. The term *contingency* refers to covariation between two or more *binary* variables (i.e., discrete variables with only two possible values; in this sense, the study of contingency learning can be viewed as the study of the simplest instance of correlation learning, since *correlation* implies covariation between two or more *continuous* variables). In a typical contingency learning experiment, participants are given presentations of two or more events, which can be either present or absent on each learning trial (i.e., discrete presentations of a set of responses and/or stimuli). Then, participants are asked to rate the contingent (or sometimes causal) relationship between these events, generally by using a numerical scale in which different values correspond to different degrees of relation. Because of this simple experimental scenario, researchers in HCL were able to collect considerable data, which rapidly opened the door for further and more refined studies.

The first studies of HCL focused on the most basic question—that is, is human learning sensitive to different contingencies? And, if so, how

sensitive is it? Initial studies, such as that of Jenkins and Ward (1965), showed evidence of such a sensitivity in instrumental learning (i.e., between a response and an outcome). This ability of humans to detect response–outcome relations has been extensively documented (e.g., Allan & Jenkins, 1980, 1983; Alloy & Abramson, 1979; Chatlosh, Neunaber, & Wasserman, 1985; Neunaber & Wasserman, 1986; Wasserman, 1990; Wasserman, Chatlosh, & Neunaber, 1983; Wasserman, Elek, Chatlosh, & Baker, 1993; for a review see Wasserman, Kao, Van Hamme, Katagiri, & Young, 1996), as has been the ability of humans to detect cue–outcome relations (e.g., López, Almaraz, Fernández, & Shanks, 1999; Shanks, 1991; for a review see Shanks, López, Darby, & Dickinson, 1996b).

Early evidence demonstrating humans' sensitivity in contingency detection prompted the development of models that explained this kind of learning from a statistical view. Although there are many potential statistical indices of contingency, Allan's (1980; also see Kelley, 1967) ΔP model is the simplest example of this family of models (see Shanks, Holyoak, & Medin, 1996a, for a review of many of these models). Despite being inspired in part by a model originally developed in the animal conditioning literature (i.e., Rescorla, 1968), the impact of the ΔP model was so profound in the HCL literature, that it still has not been totally replaced by recent and more refined statistical models, which now tend to narrow the focus from contingency learning in general, to causal learning in particular (e.g., Cheng, 1997; Cheng & Novick, 1992; Spellman, 1996). Rather, the ΔP measure was implemented, in one way or another, in these new models. Specifically, in the ΔP model information concerning the cue (or response) and outcome is acquired and stored as co-occurrence frequencies in a 2×2 matrix (see Figure 1), which results from the combination of the binary values of the cue (present vs. absent) and the outcome (present vs. absent). In this matrix, Cell *a* records the number (frequency) of trials on which both the cue and the outcome were present (*fa*), Cell *b* records the number of trials on which the cue

	Outcome present	Outcome absent
Cue present	a fa	b fb
Cue absent	c fc	d fd

Figure 1. 2×2 contingency matrix. Cells in this matrix store the frequencies of the presence and the absence of the cue and the outcome: Cell a = frequency of cue and outcome present (fa); Cell b = frequency of cue present and outcome absent (fb); Cell c = frequency of cue absent and outcome present (fc); Cell d = frequency of cue and outcome absent (fd).

was present, and the outcome was absent (fb), Cell c records the number of trials on which the cue was absent, and the outcome was present (fc), and Cell d records the number of trials on which both the cue and the outcome were absent (fd). These frequencies are used to compute ΔP :

$$\Delta P = \frac{a}{a+b} - \frac{c}{c+d} = P(O|C) - P(O|\sim C) \quad (1)$$

According to Equation 1, ΔP corresponds to the difference between two conditional probabilities: the probability of the outcome given the cue, $P(O|C)$, and the probability of the outcome in the absence of the cue, $P(O|\sim C)$. Therefore, ΔP always adopts a value ranging between +1 and -1. When the outcome occurs more frequently in the presence of the cue than in its absence, the cue–outcome contingency is positive (i.e., generative learning in the causal learning tradition and excitatory learning in the animal conditioning tradition). By contrast, when the outcome occurs more frequently in the absence of the cue than in its presence, the cue–outcome contingency is negative (i.e., preventative learning in the causal learning tradition and inhibitory learning in the animal conditioning tradition). Finally, the contingency is 0 when the cue and the outcome occur independently of each other.

Allan's (1980) ΔP model provided a simple account for much observed covariation learning.

As acknowledged even by opponents of statistical models (see Shanks et al., 1996a, for a debate between supporters of statistical and associative accounts), values of ΔP resulting from different frequencies of occurrence of the cue (or response) and the outcome closely resemble the actual ratings or nonverbal responses found in studies assessing sensitivity to different contingencies in humans (for supportive evidence, see, e.g., Wasserman et al., 1996; but for conflicting evidence, see, e.g., Wasserman et al., 1993). However, the hegemony of ΔP in HCL was soon challenged by associative theorists from the animal conditioning tradition, and the difficulty of the basic ΔP model to account for cue competition effects played an important role in this turn of events.

Cue competition effects: The emergence of associative models in HCL

In HCL, cue competition effects refer to phenomena in which ratings of a target cue paired with an outcome are influenced by the presence of a second cue during the pairings. Cue competition was first reported in the classical conditioning literature by Pavlov (1927), who observed the well-known *overshadowing effect*. In overshadowing, two cues, A and X, are presented in compound during their pairings with the outcome (i.e., AX–O trials). As a consequence of the training of X in compound with a typically more salient cue, A, weak responding is observed on test of X, relative to a condition in which X was trained alone (i.e., X–O trials).

Despite the early finding of the overshadowing effect by Pavlov (1927), Kamin's (1968) *blocking effect* is the cue competition effect that has received most attention in the literature. In blocking, as in overshadowing, two cues (A and X) are trained in compound with the outcome, also resulting in weak responding at test of the target cue, X. However, contrary to overshadowing, in blocking A does not compete with X based on its higher salience (i.e., in blocking, A is ordinarily equally or less salient than X), but on its history of additional pairings with the outcome. In a typical forward blocking experiment, A is paired alone with the outcome (i.e., A–O trials) prior to compound conditioning (AX–O trials), and responding to X at

test is usually found to be weak, relative to a control condition that received no A–O trials prior to the AX–O trials (instead, this condition usually receives B–O trials).

The blocking effect became a benchmark to be met by accounts of HCL due to the studies of Dickinson, Shanks, and Evenden (1984), who first reported this effect in a HCL task. Their observation of blocking, an effect previously studied in the animal conditioning literature, led researchers to suggest that the same mechanisms may be involved in animal conditioning and HCL. In fact, Dickinson et al. explicitly suggested that associative theories developed to explain classical conditioning in nonhuman animals (e.g., Rescorla & Wagner, 1972) might also be used to account for HCL (see also Gluck & Bower, 1988; Shanks & Dickinson, 1987). Among the many models in the associative tradition (e.g., Dickinson & Burke, 1996; Mackintosh, 1975; Pearce & Hall, 1980; Rescorla & Wagner, 1972; Van Hamme & Wasserman, 1994; Wagner, 1981), the model of Rescorla and Wagner has certainly been the most influential. Despite its having many detractors even in the associative tradition (e.g., Baker, 1974; Bouton, 1993; R. R. Miller, Barnet, & Grahame, 1995), nobody questions its relevance as the focal model that encouraged the development of the associative account of cue competition.

In Rescorla and Wagner's (1972) model, the formation and strengthening of a cue–outcome association proceeds according to Equation 2:

$$\Delta V_{\text{cue}}^n = \alpha \cdot \beta \cdot (\lambda - V_{\text{Total}}^{n-1}) \quad (2)$$

In this equation, ΔV_{cue}^n represents the change in associative strength of the cue on trial n , and α and β are learning-rate parameters representing the associability (sometimes equated with salience) of the cue and the outcome, respectively. These parameters adopt values between 0 and 1, as a function of their corresponding salience. The parenthetical term (i.e., $\lambda - V_{\text{Total}}^{n-1}$) represents the discrepancy between the amount of associative strength that can be supported by the outcome

(λ) and the current total associative strength acquired, through trial $n-1$, by all the cues present on trial n (V_{Total}^{n-1}). The value of λ depends on the presence or absence of the outcome on trial n : When the outcome is presented, λ adopts a value of 1; when the outcome is absent, λ adopts a value of 0. Cue–outcome associations are acquired and strengthened based on the discrepancy between the expected and actual occurrence of the outcome (i.e., $\lambda - V_{\text{Total}}^{n-1}$). Therefore, in acquisition training (i.e., cue–outcome pairings), the total of the strengths of all of the cue–outcome associations will increase toward unity based on a positive discrepancy (i.e., $1 - V_{\text{Total}}^{n-1}$), whereas in extinction and inhibition training (i.e., cue-alone trials), the total of the strengths of all of the cue–outcome association will decrease toward zero based on a negative discrepancy (i.e., $0 - V_{\text{Total}}^{n-1}$).

Cue competition effects are explained by Rescorla and Wagner's (1972) model through the discrepancy in the parenthetical term ($\lambda - V_{\text{Total}}^{n-1}$), which represents what Kamin (1968) referred to as the surprisingness of the outcome (i.e., the occurrence of the outcome on a given trial is assumed to be more or less surprising based on the associative strength of the cues present on that trial). Specifically, cue competition arises because the discrepancy term for each cue depends on the total associative strength of all the cues present on trial n (V_{Total}^{n-1}). Based on this feature of the model, the associative strength of a cue can have an impact upon the acquisition of associative strength by a second cue with which it was presented. For example, in forward blocking A first gains associative strength during the A–O trials. Then, during the subsequent AX–O trials, the summated associative strength of A and X will be close to the value of λ even on the first AX–O trial, due to A's previously acquired associative strength, thereby resulting in little learning to X. In Kamin's terminology, the outcome is not surprising due to its being already expected based on the presence of A. As a consequence of this small discrepancy, the amount of associative strength available for cue X to acquire will be small, and, hence, a weak X–O association will be formed.

The ΔP model cannot explain the blocking effect. According to this model, the contingency between X and O decreases on trials on which the outcome occurs without X, regardless of whether these trials consist of A–O pairings (i.e., blocking condition) or B–O pairings (control condition). In other words, this model does not address the fact that blocking, and other cue competition effects, are due to outcome presentations, not only in the absence of the target cue, but also in the presence of the target cue's companion stimulus. This inability of the ΔP model to account for cue competition effects in HCL encouraged the development of more sophisticated statistical models, such as Cheng and Novick's (1992) focal-set theory, Cheng's (1997) Power-PC theory, and Spellman's (1996) use of conditional contingencies (see also Spellman, Price, & Logan, 2001). At the same time, new data concerning cue competition in HCL also encouraged the revision of some traditional associative models, such as Van Hamme and Wasserman's (1994) revision of Rescorla and Wagner's (1972) model, and Dickinson and Burke's (1996) revision of Wagner's (1981) SOP model. Specifically, these revisions were prompted by the observation of retrospective revaluation in HCL (e.g., backward blocking, Shanks, 1985), which demonstrated that, contrary to the assumptions of traditional associative models, a cue's associative status could be updated on trials on which this cue was absent, but associatively activated by the presentation of a companion stimulus. Because retrospective revaluation was subsequently also detected in animal conditioning preparations (see, e.g., R. R. Miller & Matute, 1996) the study of HCL, which imported the associative account from the animal learning tradition, soon returned the favour by stimulating the identification of new animal learning phenomena. But the debate between associative and statistical models, despite its being so active during the nineties (e.g., Allan, 1993; Shanks, 1995; Shanks et al., 1996), lost some of its strength over the last few years. Recent exchanges between associative and statistical researchers are no longer aimed at globally comparing the two

families of accounts of HCL, perhaps due to the recognition that whole families of models could not be contrasted due to the open-ended generation of revised models within each family (e.g., R. R. Miller & Escobar, 2001). Instead, HCL studies focused on testing specific models, such as Cheng's Power-PC theory (e.g., Allan, 2003; Buehner, Cheng, & Clifford, 2003; Lober & Shanks, 2000). Moreover, during the last few years a new view of HCL was proposed that introduced a wholly new orientation: the inferential reasoning account.

Effects of outcome additivity and maximality on cue competition: The challenge of the inferential reasoning view

Tracking down the theoretical origins of the inferential reasoning view as applied to HCL is not an easy task. One precursor was probably the propositional logic proposed for higher order information processing in some of the cognitive literature (e.g., Anderson, 1980; Johnson-Laird, 1988; Lindsay & Norman, 1972; G. A. Miller, Galanter, & Pribram, 1960; Wason & Johnson-Laird, 1972), in which propositions are used to solve behavioural tasks. The inferential reasoning view also has roots in the theoretical tradition concerned with causal reasoning, which is represented in the study of HCL by the work of Waldmann (1996, 2000, 2001; see also Waldmann & Holyoak, 1992) and by the philosophically oriented work of Pearl (e.g., Pearl, 1993, 2000). This causal reasoning approach has recently developed into the Bayesian net models (e.g., Glymour, 2003; Waldmann & Hagmayer, 2005; Waldmann & Martignon, 1998).

The inferential reasoning account of HCL has been encouraged by recent studies performed by De Houwer and his collaborators (e.g., Beckers, De Houwer, Pineño, & Miller, 2005; De Houwer, 2002; De Houwer & Beckers, 2003; De Houwer, Beckers, & Glautier, 2002) as well as by Lovibond and his colleagues (e.g., Lovibond, Been, Mitchell, Bouton, & Frohardt, 2003; Mitchell & Lovibond, 2002; Mitchell, Lovibond, & Condoleon, 2005; see also Waldmann & Walker, 2005; Wu & Cheng, 1999, for related

proposals). As a whole, these studies have tried to demonstrate that cue competition effects, specifically forward and backward blocking, might arise from inferential processes. Two specific findings can be regarded as the most compelling support for this view: effects of outcome additivity and effects of outcome maximality.

The *outcome additivity effect* (Beckers et al., 2005; Lovibond et al., 2003; Mitchell & Lovibond, 2002; Mitchell et al., 2005) refers to the observation that, prior to target training in cue competition situations, experience with two effective (i.e., reinforced) nontarget cues alone and with an effective compound of the two cues with the same outcome attenuates subsequent cue competition. The inferential account of this phenomenon is based on the notion that, in order for blocking to occur, the participant must assume that multiple effective causes have additive effects on the occurrence (i.e., frequency and/or intensity) of the outcome. This assumption concerning additivity of effective causes upon the outcome, together with the observation in a blocking experiment that the addition of the target cause to the blocking cause does not actually affect the intensity of the outcome, is what, from the inferential reasoning view, produces the blocking effect. More specifically, the blocking effect is postulated to be the result of the following counterfactual reasoning process, akin to a syllogism: (a) if both cues A and X are potential causes of the outcome, then the outcome should be stronger when A and X are presented together than when they are presented alone; (b) the outcome following the AX compound (AX-O trials) is not stronger than that following A alone (A-O trials); (c) thus, A and X are not both effective causes of the outcome; (d) since A alone causes the outcome, and information on the effectiveness of X alone was not provided, it is logical to assume that X is not an effective cause of the outcome. The inferential reasoning account has received support from experiments showing that, if participants are given explicit additivity pretraining with cues different from those involved in the blocking procedure (e.g., C and D separately paired with a moderate outcome; CD compound paired with

an intense outcome, then A-O trials followed by AX-O trials), robust blocking of X is found, whereas explicit nonadditivity pretraining (e.g., C, D, and CD each paired with an outcome of identical intensity, either moderate or intense, then A-O trials followed by AX-O trials), little or no blocking of X is observed (e.g., Beckers et al., 2005; Livesey & Boakes, 2004; Lovibond et al., 2003). This implies that although additivity of causes appears to be the default expectation, experience with nonadditivity of causes can attenuate this expectation.

Related to the outcome additivity effect, the *outcome maximality effect* has also yielded support for the inferential reasoning account. Outcome maximality (Beckers et al., 2005; De Houwer et al., 2002) refers to the notion that, in order for participants to assess outcome additivity, it is critical for them to know that the outcome experienced was not limited by some physical constraint such as a maximum possible frequency or intensity (i.e., a ceiling effect). Additivity is not expected if it requires an outcome that is greater than what the subject believes to be possible. More specifically, in a blocking procedure, participants are only able to determine whether the outcome following an AX compound cue is greater than that produced by A alone or X alone (consistent with the outcome, default, additivity assumption) if the outcome presented on A-O trials is known to be of a submaximal intensity such that more intense outcomes could have been observed. Simply put, if the presentation of A results in a maximal outcome, it is not possible to determine the effectiveness of X as a cause of the outcome based on trials on which the AX compound is also followed by a maximal outcome. By contrast, if A and AX each produce the same submaximal outcome, participants can then logically infer that the presence of cue X does not influence the outcome's intensity because, if it did, a stronger outcome would have occurred following the AX compound.

Other evidence seemingly supportive of the inferential reasoning account includes the observation that increasing the difficulty of a secondary task performed during cue competition treatment attenuates cue competition (De Houwer &

Beckers, 2003; Vandorpe, De Houwer, & Beckers, 2005; Waldmann & Walker, 2005). This suggests that cue competition occurs only when sufficient amounts of limited cognitive faculties can be devoted to inference. It has also been claimed that evidence of inference in cue competition has been provided by studies in which different cover stories were presented in order to provide either a predictive or a causal scenario for the cue–outcome relations. In these studies, competition was obtained only with the causal scenario (De Houwer et al., 2002; for related findings, see Waldmann & Holyoak, 1992). The conclusion here is based on the assumption that outcome additivity is more readily applied to causation than to prediction, an assumption that leaves unexplained other reports of cue competition in clearly predictive situations (e.g., Arcediano, Matute, & Miller, 1997; Chapman, 1991; Chapman & Robbins, 1990; Williams, Sagness, & McPhee, 1994).

Regardless of the specific manipulations used in each study, the message sent by researchers working in the inferential reasoning framework is consistent: Both the associative and statistical accounts of HCL have erred in overlooking the role that higher order cognitive processes may play in cue competition effects. Moreover, the claim that these higher order cognitive processes might be involved in HCL experiments, but not in experiments using nonhuman animals, has been brought into question. As demonstrated by Beckers, Miller, De Houwer, and Urushihara (2006), outcome additivity and maximality effects can be readily found in a Pavlovian conditioning preparation using rats as subjects. Parsimony leads to the tentative conclusion that rats, as well as humans, use inferential processes or, alternatively, that neither do and some noninferential process is actually responsible for these observations.

A comparison of accounts of HCL

Despite differences among the models within each family (statistical, associative, and inferential reasoning), accounts within a family share certain features and assumptions. In the present section we compare these families of accounts

regarding: (a) their levels of explanation of learning phenomena, (b) how information is encoded and treated, (c) their focus on either acquisition or performance processes to explain HCL phenomena, and (d) their ability to explain path-dependence phenomena. As will be seen, each account has certain advantages and disadvantages. However, because these accounts are not totally exclusive, adopting an integrative strategy can solve, at least partially, some of the problems of each account. This integrative strategy has been already followed by some hybrid models.

Explanatory levels in accounts of HCL

Models differ not only in the specific mechanisms they propose to explain phenomena, but also in the level of analysis of their explanations. Although different categories of explanatory levels have been proposed, Marr's (1982) three-level distinction has become commonplace in the study of HCL (e.g., Baker, Murphy, & Vallée-Tourangeau, 1996; López, Cobos, Caño, & Shanks, 1998). According to Marr's distinction, theoretical explanations can be viewed as computational, algorithmic, or implementational. *Computational* models merely describe what the output should be, given an input, remaining silent on the process by which information is transformed going from input to output. With respect to HCL, computational models describe the goals guiding behaviour, speaking of *what* should be expected to occur in a given situation. In this sense, computational models can be viewed as normative models because they describe what is presumably functional—that is, the participant's ideal response under a given set of contingencies. *Algorithmic* models provide a deeper explanatory level than computational models by trying to answer not only what output should be expected in a given situation, but also how this output was arrived at. A model is said to be algorithmic when it provides not only representations for the input and the output, but also an algorithm (i.e., a procedure with a limited number of steps) for transforming the input into the output. Thus, algorithmic models of HCL speak of the mechanisms involved in the acquisition, storage, retrieval,

and/or expression of information regarding covarying events. Finally, *implementational* models speak to how the algorithms proposed by algorithmic models can be physically implemented. An implementational model of HCL would attempt to describe the neurobiological basis of contingency learning.

HCL accounts are generally silent on the implementational level¹ and therefore fall in either the computational (normative) or algorithmic categories. Associative models (e.g., Dickinson & Burke, 1996; Mackintosh, 1975; Pearce & Hall, 1980; Rescorla & Wagner, 1972; Van Hamme & Wasserman, 1994; Wagner, 1981) are algorithmic. Thus, associative models provide not only the equations for transforming the input (e.g., the occurrence or absence of the cues and the outcome, together with their saliences) into the output (e.g., response or, at least, associative strength), but also a well-defined set of rules for transforming information step by step. However, this is not the distinctive feature of associative models that makes them algorithmic (i.e., statistical models also provide equations and rules for transforming the input and the output). Rather, associative models are algorithmic because they attempt to describe the psychological mechanisms presumably involved in learning phenomena. Also, as algorithmic models, associative models speak of mental representations of the input (i.e., internal representations of the cues and the outcome) and the outcome (i.e., associative and/or response strengths).

At first glance, statistical models can be considered computational accounts. However, this claim must be qualified, since this assertion is open to controversy. Certainly, statistical models are usually viewed as normative by researchers in the associative tradition (e.g., Baker et al., 1996; López et al., 1998) as well as in the statistical tradition (e.g., Cheng, 1997; Glymour & Cheng, 1998). However, several studies have suggested the possibility that information might be

processed through the application of a *statistical mechanism*, which would make them algorithmic. For example, Price and Yates (1995; see also Catena, Maldonado, Megías, & Frese, 2002; Shanks, 1991) suggested that the format in which contingency information is presented (i.e., trial-by-trial vs. summarized presentation in contingency tables) might encourage processing by either an associative mechanism (trial-by-trial) or a statistical mechanism (contingency tables).

Despite occasional proposals that a statistical mechanism could be involved in HCL, however, there is a general agreement on viewing these statistical models as normative accounts. The case of inferential reasoning accounts can be said to be exactly the opposite. Authors working in the inferential reasoning framework (e.g., Beckers et al., 2005; Lovibond et al., 2003) frequently speak of these accounts as a description of cognitive mechanisms actually involved in cue competition phenomena. This is explicit in statements such as “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 judgments *can* depend on causal beliefs or on the retrospective recoding of trials” (De Houwer & Beckers, 2002a, p. 306). In a similar vein, Lovibond et al. (2003, p. 141) stated “According to such an inferential approach, causal judgments may be derived from the same sorts of inferential or deductive processes that participants employ in other complex reasoning tasks”. But is the contemporary inferential reasoning account algorithmic? Accounts of cue competition based on inferential processes certainly attempt to explain how higher order cognitive processes influence effects like blocking. Thus, inferential reasoning accounts, like associative accounts, are algorithmic in that they presumably describe the underlying processing of information by the subject. The problem with these accounts is that, at least in their current form, these inferential

¹ Although at the implementation level there are models of human learning based on neuroscience, our discussion of HCL models exclusively refers to those within the cognitive tradition.

processes are vaguely defined by a qualitative set of rules. However, we must acknowledge that this account is young and still in an early stage of development. Perhaps, within the next few years the inferential reasoning account of HCL will reach a level of formalization similar to that of associative models.

Information encoding and treatment

HCL accounts differ in the way that contingency-related information is assumed to be encoded. As previously mentioned, statistical models such as ΔP (Allan, 1980) assume that, for a single cue and outcome, subjects behave as if this information is encoded as co-occurrence frequencies for each of the cells of the 2×2 contingency matrix, or perhaps even memories of each trial. In contrast, in associative models (e.g., Rescorla & Wagner, 1972) information is encoded in the form of a few summary statistics, usually a single associative strength. Finally, inferential reasoning models are silent concerning the form of information encoding, although it is implicit that this encoding is detailed. This could be conceptualized as either a 2×2 contingency matrix or a set of rules or propositions based on each cell (e.g., "Cue A produces the outcome"). In this latter case, however, it would not be clear whether information is in fact encoded and stored in these rules or these propositions are mere verbal descriptions of information actually encoded as either frequencies or associations. For example, if information encoding was purely associative, and this information was accessible to conscious awareness, participants could verbalize this information in the form of a rule.

Differences in the way information is encoded foreshadow important advantages and disadvantages of the various HCL accounts. One of the implications of encoding information as co-occurrence frequencies is that it allows for the storage of richer memories than are assumed by associative models, minimally the frequency of each type of trial. This information-encoding strategy, in addition to allowing the computation of the different contingencies among events, allows for the retrieval of specific and detailed information, such as the number of times certain

events were conjointly present or absent (but not necessarily specific memory of each trial). Such nearly veridical memories could not be encoded and stored by associative models, which suffer from a paucity of encoded information. Frequency of co-occurrence information is demanding of memory capacity, more so than simple associative strength but less so than fully veridical memory. However, nearly veridical memories might not be necessary within current models of HCL because they merely attempt to explain the participants' response to a cue based on the cue-outcome relationship, something for which co-occurrence frequencies, associative strengths, and rules or propositions are equally valid. Without further use of co-occurrence frequencies, having encoded them is of no special merit, and one must acknowledge the greater parsimony of associative strength as being uniquely meritorious. Moreover, advocates of the associative account could claim that, although information encoding as frequencies in a contingency matrix allows for the retrieval of information like the number of occasions a specific trial type occurred, not all statistical models make use of this advantage. For example, in the ΔP model, detailed information such as the number of occasions on which different trial types took place only serves the purpose of computing the conditional probabilities of the outcome both in the presence and in the absence of the target cue. These probabilities are then used to compute the expected subject's response or verbal rating (i.e., the value of ΔP). It is therefore fair to say that the co-occurrence frequencies are unused in the framework of the ΔP model after the value of ΔP is computed. Nevertheless, due to its finer grain, information encoded in the form of frequencies in contingency tables can be acquired under more varied circumstances than can associatively encoded information. According to associative models, information can be acquired only through trial-by-trial training, whereas statistical models are open to the acquisition of new information both on a trial-by-trial basis and from frequency lists (e.g., Catena et al., 2002; Price & Yates, 1995; Shanks, 1991). However, the emphasis of associative models on the trial-by-trial encoding

of information is not necessarily a disadvantage in all situations. It is precisely this feature that allows associative models to explain some trial-order effects, particularly recency effects (i.e., dominant impact of most recently acquired information on behaviour), which cannot be explained by statistical models without additional assumptions, such as time stamps on memories and weights dependent upon these stamps (for a detailed discussion, see Pineño & Miller, 2005).

Despite the potential advantages of frequency-based encoding relative to associative encoding, it is notable that statistical models make little use of this potential advantage. Moreover, they suffer from problems arising precisely from the way they use these frequencies. Because conditional probabilities are sensitive to relative as opposed to absolute frequencies, statistical models cannot simulate acquisition curves (Baker, Mercier, Vallée-Tourangeau, Frank, & Pan, 1993; Chatlosh et al., 1985; Shanks, 1987; but see Baker, Berbrier, & Vallée-Tourangeau, 1989).² For example, these models predict asymptotic responding after only one cue–outcome pairing (i.e., ΔP is 1 after only one acquisition trial provided that it is a Cell *a* type trial). Inferential models suffer from the same weakness; one reinforced trial (X–O) should establish the proposition that X is followed by O. The prediction of graded acquisition curves is a unique success of associative models.³

Acquisition versus performance processes in HCL phenomena

HCL accounts also differ in their emphasizing different stages of information processing to explain HCL phenomena, such as cue competition. In statistical models, as discussed before,

subjects behave as if contingency information was directly encoded as co-occurrence frequencies in a contingency matrix. The acquisition of information is assumed to be free from any filter, taking place in a noncompetitive manner. In these models, cue competition arises from the computation of contingency. For example, in Allan's (1980) ΔP , blocking occurs because the A–O trials give rise to a high probability of the occurrence of the outcome in the absence of the target cue, $P(O|\sim C)$, which reduces the impact that the probability of the occurrence of the outcome in the presence of the target cue, $P(O|C)$, would otherwise have on the final contingency rating. (Admittedly this is not a compelling account of blocking because blocking is frequently seen, in a group that receives A–O followed by AX–O, relative to a control group that receives B–O followed by AX–O.) Therefore, to the extent that statistical models speak to cue competition, it is due to the impaired expression of the relationship between the target cue and the outcome. Most associative models, by contrast, assume that cue competition arises from processing during the acquisition stage. For example, the Rescorla and Wagner (1972) model assumes that the presence of the competing cue impairs learning of the target cue–outcome association. Finally, the inferential models implicitly hold a performance-focused view because, in order for participants to infer and apply rules, they must be able to process information previously acquired in a noncompetitive manner. In this vein, Beckers et al. (2005) demonstrated that outcome additivity training interpolated between the blocking treatment and test (Experiment 4) enhanced the blocking effect just like it did when it was given prior to the blocking treatment (Experiment 2). Therefore,

² The problem for statistical models of not predicting acquisition curves might be solved by assuming that the 2×2 contingency table is preexperimentally filled with noise, such that the initial values of the cells are greater than zero—that is, $(fa = fb = fc = fd) > 0$. However, it is not clear how these cells, for a specific cue and a specific outcome, could be available in memory prior to any kind of experience with these stimuli. If this approach were to be adopted, one would have to assume the existence of virtually an infinite number of cells for any possible combination of cues and outcomes.

³ One might argue that only averaged acquisition curves are gradual because individual curves usually show a step-like function (e.g., Gallistel, Fairhurst, & Balsam, 2004). However, even individual curves do not commonly show an abrupt increase to asymptote after a single cue–outcome pairing, as predicted by statistical and inferential models.

the results of these experiments by Beckers et al. suggest that inferences are compared at the time of test, thereby supporting a view of inferential models as performance focused. However, in another paper De Houwer and Beckers (2003) reported that a distractor task during training (Experiment 1) or during both training and testing (Experiment 2) attenuated blocking, which in the inferential reasoning framework suggests that inferences are initially drawn during training and then compared during testing. Thus, critical information processing appears to occur during both training and testing, which realistically is not surprising. Nevertheless, it would be interesting to examine the consequences of a distractor task only during testing to determine if it too would attenuate blocking.

The stages of information processing in which cue competition is assumed to take place is one of the differentiating features of associative and statistical models of HCL that has received considerable attention (for a review, see R. R. Miller & Escobar, 2001). As part of this debate between acquisition-focused models (most associative models) and performance-focused models (statistical and inferential models), considerable evidence has been marshalled demonstrating that responding elicited at test by a target cue can be strongly affected by treatments performed after training with the target cue that do not directly involve this cue. These effects are referred to as retrospective revaluation (e.g., Dickinson & Burke, 1996), in which following compound training with two cues (e.g., AX-O trials), further training with A influences responding to cue X at test. Specifically, when Cue A is further paired with the outcome (i.e., A-O trials) following AX-O trials, weak responding to X at test is often found—an effect known as backward blocking (e.g., Shanks, 1985; Wasserman & Berglan, 1998). By contrast, when A-alone trials are given following AX-O trials, enhanced responding to X is usually found at test—an effect known as recovery from overshadowing (e.g., Larkin, Aitken, & Dickinson, 1998; Wasserman & Berglan, 1998). Retrospective revaluation cannot be accounted for by traditional associative models

(e.g., Mackintosh, 1975; Pearce & Hall, 1980; Rescorla & Wagner, 1972; Wagner, 1981) because, according to them, the associative strength of a cue can be updated only on trials on which the cue is present. Specifically, according to these models, any change in associative strength undergone by a cue is a direct function (among other factors) of the cue's salience, which adopts a positive value when the cue is present and a zero value when it is absent. Because, in retrospective revaluation effects, the target cue is absent during the revaluation treatment of its companion cue, no change in its associative status is predicted to occur from these models. That is, in this framework any treatment following training with the AX compound that does not include the presence of the target cue, X (e.g., A-O or A-alone trials) is expected to have no impact on X's associative strength. For some years, until the revision of some traditional associative models (e.g., Dickinson & Burke, 1996; Van Hamme & Wasserman, 1994), retrospective revaluation effects and other treatments not directly involving the target cue were considered to be exclusive support for performance-focused models (e.g., statistical models among others). Because, in these models all information is encoded and stored in order to be expressed in a competitive manner at the time of testing, the order in which different trial types take place during training is completely irrelevant. Therefore, A-O or A-alone trials are expected to have the same effect upon responding to X regardless of whether they are given before, during, or after training with the AX compound trials. On the one hand, this readily accounts for retrospective revaluation effects such as backward blocking. On the other hand, that statistical models ignore trial order is a mixed blessing because trial order is not irrelevant despite the occurrence of retrospective revaluation phenomena. For example, forward blocking is most often more robust than backward blocking all other factors being equal (e.g., Melchers, Lachnit, & Shanks, 2004; but see De Houwer et al., 2002).

Retrospective revaluation effects are still a fertile field for research. Recent studies have

shown evidence of higher order retrospective revaluation (e.g., De Houwer & Beckers, 2002b, 2002c; Macho & Burkart, 2002), an effect consisting of changes in responding to an absent cue, X, due to training with indirect associates of X. For example, in second order retrospective revaluation, a cue, X, is paired with another cue, A, which in turn is paired with a third cue, B. The effect of further training of Cue B alone is two-fold: (a) a change in responding to Cue A, contrary to the change undergone by B (first order retrospective revaluation) and (b) a change in responding to cue X, in the same direction as the change undergone by B (second order retrospective revaluation). Findings on higher order retrospective revaluation have posed important difficulties even to models that were able to account for first order retrospective revaluation (e.g., Dickinson & Burke, 1996; Van Hamme & Wasserman, 1994). However, these effects can be readily explained by an extension of the comparator hypothesis (Denniston, Savastano, & Miller, 2001), developed to explain similar phenomena in animal conditioning (Denniston, Savastano, Blaisdell, & Miller, 2003), as well as some statistical models implementing recursive probability contrasts (Macho & Burkart, 2002).

Path dependence

Another manner in which statistical models might be viewed as superior to associative models is in their ability to anticipate path dependence. This prediction stems from statistical models' assumption of nearly veridical memory, rather than retention of only summary statistics as is assumed by associative models (see earlier section *Information Encoding and Treatment*). Path dependence refers to the influence of a cue's complete training history on its current behavioural control. Path dependence is usually demonstrated by taking two cues that elicit the same responding despite different training histories and showing that radically different behaviours emerge as a function of a common added treatment. A compelling example of path dependence was provided in a study by Bouton (1984) concerned with reinstatement (i.e., recovery of responding after extinction treatment due to

outcome-alone presentations). In this study (Experiment 5), steps were made to obtain comparable responding to a cue in conditions receiving different treatments with the cue, either mere acquisition training or acquisition training followed by extinction treatment. Despite comparable responding prior to outcome-alone presentations, reinstatement was found to occur only in the condition given extinction treatment with the cue.

Evidence of a cue's path dependence poses a problem to associative models because, in these models, the only information that is retained is the current associative strength (and sometimes the current associability) of the cue, regardless of the training conditions that produced the association. The problem of path independence in most associative models (e.g., Rescorla & Wagner, 1972) is profound because, in these models, effects like extinction are due to the destruction of the previously acquired cue–outcome association. In other words, these models hypothesize active erasure of old memories (i.e., catastrophic interference, e.g., McCloskey & Cohen, 1989). Statistical models avoid anticipating catastrophic interference because, according to these models, no memory is assumed to be erased during interference treatments such as extinction. Rather, statistical models explain interference, either between cues (e.g., blocking) or between outcomes (e.g., extinction), as the addition of new information to the contingency matrix, resulting in a weakened contingency estimation without any loss of memory of pairings between the target cue and the outcome. For example, extinction training would increase the cue–no-outcome frequency (i.e., Cell *b* in the contingency table), thereby directly reducing the probability of the outcome in the presence of the cue, $P(O|C)$. By contrast, cue competition such as blocking would be due to an increase in the no-cue–outcome frequency (i.e., Cell *c* in the contingency table). Although this treatment would not directly affect $P(O|C)$, the final contingency rating would be diminished due to the increased probability of the outcome in the absence of the cue, $P(O|\sim C)$. Importantly, within the statistical account, no memory is assumed to be erased.

Rather, it is the addition of new memories that results in interference phenomena.

Like statistical models, inferential models are free from the problem of catastrophic interference. According to inferential models, interference (either between cues or between outcomes) is explained as due to the development of new rules that, when successfully applied at the time of testing, result in low ratings of the target cue. For example, extinction training would allow the participants to infer a new rule of the kind “the cue is no longer a predictor or cause of the outcome”, without necessarily forgetting the previous cue–outcome pairings. An alternative new inference is “the cue is not consistently a predictor or cause of the outcome”. The uncertainty as to which of these alternatives would constitute the new inference illustrates the lack of full specification of inferential theory at this time. Notably, neither inference can readily account for spontaneous recovery as a function of the retention interval.

Hybrid models: The integrative approach

In the previous section we discussed some advantages and problems of the associative, statistical, and inferential reasoning accounts of HCL. It is obvious from this discussion that each family of models has something unique to offer, but that there are flaws intrinsic to each account that cannot be easily overcome without incorporating some assumptions from other families.⁴ This is precisely the strategy followed by some researchers: to incorporate some of the notions from different accounts into a single model. This integrative approach, resulting in hybrid models of HCL, has proven able to partially overcome the difficulties of pure models (e.g., strictly associative or statistical models).

Hybrid models in the associative tradition

Some theorists have developed hybrid models that incorporate notions from both statistical and

inferential reasoning accounts (e.g., causal Bayes nets, see Glymour, 2003; Waldmann & Hagmayer, 2005; Waldmann & Martignon, 1998). Other types of hybrid model have been developed in the associative tradition in order to account for various types of Pavlovian phenomena. Such is the case of R. R. Miller and Matzel's (1988) comparator hypothesis and Bouton's (1993) retrieval failure model, which are readily applicable to HCL, although each was initially developed to account for specific findings in the Pavlovian paradigm. These two models can be viewed as essentially associative (i.e., they assume that contingency information is encoded and stored as associative strengths); however, they incorporate some mechanisms typical of statistical models. Like statistical theories, in these models interference phenomena are not due to acquisition deficits or unlearning (i.e., catastrophic interference), but to learning of associations that interfere with either retrieval or expression of the target association. In the case of Bouton's model, which was developed to account for interference between outcomes trained apart (e.g., extinction), these effects are due to an inhibitory cue–outcome association that coexists with the excitatory cue–outcome association and interferes with retrieval of the excitatory cue–outcome association at test. By contrast, in R. R. Miller and Matzel's comparator hypothesis, cue competition effects such as blocking are due to the impaired expression of the association between the target cue and the outcome. This expression deficit occurs because the presentation at test of the target cue (X) not only directly activates the outcome (O) representation (i.e., X–O association), but also indirectly activates a second representation of the outcome (i.e., X activates the representation of the competing cue, A, through the X–A within-compound association, and, in turn, this activation of Cue A activates the representation of the outcome through the A–O association).

Through this integrative approach, this hybrid model is able to overcome many of the problems

⁴ Even a traditional model like ΔP can be viewed as a hybrid model in that it uses an associative structure to represent the cue–outcome pairings chronicled in Cell *a*.

of each purebred account. For example, the comparator hypothesis proposes a kind of information encoding (i.e., associative) that is economical in comparison to that of statistical models (i.e., frequencies of trial types), but rich enough to allow interactions between different memories at the time of behavioural expression. Therefore, despite this model summarizing information in terms of associative strengths, it bases its explanations of many learning phenomena on specific rules by which associations interact in order to produce responding at test, an approach similar to that of statistical models in their use of frequencies and the corresponding conditional probabilities in computing expected contingency ratings at test. But this integrative approach also involves a trade-off. As occurred with statistical models, because much or all information is acquired and retained as different associations and because associations only interact at the time of testing, this associative–statistical hybrid model is unable, without additional assumptions, to simulate trial-order effects, something that traditional associative models could explain (e.g., Rescorla & Wagner, 1972; for a detailed explanation, see Pineño & Miller, 2005). Notably, Bouton's (1993) model clearly predicts trial-order effects as a result of an additional assumption. Specifically, Bouton suggests that both inhibitory associations and second-learned associations wane with time faster than excitatory associations or first-learned associations. This assumption readily accounts for trial-order phenomena such as spontaneous recovery.

Putting top-down and bottom-up processes together in a single model

Associative accounts of HCL, as previously discussed, assume bottom-up processing of information. Using simple automatic processes, these models explain how humans acquire and transform information on event covariation into adaptive responding. Similarly, statistical accounts of HCL assume that participants behave as if they were doing bottom-up information processing. In contrast, inferential reasoning accounts assume that participants actively contrast hypotheses at

the time of testing, and that behaviour indicative of phenomena such as cue competition arises when these hypotheses are either confirmed or disconfirmed by experience. In these top-down inferential reasoning accounts, higher order cognitive processes control or at least modulate HCL. Moreover, it has been suggested that these processes may even be the basis of classical conditioning phenomena with both humans and nonhumans (Beckers et al., 2006).

However, as we previously mentioned, inferential models are incompletely specified, and their predictions are often vague. Hence, an account that includes both bottom-up processing (either associative or statistical) and top-down processing (inferential reasoning) might provide a more complete picture of HCL. Bottom-up processing is necessary to explain how information concerning covarying events is encoded and stored in the system, something that the inferential reasoning account currently fails to do. Top-down processing would allow covariation learning to be under the influence of more complex cognitive processes. This integrative approach between bottom-up and top-down processes, which has already been followed in the aforementioned Bayes nets, could also be used to account for outcome additivity and maximality effects (e.g., Beckers et al., 2005; Lovibond et al., 2003) in the framework of an associative or statistical model.

Any integration of an associative or statistical model with the inferential reasoning view, however, will face some theoretical difficulties, such as those arising from the mathematical formalization of inferences based on outcome additivity and maximality. However, such a formalization of inferential rules into quantitatively defined algorithms could be achieved by taking advantage of inference's close similarity to Boolean logic, which is already implemented in numerous computer programming languages. Perhaps more challenging, theoretical problems would arise from subtle interactions among the features of each type of model in a single framework. Let us assume that an associative model such as Rescorla and Wagner's (1972) model and an inferential reasoning model were to be the

basis for building a hybrid. Some phenomena are explicable within each contributing model, which is unparsimonious, and, more problematic, the relative contributions of each contributing model toward explaining these phenomena would be unclear. For example, it would be necessary to determine what causes cue competition effects in such a hybrid model. That is, are effects like blocking due to an acquisition deficit (Rescorla & Wagner) or to successful acquisition but a failure to observe outcome additivity (inferential process)? If blocking is assumed to be due exclusively to the inferential process, some modifications will have to be made to the Rescorla and Wagner model in order to allow for noncompetitive acquisition of different cue–outcome associations (i.e., akin to the earlier model by Bush & Mosteller, 1951). However, even in this case, it would be necessary to refine the rules for the acquisition of associative strength and rules for drawing inference so that the strength of the associations could somehow be related to the additivity and maximality inferential rules. For example, the maximality rule (i.e., participants can only assess outcome additivity when the outcome is submaximal) applies only in cases in which the associative strength of the blocking cue, A , is already asymptotic. Consider the case in which V_A is 0.5 and V_X is 0.0. Even if a submaximal outcome (e.g., $\lambda = .7$, following Rescorla and Wagner's assumption of 1.0 as the maximal value of λ) were paired with the AX compound, there would still be room for participants to assess if cue X has any impact on the outcome (i.e., because $\lambda - V_{AX} = .2$). In summary, any formalization of a hybrid model that includes inferential reasoning processes will be a difficult task.

Some critical issues will have to be addressed before such a hybrid model can be formalized. For example, although it has been demonstrated that pretraining on outcome additivity enhances the blocking effect (e.g., Beckers et al., 2005), it is well known, since Dickinson et al.'s (1984) original report, that blocking in HCL can be observed with no need of such pretraining. Are we to conclude that humans and other animals presume as a default assumption the additivity of the causal

strengths of different cues presented in compound? The answer to this question is important because such an assumption would open a wide field for new questions such as, are these a priori assumptions innate or are they learned? If so, when and how are they learned? Given that outcome additivity effects have been observed not only in humans but also in rats (Beckers et al., 2006), answering those questions promises to be a difficult but potentially illuminating enterprise.

As an alternative to the possibility of a priori assumptions concerning outcome additivity and maximality, it is possible, although perhaps unparsimonious, that cue competition effects can take place at multiple stages in information processing. For example, such effects could occur due to a low-order process (i.e., acquisition, retrieval, and/or expression failure) as well as inferential reasoning processes. In this case, outcome additivity pretraining could enhance the observed cue competition effect by adding to a low-level process already taking place the impact of higher order reasoning (e.g., outcome additivity). However, the lower level process would have to result in an expression deficit because an acquisition deficit would not provide the knowledge base necessary for the inferential process. In sum, although an integrative approach involving both top-down and bottom-up processes would be able to account for most currently chronicled phenomena in the HCL literature, details are yet to be specified, and the product will suffer the complexities of any hybrid model that tends to thwart unambiguous predictions.

Summary and conclusion

The present paper has offered a brief comparison of the associative, statistical, and inferential reasoning accounts of HCL. Researchers can adopt one or another account of HCL, but it is fair to admit at this time that each account provides different insights on the conditions and processes underlying HCL phenomena. Because of this, we have suggested that, for the short term, future understanding of HCL may depend upon the development of hybrid models that integrate

specific notions from different accounts. In fact, some of the already existing hybrid models have proven successful in dealing with those phenomena for which they were developed, relative to the theoretically pure accounts on which they are based. We anticipate that research in HCL during the next few years will prove to be fertile thanks in part to the addition of inferential reasoning accounts to the long-standing debate between associative and statistical accounts.

In the Introduction we mentioned that the field of HCL is rapidly evolving. The inferential reasoning account has sent a message, and we certainly have to take it seriously. Time and research will determine if the inferential reasoning view survives new evidence and further revisions of associative and statistical approaches. Will the inferential account stand on its own or will it symbiotically share assumptions with other accounts (i.e., associative or statistical)? This question will not be answered until the inferential view is more clearly specified. In the meantime, the fresh air brought by the inferential reasoning view to our field will reinvigorate the study of HCL and will encourage revision and development of models of HCL.

REFERENCES

- Allan, L. G. (1980). A note on measurement of contingency between two binary variables in judgement tasks. *Bulletin of the Psychonomic Society*, *15*, 147–149.
- Allan, L. G. (1993). Human contingency judgments: Rule based or associative? *Psychological Bulletin*, *114*, 435–448.
- Allan, L. G. (2003). Assessing Power PC. *Learning & Behavior*, *31*, 192–204.
- Allan, L. G., & Jenkins, H. M. (1980). The judgment of contingency and the nature of the response alternatives. *Canadian Journal of Psychology*, *34*, 1–11.
- Allan, L. G., & Jenkins, H. M. (1983). The effect of representations of binary variables on judgment of influence. *Learning and Motivation*, *14*, 381–405.
- Allan, L. G., & Tangen, J. M. (2005). Judging relationships between events: How do we do it? *Canadian Journal of Experimental Psychology*, *59*, 22–27.
- Alloy, L. B., & Abramson, L. Y. (1979). Judgment of contingency in depressed and nondepressed students: Sadder but wiser? *Journal of Experimental Psychology: General*, *108*, 441–485.
- Anderson, J. R. (1980). *Cognitive psychology and its implications*. New York: W. H. Freeman and Co.
- Arcediano, F., Matute, H., & Miller, R. R. (1997). Blocking of conditioned responding in humans. *Learning and Motivation*, *28*, 188–199.
- Baker, A. G. (1974). Conditioned inhibition is not the symmetrical opposite of conditioned excitation: A test of the Rescorla–Wagner model. *Learning and Motivation*, *5*, 369–379.
- Baker, A. G., Berbrier, M. W., & Vallee-Tourangeau, F. (1989). Judgments of a 2 × 2 contingency table: Sequential processing and the learning curve. *Quarterly Journal of Experimental Psychology*, *41B*, 65–97.
- Baker, A. G., Mercier, P., Vallée-Tourangeau, F., Frank, R., & Pan, M. (1993). Selective associations and causality judgments: Presence of a strong causal factor may reduce judgments of a weaker one. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *19*, 414–432.
- Baker, A. G., Murphy, R. A., & Vallee-Tourangeau, F. (1996). Associative and normative models of causal induction: Reacting to versus understanding cause. In D. R. Shanks, K. J. Holyoak, & D. L. Medin (Eds.), *The psychology of learning and motivation* (Vol. 34, pp. 1–45). New York: Academic Press.
- Beckers, T., De Houwer, J., Pineño, O., & Miller, R. R. (2005). Outcome additivity and outcome maximality influence cue competition in human causal learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *31*, 238–249.
- Beckers, T., Miller, R. R., De Houwer, J., & Urushihara, K. (2006). Reasoning rats: Forward blocking in Pavlovian animal conditioning is sensitive to constraints of causal inference. *Journal of Experimental Psychology: General*, *135*, 92–102.
- Bouton, M. E. (1984). Differential control by context in the inflation and reinstatement paradigms. *Journal of Experimental Psychology: Animal Behavior Processes*, *10*, 56–74.
- Bouton, M. E. (1993). Context, time, and memory retrieval in the interference paradigms of Pavlovian learning. *Psychological Bulletin*, *114*, 80–99.
- Buehner, M. J., Cheng, P. W., & Clifford, D. (2003). From covariation to causation: A test of the assumption of causal power. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *29*, 1119–1140.

- Bush, R. R., & Mosteller, F. (1951). A mathematical model for simple learning. *Psychological Review*, *58*, 313–323.
- Catena, A., Maldonado, A., Megias, J. L., & Frese, B. (2002). Judgement frequency, belief revision, and serial processing of causal information. *Quarterly Journal of Experimental Psychology*, *55B*, 267–281.
- Chapman, G. B. (1991). Trial order affects cue interaction in contingency judgment. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *17*, 837–854.
- Chapman, G. B., & Robbins, S. J. (1990). Cue interaction in human contingency judgment. *Memory & Cognition*, *18*, 537–545.
- Chatlosh, D. L., Neunaber, D. J., & Wasserman, E. A. (1985). Response–outcome contingency: Behavioral and judgmental effects of appetitive and aversive outcomes with college students. *Learning and Motivation*, *16*, 1–34.
- Cheng, P. W. (1997). From covariation to causation: A causal power theory. *Psychological Review*, *104*, 367–405.
- 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). A review of recent developments in research and theories on human contingency learning. *Quarterly Journal of Experimental Psychology*, *55B*, 289–310.
- 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. (2002c). Second-order backward blocking and unovershadowing in human causal learning. *Experimental Psychology*, *49*, 27–33.
- De Houwer, J., & Beckers, T. (2003). Secondary task difficulty modulates forward and backward blocking in human contingency learning. *Quarterly Journal of Experimental Psychology*, *56B*, 345–357.
- De Houwer, J., Beckers, T., & Glautier, S. (2002). Outcome and cue properties modulate blocking. *Quarterly Journal of Experimental Psychology*, *55A*, 965–985.
- Denniston, J. C., Savastano, H. I., Blaisdell, A. P., & Miller, R. R. (2003). Cue competition as a retrieval deficit. *Learning and Motivation*, *34*, 1–31.
- 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*. Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.
- 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. R., & Evenden, J. (1984). Judgement of act–outcome contingency: The role of selective attribution. *Quarterly Journal of Experimental Psychology*, *36A*, 29–50.
- Gallistel, C. R., Fairhurst, S., & Balsam, P. (2004). The learning curve: Implications of a quantitative analysis. *PNAS*, *101*, 13124–13131.
- Gluck, M., & Bower, G. H. (1988). From conditioning to category learning: An adaptative network model. *Journal of Experimental Psychology: General*, *117*, 227–247.
- Glymour, C. (2003). Learning, prediction and causal Bayes nets. *Trends in Cognitive Science*, *7*, 43–48.
- Glymour, C., & Cheng, P. J. (1998). Causal mechanism and probability: A normative approach. In M. Oaksford & N. Chater (Eds.), *Rational models of cognition*. Oxford, UK: Oxford University Press.
- Jenkins, H., & Ward, W. (1965). Judgment of contingency between responses and outcomes. *Psychological Monographs*, *7*, 1–17.
- Johnson-Laird, P. N. (1988). *The computer and the mind: An introduction to cognitive science*. Cambridge, MA: Harvard University Press.
- Kamin, L. J. (1968). “Attention-like” processes in classical conditioning. In M. R. Jones (Ed.), *Miami symposium on the prediction of behavior: Aversive stimulation* (pp. 9–31). Miami, FL: University of Miami Press.
- Kelley, H. H. (1967). Attribution theory in social psychology. *Nebraska Symposium on Motivation*, *15*, 192–238.
- 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.
- Lindsay, P. H., & Norman, D. A. (1972). *Human information processing. An introduction to psychology*. New York: Academic Press.
- Livesey, E. J., & Boakes, R. A. (2004). Outcome additivity, elemental processing and blocking in human

- causality judgements. *Quarterly Journal of Experimental Psychology*, 57B, 361–379.
- Lober, K., & Shanks, D. R. (2000). Is causal induction based on causal power? Critique of Cheng (1997). *Psychological Review*, 107, 195–212.
- López, F. J., Almaraz, J., Fernández, P., & Shanks, D. (1999). Adquisición progresiva del conocimiento sobre relaciones predictivas: Curvas de aprendizaje en juicios de contingencia. *Psicothema*, 11, 337–349.
- López, F. J., Cobos, P. L., Caño, A., & Shanks, D. R. (1998). The rational analysis of human causal and probability judgement. In M. Oaksford & N. Chater (Eds.), *Rational models of cognition* (pp. 314–352). Oxford, UK: Oxford University Press.
- Lovibond, P. F., Been, S., Mitchell, C. J., Bouton, M. E., & Frohardt, R. (2003). Forward and backward blocking of causal judgment is enhanced by additivity of the effect magnitude. *Memory & Cognition*, 31, 133–142.
- Macho, S., & Burkart, J. (2002). Recursive retrospective revaluation of causal judgments. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 28, 1171–1186.
- Mackintosh, N. J. (1975). A theory of attention: Variations in the associability of stimuli with reinforcement. *Psychological Review*, 82, 276–298.
- Marr, D. (1982). *Vision: A computational approach*. San Francisco: Freeman & Co.
- McCloskey, M., & Cohen, N. (1989). Catastrophic interference in connectionist networks: The sequential learning problem. In G. H. Bower (Ed.), *The psychology of learning and motivation* (Vol. 24, pp. 109–165). San Diego, CA: Academic Press.
- Melchers, K. G., Lachnit, H., & Shanks, D. R. (2004). Within-compound associations in retrospective revaluation and in direct learning: A challenge for comparator theory. *Quarterly Journal of Experimental Psychology*, 57B, 25–53.
- Miller, G. A., Galanter, E., & Pribram, K. H. (1960). *Plans and the structure of behavior*. New York: Holt, Rinehart & Winston.
- Miller, R. R., Barnet, R. C., & Grahame, N. J. (1995). Assessment of the Rescorla–Wagner model. *Psychological Bulletin*, 117, 1–24.
- 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., & Matute, H. (1996). Animal analogues of causal judgment. In D. R. Shanks, K. J. Holyoak, & D. L. Medin (Eds.), *The psychology of learning and motivation* (Vol. 34, pp. 133–166). San Diego, CA: Academic Press.
- 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.
- Mitchell, C. J., & Lovibond, P. F. (2002). Backward and forward blocking in human electrodermal conditioning: Blocking requires an assumption of outcome additivity. *Quarterly Journal of Experimental Psychology*, 55B, 311–329.
- Mitchell, C. J., Lovibond, P. F., & Condoleon, M. (2005). Evidence for deductive reasoning in blocking of causal judgments. *Learning and Motivation*, 36, 77–87.
- Neunaber, D. J., & Wasserman, E. A. (1986). The effects of unidirectional versus bidirectional rating procedures on college students' judgments of response–outcome contingency. *Learning and Motivation*, 17, 162–179.
- Pavlov, I. P. (1927). *Conditioned reflexes*. London: Clarendon Press.
- Pearce, J. M., & Hall, G. (1980). A model for Pavlovian learning: Variations in the effectiveness of conditioned but not of unconditioned stimuli. *Psychological Review*, 87, 532–552.
- Pearl, J. (1993). Graphical models, causality, and interventions. *Statistical Science*, 8, 266–269.
- Pearl, J. (2000). *Causality: Models, reasoning, and inference*. Cambridge, UK: Cambridge University Press.
- Pineño, O., & Miller, R. R. (2005). Primacy and recency effects in extinction and latent inhibition: A selective review with implications for models of learning. *Behavioural Processes*, 69, 223–235.
- Price, P. C., & Yates, J. F. (1995). Associative and rule-based accounts of cue interaction in contingency judgement. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 21, 1639–1655.
- 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-Century-Crofts.

- Shanks, D. R. (1985). Forward and backward blocking in human contingency judgment. *Quarterly Journal of Experimental Psychology*, *37B*, 1–21.
- Shanks, D. R. (1987). Acquisition functions in contingency judgment. *Learning and Motivation*, *18*, 147–166.
- Shanks, D. R. (1991). On similarities between causal judgements in experienced and described situations. *Psychological Science*, *2*, 341–350.
- Shanks, D. R. (1995). Is human learning rational? *Quarterly Journal of Experimental Psychology*, *48A*, 29–50.
- Shanks, D. R., & Dickinson, A. (1987). Associative accounts of causality judgement. In G. H. Bower (Ed.), *The psychology of learning and motivation* (Vol. 21, pp. 229–261). San Diego, CA: Academic Press.
- Shanks, D. R., Holyoak, K. J., & Medin, D. L. (1996a). *Causal learning: The psychology of learning and motivation* (Vol. 34). New York: Academic Press.
- Shanks, D. R., López, F. J., Darby, R. J., & Dickinson, A. (1996b). Distinguishing associative and probabilistic contrast theories of human contingency judgment. In D. R. Shanks, K. J. Holyoak, & D. L. Medin (Eds.), *The psychology of learning and motivation: Causal learning* (pp. 265–311). San Diego, CA: Academic Press.
- Spellman, B. A. (1996). Acting as intuitive scientists: Contingency judgments are made while controlling for alternative potential causes. *Psychological Science*, *7*, 337–342.
- Spellman, B. A., Price, C. M., & Logan, J. M. (2001). How two causes are different from one: The use of (un)conditional information in Simpson's paradox. *Memory & Cognition*, *29*, 193–208.
- Vandorpe, S., De Houwer, J., & Beckers, T. (2005). Further evidence for the role of inferential reasoning in forward blocking. *Memory & Cognition*, *33*, 1047–1056.
- Van Hamme, L. J., & Wasserman, E. A. (1994). Cue competition in causality judgments: The role of non-presentation of compound stimulus elements. *Learning and Motivation*, *25*, 127–151.
- 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.
- Waldmann, M. E. (1996). Knowledge-based causal induction. In D. R. Shanks, K. J. Holyoak, & D. L. Medin (Eds.), *The psychology of learning and motivation: Vol. 34. Causal learning* (pp. 47–88). San Diego, CA: Academic Press.
- 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., & Hagmayer, Y. (2005). Seeing versus doing: Two models of accessing causal knowledge. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *31*, 216–227.
- Waldmann, M. R., & Holyoak, K. J. (1992). Predictive and diagnostic learning within causal models: Asymmetries in cue competition. *Journal of Experimental Psychology: General*, *121*, 222–36.
- Waldmann, M. R., & Martignon, L. (1998). A Bayesian network model of causal learning. In M. A. Gernsbacher & S. J. Derry (Eds.), *Proceedings of the Twentieth Annual Conference of the Cognitive Science Society*. Mahwah, NJ: Lawrence Erlbaum Associates, Inc.
- Waldmann, M. R., & Walker, J. M. (2005). Competence and performance in causal learning. *Learning & Behavior*, *33*, 211–229.
- Wason, P. C., & Johnson-Laird, P. N. (1972). *Psychology of reasoning: Structure and content*. Cambridge, MA: Harvard University Press.
- Wasserman, E. A. (1990). Detecting response–outcome relations: Toward an understanding of the causal texture of the environment. In G. H. Bower (Ed.), *The psychology of learning and motivation* (Vol. 26, pp. 27–82). San Diego, CA: Academic Press.
- Wasserman, E. A., & Berglan, L. R. (1998). Backward blocking and recovery from overshadowing in human causal judgment: The role of within-compound associations. *Quarterly Journal of Experimental Psychology*, *51B*, 121–138.
- Wasserman, E. A., Chatlosh, D. L., & Neunaber, D. J. (1983). Perception of causal relations in humans: Factors affecting judgments of response–outcome contingencies under free-operant procedures. *Learning and Motivation*, *14*, 406–432.
- Wasserman, E. A., Elek, S. M., Chatlosh, D. L., & Baker, A. G. (1993). Rating causal relations: Role of probability in judgments of response–outcome contingency. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *19*, 174–188.
- Wasserman, E. A., Kao, S.-F., Van Hamme, L. J., Katagiri, M., & Young, M. E. (1996). Causation and association. In D. R. Shanks, K. J. Holyoak, &

- D. L. Medin (Eds.), *The psychology of learning and motivation: Causal learning* (pp. 207–264). San Diego, CA: Academic Press.
- 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.
- Wills, A. J. (Ed.). (2005). *New directions in human associative learning*. Mahwah, NJ: Lawrence Erlbaum Associates, Inc.
- Wu, M., & Cheng, P. W. (1999). Why causation need not follow from statistical association: Boundary conditions for the evaluation of generative and preventive causal powers. *Psychological Science*, *10*, 92–97.