
How One Cause Discounts or Augments Another: A Connectionist Account of Causal Competition

Frank Van Overwalle

Dirk Van Rooy

Vrije Universiteit Brussel, Belgium

The authors investigated the degree of discounting and augmentation of a target cause by an alternative cause given a varying number of observations on the alternative cause while holding its degree of covariation constant. Two experiments showed that more observations of the alternative cause resulted in greater discounting or augmentation of a target cause. This sample size effect cannot be explained by current attribution theories based on statistical notions or belief updating but can be accounted for by a connectionist framework. In addition, the authors found that the sample size effect was stronger when the information was presented in a sequential trial-by-trial format as opposed to a summarized format but found no effect of information order. Possible extensions of statistical models with confidence weights that take account of sample size were considered and simulated but none accommodated the data as well as connectionist models.

How does a perceiver learn which one among multiple factors caused an event? According to Kelley (1967, 1971), perceivers take into account not only how a potential cause covaries with the event but also how this cause competes with rival explanations. The principle of covariation prescribes that an event is attributed to a cause with which it covaries over time (Kelley, 1967). Despite the central place accorded to this principle in attribution theory, Kelley (1971) also argued that the covariation principle in itself is insufficient to explain how perceivers select between alternative causes. For instance, although both speeding and bad weather may covary with a car accident, perceivers often discredit or discount the effect of one cause in favor of the other. To account for competition of multiple possible causes that goes beyond the covariation principle, Kelley (1971) proposed two complementary principles of discounting and augmentation.

The discounting principle specifies that “the role of a given cause in producing a given effect is discounted if other plausible causes are also present” (Kelley, 1971, p. 8). One of the most common examples of discounting in social cognition is when internal attributions to the actor are discounted given evidence on the potent influence of external pressures. The opposite tendency is described in the augmentation principle. This principle specifies that “if for a given effect, both a plausible inhibitory cause and a plausible facilitatory cause are present, the role of the facilitative cause in producing the effect will be judged greater” (Kelley, 1971, p. 12). For instance, a person’s success will be more strongly attributed to internal capacities when the task was hard rather than easy (for an overview, see McClure, 1998).

Although Kelley (1971) initially reserved the terms *discounting* and *augmentation* to describe competition between alternative explanations based on prior knowledge and minimal causal information (see also Morris & Larrick, 1995; Read & Montoya, 1999), other researchers applied these terms more broadly to denote causal competition also during causal induction, that is, in conjunction with covariation information (e.g., Hansen & Hall, 1985; Kruglanski, Schwartz, Maides, & Hamel, 1978; Rosenfield & Stephan, 1977; Van Overwalle & Van Rooy, 1998; Wells & Ronis, 1982). In line with this latter approach, the present studies focus on the effect of dis-

Authors’ Note: This research was in part supported by the Belgian Foundation of Scientific Research–Flanders (F.W.O.) under Grants 8.0192.95 and G.0128.97. We are grateful to Bert Timmermans for his comments on an earlier version of this article. Address correspondence to Frank Van Overwalle, Department of Psychology, Vrije Universiteit Brussel, Pleinlaan 2, B-1050 Brussels, Belgium; e-mail: frank.vanoverwalle@vub.ac.be.

PSPB, Vol. 27 No. 12, December 2001 1613-1626

© 2001 by the Society for Personality and Social Psychology, Inc.

counting and augmentation during covariation detection and causal induction.

Specifically, the aim of the present article is to evaluate current attribution models that provide computational accounts of the principles of discounting and augmentation in relation to the principle of covariation. Some of these computational theories are based on statistical principles (Anderson & Sheu, 1995; Cheng & Holyoak, 1995; Fales & Wasserman, 1992; Försterling, 1989), others on an anchoring and adjustment analogy (Busemeyer, 1991; Hogarth & Einhorn, 1992), whereas still others are inspired on associative or connectionist principles (Read & Montoya, 1999; Van Overwalle, 1998; Van Overwalle & Van Rooy, 1998).

To test the validity of these models, we induced changes in the discounting and augmentation of a target cause by varying only the number of observations of the competing (or alternative) cause while keeping its degree of covariation constant. Research has amply documented that people make more pronounced causal judgments after receiving more covariation information (Baker, Berber, & Vallée-Tourangeau, 1989; Försterling, 1992; Shanks, 1985, 1987, 1995; Shanks, Lopez, Darby, & Dickinson, 1996; Van Overwalle & Van Rooy, 2000). The specific hypothesis tested in this article is that as the number of observations for a competing explanation goes up, a target explanation will be more discounted or augmented (Hansen & Hall, 1985; Wells & Ronis, 1982). For instance, when there is growing evidence indicating that a successful task was easy, then the person's ability will be more discounted. Conversely, the greater the evidence that the task was hard, the more the person's ability will be augmented.

Why do we test a hypothesis that seems so intuitively plausible and straightforward? The reason is that, surprisingly, all prominent attribution models in social psychology that use a statistical formulation (e.g., Anderson & Sheu, 1995; Busemeyer, 1991; Cheng & Holyoak, 1995; Fales & Wasserman, 1992; Försterling, 1989; Hogarth & Einhorn, 1992) actually contradict this hypothesis and predict no variation in competition due to sample size differences. Only the connectionist approach makes this prediction (Read & Montoya, 1999; Van Overwalle, 1998; Van Overwalle & Van Rooy, 1998). This has far-reaching implications on how we view the causal attribution process. In contrast to statistical models that describe attribution as a complex and laborious application of statistical rules, connectionist models assume that causality is often an implicit process outside awareness and that only the outcome of that process is open to consciousness (Smith & DeCoster, 1999). This latter view seems to describe much better how causes are learned and understood in the hustle of everyday social life. Before turning to this connectionist framework, we

begin by explaining briefly why existing statistical and adjustment models fail to predict variations due to sample size.

STATISTICAL MODELS

The majority of attribution models in social psychology are rooted in the idea that people are intuitive statisticians who apply a kind of statistical rule to estimate causality. One of the most popular rules was formalized in the probabilistic contrast model of Cheng and Novick (1990, 1992; Cheng & Holyoak, 1995) and "requires that people . . . estimate and compare proportions" (p. 549). As pointed out by Cheng and Novick (1992), the probabilistic contrast approach is computationally superior to similar statistical formulations that make anomalous predictions in some cases (Einhorn & Hogarth, 1986; Morris & Larrick, 1995; Schustack & Sternberg, 1981). The model accounts for discounting and augmentation by assuming that probabilities are calculated conditional on the presence or absence of the competing cause (Cheng & Holyoak, 1995). This is analogous to experimental designs, where extraneous variables are kept constant across conditions. Other valuable models that can account for discounting and augmentation in addition to covariation are Försterling's (1992) formulation of Kelley's ANOVA model that is based on the statistical ANOVA parameter of effect size (i.e., η^2), and the Bayesian model proposed by Fales and Wasserman (1992).

However, none of these statistical models can account for differences in discounting and augmentation when the degree of covariation is kept constant and only the number of observations for the alternative cause is varied. The main reason is that these models are all based on proportions between frequencies of the causes and the effects rather than on the raw frequencies themselves. Consequently, because the degree of covariation is kept constant and, in addition, because some critical cause-effect frequencies are set to zero to induce straightforward discounting or augmentation, these proportions remain unchanged and thus so do the predicted causal estimates (see Appendix A for a mathematical demonstration). Thus, although these statistical models may adequately describe attribution judgments at asymptote (i.e., given a sufficient large size), they clearly fail to account for the gradual increase or decrease in judgments while learning is still going on and asymptote has not been reached.

JUDGMENT UPDATING MODELS

Recently, a number of models have been proposed that avoid these limitations of earlier statistical models by incorporating some form of anchoring and updating

rule (the step-by-step belief-adjustment model of Hogarth & Einhorn, 1992; the serial averaging strategy of Busemeyer, 1991), which makes them sensitive to sample size. However, a serious restriction of these proposals is that the adjustment rules involve only a single cause and do not take into account the influence of alternative causes, which is necessary to predict discounting and augmentation. Therefore, we will ignore these models and immediately turn to the connectionist approach.

CONNECTIONIST MODELS

Recently, several authors moved away from statistical approaches and proposed an entirely different, connectionist framework that is inspired by basic neurological properties of the human brain (Read & Montoya, 1999; Smith, 1996; Van Overwalle, 1998; see also Allan, 1993; McClelland & Rumelhart, 1988; Shanks, 1995). We claim that unlike earlier models, adaptive connectionist models with an error-correcting learning algorithm predict differences in discounting and augmentation given changes in sample size of the competing cause.

To explain the properties of adaptive networks, we focus on one of the simplest architectures: the "feedforward" model. In the feedforward model, causal estimates are represented by the weight of the connections between input nodes representing causes and an output node representing the effect. Activation in the network typically runs from causes to effect (hence the name feedforward). Whenever a cause is present, its input node is activated and this activation is then automatically propagated to the output node in proportion to the weight of the connection (i.e., the causal strength at that moment). All input activations received at the output node are linearly summed to determine the output activation, which reflects the effect predicted by the network on the basis of the causal input given.

A key feature of adaptive connectionist networks is that the connection weights are adapted in response to information on new co-occurrences between causes and effect, using a learning algorithm. The learning algorithm we focus on here is the delta algorithm (McClelland & Rumelhart, 1988). The goal of this algorithm is to reduce the error between the mental representation of the effect as predicted by (the output node of) the network and the actual effect. This error is reduced by adjusting the weights of the cause-effect connections. When the occurrence of the effect is underestimated, the weights are adjusted upward; when the occurrence of the effect is overestimated, adjustments are made downward. Thus, as Dennett (1998) noted, this learning algorithm mimics important aspects of reasoning and intelligence, because "intelligence is . . . for improving the fidelity of your expectations about what is going to happen next" (p. 184).

The delta learning algorithm has been applied in many investigations on human categorization and causality (for reviews, see Allan, 1993; Shanks, 1995; Van Overwalle & Van Rooy, 1998) and is formally identical to Rescorla and Wagner's (1972) associative model. Although it does not tally frequencies or compute statistical probabilities, it forces the weights to converge to the probabilistic norm (Cheng & Novick, 1990) after a sufficient number of observations (Chapman & Robbins, 1990; Van Overwalle, 1996). Thus, the feedforward network respects the statistical principle of covariation.

Most important, the delta algorithm can easily account for our manipulation, in which the sample size of the competing cause is increased. The reason is that the network's output is computed on the summed activation of all causes present. Thus, adjustments are driven by competition for predictive strength. To illustrate, Figure 1 depicts the weights of two tennis partners, Theo (the target explanation T) and Xavier (the alternative explanation X), who together win a doubles game.

The top panel illustrates discounting. Consider a network that is provided with information that Xavier won either five single games (large size) or one single game (small size). As can be seen, after five games, Xavier's weight is much stronger than after one game. Next, when the network learns the new information that Xavier and Theo win their double game, competition between the two explanations arises. If Xavier takes already a great part of the available weight, there is less weight left for Theo to gain. This results in stronger discounting of Theo's weight than if Xavier takes a smaller part of the available weight.

The bottom panel illustrates the reverse mechanism of augmentation, where the network receives the same information, except that Xavier first lost five or one single game so that he acquired inhibitory weight. As can be seen, to reduce the inhibitory impact of Xavier in winning the double game, augmentation of Theo's contribution is much stronger when Xavier lost five times rather than only once.

DESIGN AND HYPOTHESES

The design of the experiments reported here was modeled after the example illustrated above. In a first experiment, we induced competition between two causal explanations (e.g., the contribution of Theo versus Xavier in winning a game) and manipulated the strength of the competing cause by varying how often this cause was presented alone: either one time (small size) or five times (large size). Thus, for instance, participants learned that Theo's tennis partner, Xavier, won five single games in the large size condition as opposed to only one game in the small size condition. However, this manipulation confounds size with consistency infor-

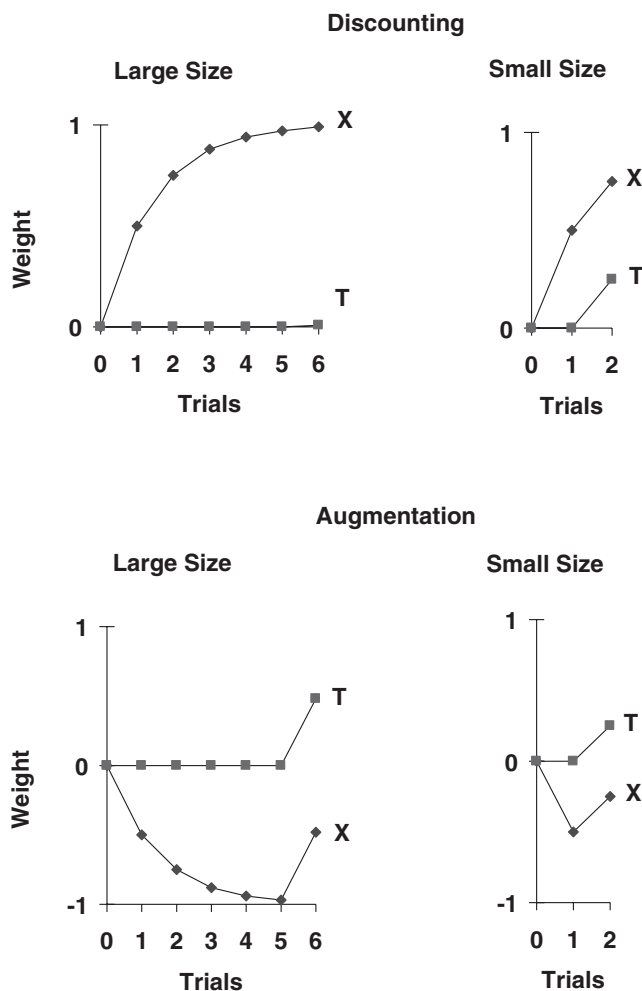


Figure 1 A feedforward simulation of competition between a target cause T (e.g., Theo) and a competing cause X (e.g., Xavier), given positive (+) or negative (-) outcomes.

NOTE: For discounting, there were five (large size) or one (small size) X+ trials, indicating that the competing cause (Xavier) won his single games five times or one time, respectively. This was followed by one compound TX+ trial, indicating that Theo and Xavier together won their doubles game. For augmentation, there were five (large size) or one (small size) X- trials, indicating that the competing cause (Xavier) lost one or five single game(s), followed by the same compound TX+ trial, indicating that Theo and Xavier together won their doubles game. The figure illustrates forward competition. In backward competition, the compound TX+ trial, indicating that both players won their doubles game, comes first.

mation (or covariation with time, see Kelley, 1967). For instance, the information that Xavier won five times also increases the consistency of Xavier's performance on different occasions. To assess whether the size effect is not solely due to increased consistency, in a second experiment, increasing size was manipulated by proving information that the target cause competed against many alternative actors or stimuli (each on a single occasion) rather than against one alternative on several occasions

(as in the first experiment). This should rule out consistency information as an alternative explanation for the size effect. Consistent with the connectionist framework, we predict for the two experiments that a large as opposed to small size will result in a greater facilitatory or inhibitory strength of the competing cause, leading to, respectively, more discounting and more augmentation of the target cause.

We also explored two additional questions. One question was whether the order in which the competing and target causes are presented would influence the predicted size effects. The competing cause was either presented before the target (forward order) or after the target (backward order). In associative research with animals (where discounting and augmentation are analogous to blocking and superconditioning, respectively), the typical finding is that only forward competition effects occur (Kamin, 1968; Rescorla, 1969). However, more recent studies with humans indicate that forward and backward competition are generally equally effective (Shanks, 1985; Van Hamme, 1994; Williams & Docking, 1995; Williams, Sagness, & McPhee, 1994). Therefore, we expect that information order will have little effect on our size manipulation.

Another exploratory question was whether our predictions would apply on different presentation formats. In associative research, covariation information is typically presented in a sequential trial-by-trial format in which each occurrence of one or more causes is described in short sentences at successive trials. This format seems to reflect people's incidental learning during everyday life and also is consistent with a connectionist approach where input information is received and processed on-line (i.e., after each piece of information is received). Conversely, in social research, covariation information is typically presented in a prepacked format where the occurrences of causes are summarized in one single sentence. This format seems to capture people's verbal interactions with one another but is less suitable for sequential processing as required by connectionist models. Consequently, we anticipated that the size effect would be stronger in a sequential format than in a summary format.

EXPERIMENT 1: COMPETING AGAINST ONE ALTERNATIVE

Method

PARTICIPANTS

Participants were 106 male and female students from the Vrije Universiteit Brussel who participated for a partial course requirement. They were tested in groups of one to five. Approximately half of the participants

received the sequential format ($n = 51$), whereas the other half received the summary format ($n = 55$).

MATERIAL

The overall design of the experiment involved one between-subjects Format factor (sequential or summarized) and three within-subject factors, including Order (forward or backward), Type (discounting or augmentation), and Size (small or large). We created two stories for each combination of the three within-subjects factors, resulting in 16 stories. The stories were counterbalanced across participants between Size and Order. The target and competing causes involved either actors or stimuli (objects or persons). To make sure that the actors and stimuli would be seen as causally related to the effect (which is necessary for discounting and augmentation to occur, see McClure, 1998), the stories involving actors used action verbs (which tend to imply the actor as the cause) and the stories involving stimuli used state verbs (which tend to imply the stimulus as the cause; see Rudolph & Försterling, 1997).

Size. Each story consisted of five trials (large size) or one trial (small size) in which a competing cause was present alone. In addition and regardless of size, there were another five compound trials in which both the target and competing cause were present together.

Order. In the forward order, the trials describing the sole presence of the competing cause were given first, whereas in the backward order, these competing trials were given last.

Type. To induce discounting, both target and competing trials were followed by the same outcome. In contrast, to induce augmentation, the outcome of the competing cause alone was opposite to the focal outcome when also the target cause was present. To make sure that the participants would encode this opposite information correctly, the semantic negation of an outcome was always indicated in capitals.

The manipulation of these three within-subjects factors is illustrated in the next example with “An” as the discounted target actor and “Elena” as the discounting competing actor. The example illustrates forward order (i.e., competing cause first), with the large size phrasing given between straight brackets. This gives for the summary presentation format:

Elena passed the first [five] selection round[s] in single scull, and An and Elena passed the five selection rounds in double scull.

The sequential format of this example is illustrated with 10 consecutive trials between slashes. The trials between parentheses were omitted for the small size:

Elena passed the first [/ second / third / fourth / fifth] selection round in single scull;

An and Elena passed the first / second / third / fourth / fifth selection round in double scull.

Similarly, the augmentation manipulation of a stimulus is illustrated below in a backward order (i.e., competing cause last) given a summary format, with “mackerel” as the augmented target stimulus and “salmon” as the augmenting competing stimulus:

Ella and four other women felt sick after eating mackerel and salmon but five [one] other women [woman] did NOT feel sick after eating salmon.

PROCEDURE

Participants were seated in front of an IBM-compatible PC and the experiment was monitored by MEL software. Instructions appeared on the screen and the use of the rating scale was practiced. Participants read 16 stories, which appeared in a different random order for each subject. In the summary format, each story was presented during one single trial. In the sequential format, each story was presented during several consecutive trials. The computer randomized for each subject the order in which specific actors or stimuli appeared, with the provision that forward and backward order was not affected.

After reading each story, participants had to rate the causal influence of the target factor and the competing factor: *something special about [actor or stimulus]*. They rated the causal influence of each factor on the outcome using an 11-point rating scale ranging from 0 (*absolutely no influence*) to 100 (*very strong influence*), with midpoint 50 (*partial influence*). This question phrasing is standard in most attribution research on the influence of covariation information (e.g., Cheng & Novick, 1990; Försterling, 1989). For example, in the story with An and Elena as actors, participants rated the causal influence of *something special about An* and *something special about Elena*. Similarly, in the story with mackerel and salmon as stimuli, participants rated the causal influence of *something special about mackerel* and *something special about salmon*.

Results

Because we made predictions irrespective of whether the causes involved an actor or stimulus, the ratings were collapsed across actor and stimulus. This was justified by a multivariate analysis of variance (MANOVA) with Measure (actor or stimulus), Order (forward or backward), Type (discounting or augmentation), and Size (small or large), which revealed that Measure did not interact with Size or with any higher order interaction including Size. However, because we found the expected difference between presentation formats, as the triple interaction




Figure 2 Experiment 1: Causal ratings in function of size, type, and format.

NOTE: Black data points denote a significant size effect ($p < .01$).

between Format, Type, and Size was significant for the target and competing ratings, $F(1, 104) = 6.58$ – 17.44 , $p < .05$, we will report these results separately.

SEQUENTIAL FORMAT

Figure 2 (left panel) depicts the average ratings for the target and competing causes. We begin with the competing cause. Consistent with our sample size hypothesis, we found a significant interaction between Type and Size, $F(1, 50) = 30.31$, $p < .0001$. Simple effect tests revealed that in the discounting condition, the (facilitatory) competing cause acquired more strength when size was large rather than small, $F(1, 50) = 22.03$, $p < .0001$. The predicted reversed pattern was found in the augmentation condition, where the (inhibitory) competing cause acquired less strength when size was large rather than small, $F(1, 50) = 11.74$, $p < .01$.

Our critical hypothesis concerned the target cause, which should be more discounted and more augmented given a larger sample size. This prediction also was supported because the critical interaction between Type and Size was significant, $F(1, 50) = 29.76$, $p < .0001$. Simple effect tests revealed that when size was large as opposed to small, the target cause was more discounted, $F(1, 50) = 13.66$, $p < .001$, and more augmented, $F(1, 50) = 11.57$, $p < .01$.¹ In addition, our prediction that order would have little effect on our size manipulation was confirmed, because the interaction between Order and Size was not significant, $F < 1$.

SUMMARY FORMAT

Figure 2 (right panel) depicts the average ratings for the target and competing ratings. For the competing rating, the expected interaction between Type and Size approached significance, $F(1, 54) = 3.99$, $p = .051$. Simple effect tests revealed that although the size manipulation led to the predicted increase of the (facilitatory) competing cause for discounting, $F(1, 54) = 13.19$, $p < .001$, it failed to show any change for augmentation, $F < 1$, *ns*. More important, for the target cause, neither the expected interaction nor any of the simple effect tests reached significance, $F_s < 1$, *ns*. Thus, contrary to our hypothesis, size did not increase discounting or augmentation of the target cause. Consistent with our prediction, however, the interaction between Order and Size again failed to reach significance for the target cause, $F < 1$.

Discussion

The most important finding of this experiment is that, at least in the sequential format, the sample size of the competing cause substantially influenced the perceived causality of the target cause. As more evidence on the competing cause was accumulated, the target cause was more discounted or augmented. Overall, these results lend some support to a connectionist account of causal competition. However, the size effect did not appear in the summary format. This is consistent with the connectionist conception that information uptake and encoding is an incremental process that works best on the basis of trial-by-trial information.

In addition, we found that the order in which the competing cause was presented (before or after the target causes) did not influence this size effect. This suggests that most often both forward and backward competition are effective in social explanation, consistent with associative research on human learning (e.g., Shanks, 1985; Van Hamme, 1994; Williams & Docking, 1995; Williams et al., 1994).

EXPERIMENT 2: COMPETING
AGAINST MANY ALTERNATIVES

A limitation of the first experiment, mentioned earlier, is that by increasing the number of observations of the competing cause, we also increased its perceived consistency (or covariation with time). Although Försterling (1989) actually documented that high consistency decreased attributions to the actor or the stimulus (whereas our results show the opposite pattern), we wanted to rule out consistency as an explanation for the results in the first experiment. Therefore, in a second experiment, the target did not compete against one single actor or stimulus but rather against one (small size) or five different actors or stimuli (large size). Because individual actors and stimuli appear only once in this design, consistency is kept constant and thus ruled out as a potential alternative explanation. Moreover, by competing against many alternatives rather than a single alternative repeated several times, the predicted competition effects might be stronger in this experiment.

We make essentially the same predictions as in the previous experiment. To capture the idea of competition by several alternative actors or stimuli, we borrowed Weiner's (1985) terminology and referred to them as the *external* or *general* context, respectively (see also Van Overwalle, 1997). Thus, an explanation in terms of a target actor must compete against an alternative explanation in terms of an external context. For instance, we expect that people will attribute high sales figures to a (target) clerk's capacities when there is only one other clerk with equally high sales. Conversely, when there are five other clerks with high sales, attributions to alternative causes such as external conditions (e.g., convenient location of the shop) are more likely. We make the same prediction for target explanations in terms of a specific stimulus (e.g., salmon) as opposed to competing explanations in terms of a general context (e.g., a restaurant's food quality).

Method

PARTICIPANTS

Participants were 115 male and female students from the Vrije Universiteit Brussel who participated for a partial course requirement. They were tested in groups of one to five. Approximately half of the participants received the sequential presentation format ($n = 64$) and the other half received the summary format ($n = 51$).

MATERIAL AND PROCEDURE

The overall design and procedure was similar to the previous experiment, with the following modifications. We developed 16 novel stories in which the competing causes involved different comparison actors or stimuli rather than a single competing actor or stimulus. To

make things easier, the focal event (in which the target is present, together with the external or general context) was presented only once.

The next example illustrates the discounting manipulation of Annie as target actor in a forward order (with sentence parts describing the large size between straight brackets). For the summary and sequential format respectively, this gives the following:

One [Five] other salesgirl[s] and Annie attained high sales figures for perfumes
Sandra [/ Daniela / Katharina / Sabrina / Katrien] /
Annie attained high sales figures for perfumes

Likewise, the augmentation manipulation of Danny as target stimulus is illustrated below in a backward order:

Jos respected his peer student Danny but he did NOT respect one [five] other peer student[s].
Jos respected his peer student, Danny;
Jos did NOT respect his peer student, Luis [/ Gaston / Ruben / Ricardo / David]

To assess the causal influence of the target causes, similar to Experiment 1, participants rated the causal influence of *something special about [actor or stimulus]*. To assess the causal influence of the external or general context, they rated the causal influence of *something external (outside [actor])* or *something general (plays a role with many [stimulus category])*.

Results

The same preliminary MANOVA was performed as in the previous experiment and again indicated that the ratings of the actor and stimulus could be safely collapsed.² Although there were no significant interactions between Format and Size, $F_s(1, 113) < 2.25$, *ns*, to remain consistent with the previous experiment, we will analyze the presentation formats separately.

SEQUENTIAL FORMAT

Figure 3 (left panel) depicts the average context and target ratings. As can be seen, the results depict the same predicted pattern as in the sequential format of the previous experiment. This was confirmed by the significant interaction between Type and Size on the target and context ratings, $F(1, 63) = 28.66\text{--}43.06$, $p < .0001$. Simple effect tests showed that given a larger size, the competing context was rated stronger in the discounting condition, $F(1, 63) = 10.35$, $p < .0001$, and weaker in the augmentation condition, $F(1, 63) = 26.33$, $p < .0001$. More important, given a larger size, the target cause was more discounted, $F(1, 63) = 20.93$, $p < .01$, and more augmented, $F(1, 63) = 30.32$, $p < .0001$. In addition, as predicted,




Figure 3 Experiment 2: Causal ratings in function of size, type, and format.

NOTE: Black data points denote a significant size effect ($p < .01$).

there was no significant interaction between Order and Size on the target, $F < 1$.

SUMMARY FORMAT

Figure 3 (right panel) depicts the average target and context ratings. As we had hoped for, the hypothesized size effect now also appeared in the summary presentation format, although competition was generally weaker than in the sequential format (as can be seen, for instance, from the ratings of the discounting and augmentation conditions that did not differ given a small size). This was confirmed by a significant Type \times Size interaction on the target and context ratings, $F(1, 50) = 22.60$ – 27.54 , $p < .0001$. Simple effect tests showed that given a larger size, the context was rated higher given dis-

counting, $F(1, 50) = 12.64$, $p < .001$, and lower given augmentation, $F(1, 50) = 11.15$, $p < .01$. More important, given a larger size, the target was more discounted, $F(1, 50) = 32.06$, $p < .0001$, and more augmented, $F(1, 50) = 9.91$, $p < .01$. As predicted, there was again no significant interaction between Order and Size on the target rating, $F < 1$.

Discussion

The findings of this experiment confirmed that as the number of the comparison cases increased, the target cause was more discounted or augmented. Because the frequency of each individual comparison cause was kept constant, consistency is ruled out as an explanation for these results. Similar to the previous experiment, there was no effect of order.

Perhaps the more interesting finding of this experiment is that this sample size effect was now also observed in the summary format. The most plausible explanation for the more robust size effect in the summary format is that the target competed against five different actors or stimuli rather than only one repeated over five trials as in the first experiment. Alternatively, it is also possible that the use of broad contextual categories to measure the comparison causes may have allowed for a clearer separation between target and competing causes, leading to increased competition and a more reliable effect of sample size.

MODEL SIMULATIONS

The present experiments demonstrate that there was more discounting and augmentation after increasing the number of competing cases. To evaluate how closely a connectionist formulation can predict our data, we ran a feedforward simulation and correlated the simulated results with the observed attribution ratings. To incorporate the finding that both forward and backward competition are effective, we implemented Van Hamme and Wasserman's (1994) suggestion that absent causes are encoded with a negative activation (for more details on the model specification, see Appendix B).

As a way of comparison, we did the same for the most prominent statistical models in social psychology: the probabilistic contrast model (Cheng & Holyoak, 1995) and the ANOVA model (Försterling, 1992). However, it would be of little interest to reiterate the fact that these models fail to be sensitive to sample size. Therefore, we extended these existing models with additional parameters that take into account sample size to evaluate whether these extensions would be sufficient to account for the observed size effect.

Specifically, we followed the suggestion by Cheng and Holyoak (1995) that "confidence in the assessment of a contrast is presumed to increase monotonically with the

number of cases observed” (p. 273). That is, we weighted the major theoretical variables of the models (conditional probabilities: Cheng & Novick, 1990; sum of squares: Försterling, 1989) in proportion to the number of observations available. We allowed two such different confidence weights for frequencies that involved the presence and absence of the target cause (ω_t and ω_x , respectively). This procedure is identical to the one used recently by Lober and Shanks (2000, p. 207) and parallels that of the connectionist models, where we also allowed two different learning rate parameters for target and comparison factors (ϵ_t and ϵ_x , respectively). More technical details on the model specifications are given in Appendix B.

Method

The models were run using exactly the same order of trials and blocks as in the experiments. The connectionist model was updated after each trial. We sought the best overall fitting parameters of each model by searching for the maximum correlation between simulated and observed data across all conditions, given all admissible parameter values (see Gluck & Bower, 1988; Nosofsky, Kruschke, & McKinley, 1992). We did not attempt to find common best-fit parameters for all data sets because the presentation formats and stories differed too much between conditions and experiments. The best-fit parameter values were generally quite robust, and small deviations of 0.10 in the values decreased the maximum correlations only minimally.

To evaluate the performance of the models with respect to stronger discounting and augmentation given an increased sample size, we then computed separate correlations between observed and simulated data within each discounting and augmentation condition (using the same overall best-fitting model parameters). Next, we averaged these correlations across target and competing causes as well as across forward and backward order. Table 1 depicts the results. An average correlation of +1 reflects the predicted size effect in all cases, a correlation of 0 indicates that the sample size effect is absent in all cases, and a correlation of -1 indicates that the predicted sample size effect is reversed in all cases.

Results and Discussion

As can be seen in Table 1, the fit for discounting and augmentation was generally poor for the statistical models because they showed many zero correlations indicating that they failed to predict the observed sample size effect. To illustrate, the probabilistic contrast model predicts that, for the present augmentation design, the probability of a target cause is always 1 and for a competing cause always 0 (see Equation A4), so that an increasing confidence weight has no effect at all on these esti-

TABLE 1: Fits of the Models to the Data

Model	Experiment 1		Experiment 2	
	Seq	Sum	Seq	Sum
Weighted ANOVA				
Discounting	.000	.000	.000	.000
Augmentation	.000	.500	.000	.000
	$\omega_t > .00$ $\omega_x = .00$	$\omega_t > .06$ $\omega_x = .70$	$\omega_t > .00$ $\omega_x = .00$	$\omega_t > .00$ $\omega_x = .00$
Weighted probabilistic				
Discounting	1.000	1.000	1.000	1.000
Augmentation	.000	.000	.000	.000
	$\omega_t = .98$ $\omega_x = .72$	$\omega_t = .20$ $\omega_x = .82$	$\omega_t = 1.00$ $\omega_x = .60$	$\omega_t = .60$ $\omega_x = .10$
Feedforward connectionist				
Discounting	1.000	1.000	1.000	1.000
Augmentation	1.000	-.500	1.000	1.000
	$\epsilon_t = .95$ $\epsilon_x = .50$	$\epsilon_t = .81$ $\epsilon_x = .72$	$\epsilon_t = .42$ $\epsilon_x = .20$	$\epsilon_t = .68$ $\epsilon_x = .20$

NOTE: Cell entries are correlations. Seq = sequential format, sum = summarized format, ω_t = weight for target factors, ω_x = weight for competing factor, ϵ_t = learning rate for target factors, and ϵ_x = learning rate for competing factors. The activation level for absent causes in the feedforward model was set to $\alpha_n = -.5$.

mates, resulting in correlations of 0. In contrast, the connectionist model was capable of simulating most observed size effects. The correlations were all +1 (except in the summary format of Experiment 1 where some conditions showed a nonsignificant opposite size effect). To emphasize that the obtained results are not due to the feedforward architecture used here but rather to the delta learning mechanism, we also conducted simulations with a recurrent network (Read & Montoya, 1999).³ As would be expected, the results for the discounting and augmentation conditions were exactly the same as for the feedforward network.

We explored other solutions to improve the performance of the statistical models, but they all failed. Measures such as the *F* value or the *F* test in Försterling’s ANOVA model that incorporate the number of observations (i.e., by the degrees of freedom) fail because in our discounting and augmentation designs, the error variance was always zero (see Appendix A). Other possibilities are to add a sort of anchoring and adjustment notion to statistical models. This is exactly what Hogarth and Einhorn (1992) proposed in their step-by-step belief-adjustment model and Busemeyer (1991) in his serial averaging strategy. However, the proposed adjustment rules are mathematically identical to a simplified version of the delta algorithm, one that deals with only one cause at a time (see Wasserman, Kao, Van Hamme, Katagiri, & Young, 1996). Hence, these extensions cannot account for competition between alternative explanations. Interestingly, if one would extend these adjustment rules to

include more causes, they would become identical to the delta algorithm of connectionist models.

GENERAL DISCUSSION

The general finding of the present experiments was that increasing the sample size of a competing cause reliably increased discounting and augmentation of a target cause. This provides support for adaptive connectionist models with an error-correcting delta algorithm because they were the only models that made this prediction. These connectionist models suggest that causal learning on the basis of relevant information is an on-line incremental process in which estimates of causality are continuously updated on-line.

Our experiments have, however, a shortcoming in that they did not include a baseline condition in which competition by an alternative explanation was absent. Hence, we cannot be sure in the small size conditions whether differences in causal estimates were due to discounting or augmentation, or both. However, it is difficult to escape the interpretation that discounting and augmentation took place in the large size conditions where the difference between estimates became even greater.

The experiments further confirmed that forward and backward competition are equally effective, in line with earlier associative learning research (Shanks, 1985; Van Hamme, 1994; Williams & Docking, 1995; Williams et al., 1994). In addition, they demonstrated that information provided in a summary sentence has much weaker size and competitive effects. A possible explanation is that summary information is often encoded more shallowly than sequential information and so leads to weaker effects. Alternatively, according to a connectionist view that allows only for trial-by-trial input, summary information must be converted for sequential input before it can be appropriately encoded. Perhaps this conversion is far from perfect and so results in less sample size and competition effects. We return to this issue later in the discussion.

SAMPLE SIZE ADJUSTMENTS IN STATISTICAL MODELS?

Because statistical models do not incorporate a learning mechanism that allows adjusting causal estimates when novel information comes in (Anderson & Sheu, 1995; Cheng & Holyoak, 1995; Fales & Wasserman, 1992; Försterling, 1989), sample size effects are particularly problematic for these models. In the simulations, we demonstrated that many attempts to extend these models with supplementary theoretical notions (e.g., adding confidence weights or an anchoring-and-adjustment mechanism) failed to salvage these models from this limitation.

One might argue that sample size does not so much affect causal judgments but rather the confidence with which these judgments are given. In line with this reasoning, one might suggest that facilitatory information increases confidence over trials, whereas inhibitory information decreases confidence. However, this argument simply shifts the burden of proof from causal judgments to confidence judgments. Moreover, it does not explain why the reverse confidence effects should occur for the discounted and augmented causes. Cheng (1997) provided an account of why perceivers become increasingly uncertain of the causal status of the discounted cause, but her reasoning cannot explain why they should become more certain in the case of augmentation. We see no way to solve these logical quandaries of the confidence notion within the boundaries of the existing statistical theories. If our participants had been given the possibility to express their subjective confidence in the causal ratings, perhaps this might have solved some questions, but it would have left unanswered most of the reservations that we raised.

Finally, it might be argued that the scope of statistical models is explaining judgments only at the asymptotic level (What is computed?) rather than at a pre-asymptotic level that keeps track of the development of judgment (How is it computed?). However, the fact that the connectionist approach accommodates both levels makes that argument questionable. Therefore, we tend to conclude that connectionist models may complement or even replace these earlier models by providing a low-level description of the attribution process.

IMPLICATIONS FOR SOCIAL EXPLANATION

What are the implications of a connectionist perspective in social explanation? We believe that the most important contribution is theoretical parsimony. The delta error-correcting algorithm provides a unifying mechanism that incorporates not only Kelley's (1967, 1971) principles of covariation and competition but also the statistical principle of sample size. Moreover, a connectionist model not only explains how people incrementally adjust their causal estimates (cf. delta-algorithm) but can also address how people make a quick causal judgment in situations where they do not receive a stream of new causal information. In such cases, contextual cues about the appropriate place, time, and content spread their activation automatically to a host of potential causal candidates in memory. The cause that is most strongly connected with the outcome (because of strong weights built up during prior learning) will then be selected as the most likely cause. Thus, a connectionist model is relevant for both causal learning from covariation information as well as for snapshot judgments that are based on memorized experiences

from the past. For instance, after learning that Xavier won five tennis tournaments, there is a strong connection in memory between Xavier and success in tennis. Later judgment can invoke this memory trace to make predictions on Xavier's future successes and that of his opponents.

A connectionist approach also has applied value. First, it assumes that activation spreading and weight adjustment occur at a relatively implicit and autonomous level and that only the outcome of these processes is available for conscious introspection (Smith & DeCoster, 1999). Hence, people do not need to learn complex and elaborate rules for causal induction like statistical approaches assume. Rather, they infer causality quite often even when they are not fully aware of it, much like young infants and children learn how to understand and predict their environment and how to control it. A second important implication is that biases such as sample size should not be viewed as errors of the mind but rather as a window on how the mind works. Whereas statistical models often see biases as errors against logical norms, connectionist models actually see them as illustrations of human information processing. For instance, a connectionist approach can explain other well-known biases and sample size effects such as illusory correlation, group homogeneity, and group polarization and the impact of increasing or alternative information on impression formation and attitude change (Van Overwalle, Labiouse, & French, 2000).

One question that remains largely unsolved, however, is the relationship between simple connectionist adjustments and higher level "cognitive" reasoning and language use that is typical of social interaction. In this respect, it is quite interesting that the data from the summary format generally paralleled those of the sequential format (although in a weaker form and not always for all effects). Shanks (1991; Lober & Shanks, 2000) argued that people's experience with causal acquisition in real situations gave rise to causal metabeliefs, which may be readily applied in the interpretation of verbal summary sentences. Because these metabeliefs developed from observed situations, they may show similar but weakened effects of covariation, competition, and sample size at the surface. Alternatively, Van Overwalle and Van Rooy (1998) speculated that verbal summary information is unconsciously decomposed in the form of dummy exemplars or mental models (cf. Johnson-Laird, 1983), which are then sequentially analyzed by a connectionist processor. Occasional differences between trial-by-trial and summary formats are explained by an imperfect conversion of summary information, a type of error that has been amply documented in research on deductive reasoning. Which of these two proposed mechanisms actu-

ally underlies causal induction from verbal summaries is a question for future research.

APPENDIX A

This appendix demonstrates that existing statistical models are incapable of showing a size effect of the competing cause X on the target cause T. The information given to the participants can be represented by a standard contingency table with four relevant frequencies denoted by *a-d*: when the target is present and the outcome occurs (*a*) or not (*b*) and when the target is absent and the outcome occurs (*c*) or not (*d*).

A1. Probabilistic Contrast Formulation

According to the probabilistic contrast formulation, when two or more potential causes are present, the genuine causal influence of a target T is measured *conditional* on the presence or the absence of the other competing factor X (Cheng & Holyoak, 1995). Hence, the causal influence of a target factor T on the outcome O is measured by ΔP_T , or the contrast between two conditional proportions:

$$\Delta P_T = P(O|T.X) - P(O|\sim T.X), \quad (A1)$$

where O represents the outcome, T the target cause, X the competing cause, and a tilde denotes their absence. Because the competing factor X is always present in our designs, the ΔP formulation can be reduced to the following:

$$\Delta P_T = P(O|T) - P(O|\sim T) = [a/(a + b)] - [c/(c + d)]. \quad (A2)$$

In the designs of all experiments, $b = 0$, whereas for discounting, $d = 0$ and for augmentation, $c = 0$. Hence, in the case of discounting, we can simplify the ΔP formulation as follows:

$$\Delta P_T = [a/a] - [c/c] = 0, \quad (A3)$$

and in the case of augmentation:

$$\Delta P_T = [a/a] - [0/(0 + d)] = 1. \quad (A4)$$

This indicates that the discounted target will always attain zero causal strength, whereas the augmented target will always receive full causal strength regardless of the frequencies of the competing X. It can be easily shown that in more recent versions of probabilistic theory (Cheng, 1997), the same predictions are made for augmentation, whereas discounting is undetermined (i.e., division by zero).

A2. ANOVA Formulation

The ANOVA formulation defines causal strength as an analog to the effect size η^2 of a standard ANOVA (Försterling, 1992), which is given by the following:

$$\eta^2 = SS_{\text{between}} / SS_{\text{total}} = 1 - [SS_{\text{within}} / SS_{\text{total}}]. \quad (A5)$$

If the presence of the outcome is indicated by 1, and its absence by 0, then SS_{within} can be expressed as follows:

$$SS_{\text{within}} = a[a/(a+b) - 1]^2 + b[a/(a+b)]^2 + c[c/(c+d) - 1]^2 + d[c/(c+d)]^2. \quad (\text{A6})$$

For discounting where $b = d = 0$, this formulation can be simplified as follows:

$$SS_{\text{within}} = a[a/a - 1]^2 + 0 + c[c/c - 1]^2 + 0 = 0. \quad (\text{A7})$$

For augmentation where $b = c = 0$, this reduces to the following:

$$SS_{\text{within}} = a[a/a - 1]^2 + 0 + 0 + 0 = 0, \quad (\text{A8})$$

showing that in both cases, SS_{within} equals zero. In addition, SS_{total} can be written as follows:

$$SS_{\text{total}} = (a+c) [(a+c)/(a+b+c+d) - 1]^2 + (b+d) [(a+c)/(a+b+c+d)]^2. \quad (\text{A9})$$

It can be easily shown that for discounting, $SS_{\text{total}} = 0$, so that η^2 in Equation A5 is undetermined, whereas for augmentation, $SS_{\text{total}} \neq 0$, so that $\eta^2 = 1$, irrespective of the frequencies of the competing X.

A3. Bayesian Formulation

The argument that Bayesian reasoning is not influenced by sample size is based on the assumption—shared by most models—that perceivers evaluate to what extent a causal hypothesis can explain all or most of the available data. However, this assumption can be relaxed by assuming that perceivers simultaneously entertain a number of causal hypotheses with mutually exclusive degrees of causal probability—for instance, 11 hypotheses reflecting a probability of 0, 0.1, 0.2, 0.3, and so forth to 1. In this extended case, the model is sensitive to sample size (Fales & Wasserman, 1992). Because this assumption is psychologically very implausible as it puts a great burden on human information processing, and because it is not required by any of the other models discussed, this extension is not considered here.

APPENDIX B

This appendix discusses how the models were specified for the simulations and how the statistical models were extended to account for sample size. Note that for simulating Experiment 2, the comparison causes were represented by a single context factor that was assumed to be always present in the same way as the competing causes in Experiment 1. Hence, the causal structure was essentially identical for all experiments.

B1. Weighted Probabilistic Model

To make the probabilistic model sensitive to the number of observations, we weighted each of the conditional probabilities P (see Appendix A1) with a freely estimated proportion (between 0% and 100%) of the frequencies involved. That is, if ω_t denotes a proportion of the frequencies when cause T is present and if ω_x denotes a proportion of the frequencies when cause T is absent, then the strength of a target cause may be formalized as follows:

$$\Delta P_T = \omega_t P(O|T) - \omega_x P(O|\sim T), \quad (\text{B1})$$

and, likewise, the strength of a competing cause is as follows:

$$\Delta P_X = \omega_x P(O|\sim T). \quad (\text{B2})$$

B2. Weighted ANOVA Model

To make Försterling's (1989) model sensitive to the number of observations, we weighted η^2 in Appendix A2 with a freely estimated proportion of the frequencies involved, that is, with the same ω_t and ω_x as defined above. Thus, for the target cause, this becomes as follows:

$$\eta_T^2 = [\omega_t SS_{\text{between}}] / [(\omega_t + \omega_x) SS_{\text{total}}] \quad (\text{B3})$$

and, similarly, for the competing cause as follows:

$$\eta_X^2 = [\omega_x SS_{\text{within}}] / [(\omega_t + \omega_x) SS_{\text{total}}]. \quad (\text{B4})$$

B3. Feedforward Model

The feedforward architecture consists of two input nodes representing the target cause and the competing cause and an output node representing the outcome. The input nodes are connected to the output node via weighted, unidirectional links. When a cause is present at a trial, its input node is activated to the default level 1; when a cause is missing at a trial, its input node is activated to a negative value α_n , which can be freely estimated between 0 and -1 (Van Hamme & Wasserman, 1994). However, to keep the number of free parameters in all models equal, this parameter was arbitrarily set at an intermediate value of $\alpha_n = -.5$. The positive activation of all input nodes is spread automatically to the output node in proportion to the weights of the links and then linearly summed to represent the output activation.

After each trial, the weights of the links are incrementally adjusted by reducing the error between the output activation (the outcome predicted by the network) and a teaching activation (the actual outcome), which is 1 when the focal outcome is present and -1 when absent (this coding assumes that an absent outcome was interpreted as opposite to the focal outcome, which is plausible for most of our stories). This adjustment is mathematically expressed by the following delta algorithm (McClelland & Rumelhart, 1988, p. 87):

$$\Delta w = \varepsilon (a_t - a_o) a_i, \quad (\text{B5})$$

where ε is the learning rate (freely estimated between 0 and 1) and a_t , a_o , and a_i denote, respectively, the teaching, output, and input activations. We assumed that there were separate learning rates for target factors and competing factors, denoted respectively by ε_t and ε_x .

NOTES

1. To verify that these sample size effects were not due to the within-subject nature of the experiment, we replicated the sequential format using a between-subjects design. To avoid a growing anticipation among participants that the number of trials would always be identical, as well as the experimental demand that even a minimal number of tri-

als (in the small size condition) is informative, we inserted filler stories with two and six trials. Consistent with our predictions, when size was large as opposed to small, the target was more discounted, $F(1, 75) = 3.78$, $p = .056$, and more augmented, $F(1, 45) = 5.24$, $p < .05$. Although these results are less reliable than in the present within-subjects experiment, they indicate that the sample size effect on competition survives under different experimental conditions.

2. Although most interactions with Measure and Size did not reach significance, there was a less interesting Measure \times Size \times Type interaction on the target rating, $F(1, 114) = 16.01$, $p < .001$, which indicated that the target stimulus was less augmented than the target actor given a large size.

3. The simulations of the recurrent model were run using the same specifications as the feedforward model, with the following additional recurrent parameters: $\text{istr} = \text{estr} = \text{decay} = 1$, using the linear activation rule with 1 internal processing cycle (McClelland & Rumelhart, 1988). These parameters make the recurrent model most similar to the feedforward model. In addition, we also ran the recurrent model as specified in earlier social research (Read & Montoya, 1999; Smith & DeCoster, 1999) with parameters $\text{istr} = \text{estr} = \text{decay} = .15$, using the non-linear activation rule with 10 internal processing cycles. The results given both sets of parameters were identical.

REFERENCES

- Allan, L. G. (1993). Human contingency judgments: Rule based or associative? *Psychological Bulletin*, *114*, 435-448.
- Anderson, J. R., & Sheu, C.-F. (1995). Causal inferences as perceptual judgments. *Memory and Cognition*, *23*, 510-524.
- Baker, A. G., Berber, M. W., & Valée-Tourangeau, F. (1989). Judgments of a 2×2 contingency table: Sequential processing and the learning curve. *Quarterly Journal of Experimental Psychology*, *41B*, 65-97.
- Busemeyer, J. R. (1991). Intuitive statistical estimation. In N. Anderson (Ed.), *Contributions to information integration theory: Vol. 1. Cognition*. Hillsdale, NJ: Lawrence Erlbaum.
- Chapman, G. B., & Robbins, S. J. (1990). Cue interaction in human contingency judgment. *Memory and Cognition*, *18*, 537-545.
- 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 H. L. Roiblat & J.-A. Meyer (Eds.), *Comparative approaches to cognitive science* (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.
- Dennett, D. C. (1998). *Brainchildren: Essays on designing minds*. London: Penguin.
- Einhorn, H. J., & Hogarth, R. M. (1986). Judging probable cause. *Psychological Bulletin*, *99*, 3-19.
- Fales, E., & Wasserman, E. A. (1992). Causal knowledge: What can psychology teach philosophers? *Journal of Mind and Behavior*, *13*, 1-28.
- Försterling, F. (1989). Models of covariation and attribution: How do they relate to the analogy of analysis of variance? *Journal of Personality and Social Psychology*, *57*, 615-625.
- Försterling, F. (1992). The Kelley model as an analysis of variance analogy: How far can it be taken? *Journal of Experimental Social Psychology*, *28*, 475-490.
- Gluck, M. A., & Bower, G. H. (1988). From conditioning to category learning: An adaptive network model. *Journal of Experimental Psychology: General*, *117*, 227-247.
- Hansen, R. D., & Hall, C. A. (1985). Discounting and augmenting facilitative and inhibitory forces: the winner takes all. *Journal of Personality and Social Psychology*, *49*, 1482-1493.
- Hogarth, R. M., & Einhorn, H. J. (1992). Order effects in belief updating: The belief-adjustment model. *Cognitive Psychology*, *24*, 1-55.
- Johnson-Laird, P. N. (1983). *Mental models: Towards a cognitive science of language, inference, and consciousness*. Cambridge, UK: Cambridge 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: University of Miami Press.
- Kelley, H. H. (1967). Attribution in social psychology. *Nebraska Symposium on Motivation*, *15*, 192-238.
- Kelley, H. H. (1971). Attribution in social interaction. In E. E. Jones, D. E. Kanouse, H. H. Kelley, R. E. Nisbett, S. Valins, & B. Weiner (Eds.), *Attribution: Perceiving the causes of behavior* (pp. 1-26). Morristown, NJ: General Learning Press.
- Kruglanski, A. W., Schwartz, S. M., Maides, S., & Hamel, I. Z. (1978). Covariation, discounting, and augmentation: Towards a clarification of attributional principles. *Journal of Personality*, *76*, 176-189.
- Lober, K., & Shanks, D. R. (2000). Causal induction based on causal power? Critique of Cheng (1997). *Psychological Review*, *107*, 195-212.
- McClelland, J. M., & Rumelhart, D. E. (1988). *Explorations in parallel distributed processing: A handbook of models, programs and exercises*. Cambridge, MA: Bradford.
- McClure, J. (1998). Discounting of causes of behavior: Are two reasons better than one? *Journal of Personality and Social Psychology*, *74*, 7-20.
- Morris, M. W., & Larrick, R. P. (1995). When one cause casts doubt on another: A normative analysis of Discounting in causal attribution. *Psychological Review*, *102*, 331-355.
- Nosofsky, R. M., Kruschke, J. K., & McKinley, S. C. (1992). Combining exemplar-based category representations and connectionist learning rules. *Journal of Experimental Psychology: Learning, Memory and Cognition*, *18*, 211-233.
- Read, S. J., & Montoya, J. A. (1999). An autoassociative model of causal reasoning and causal learning: Reply to Van Overwalle's critique of Read and Marcus-Newhall (1993). *Journal of Personality and Social Psychology*, *76*, 728-742.
- Rescorla, R. A. (1969). Conditioned inhibition of fear. In W. K. Honig & N. J. Mackintosh (Eds.), *Fundamental issues in associative learning*. Halifax, UK: Dalhousie University Press.
- 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. 64-98). New York: Appleton-Century-Crofts.
- Rosenfield, D., & Stephan, W. G. (1977). When discounting fails: An unexpected finding. *Memory and Cognition*, *5*, 97-102.
- Rudolph, U., & Försterling, F. (1997). The psychological causality implicit in verbs: A review. *Psychological Bulletin*, *121*, 192-218.
- Schustack, M. W., & Sternberg, R. J. (1981). Evaluation of evidence in causal inference. *Journal of Experimental Psychology, General*, *110*, 101-120.
- 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 judgments in experienced and described situations. *Psychological Science*, *2*, 341-350.
- Shanks, D. R. (1995). Is human learning rational? *Quarterly Journal of Experimental Psychology*, *48a*, 257-279.
- Shanks, D. R., Lopez, F. J., Darby, R. J., & Dickinson, A. (1996). 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* (Vol. 34, pp. 265-311). New York: Academic Press.
- Smith, E. R. (1996). What do connectionism and social psychology offer each other? *Journal of Personality and Social Psychology*, *70*, 893-912.

- Smith, E. R., & DeCoster, J. (1999). *Heuristic-systematic and other dual process models in social and cognitive psychology: An integration and connectionist interpretation*. Unpublished manuscript.
- Van Hamme, L. J. (1994). *Associative and statistical accounts of cue competition in causality judgments*. Unpublished doctoral dissertation, University of Iowa, Iowa City.
- Van Hamme, L. J., & Wasserman, E. A. (1994). Cue competition in causality judgments: The role of nonpresentation of compound stimulus elements. *Learning and Motivation, 25*, 127-151.
- Van Overwalle, F. (1996). The relationship between the Rescorla-Wagner associative model and the probabilistic joint model of causality. *Psychologica Belgica, 36*, 171-192.
- Van Overwalle, F. (1997). A test of the Joint model of causal attribution. *European Journal of Social Psychology, 27*, 221-236.
- Van Overwalle, F. (1998). Causal explanation as constraint satisfaction: A critique and a feedforward connectionist alternative. *Journal of Personality and Social Psychology, 74*, 312-328.
- Van Overwalle, F., Labiouse, C., & French, R. (2000). *Connectionist exploration in social cognition*. Unpublished manuscript.
- Van Overwalle, F., & Van Rooy, D. (1998). A connectionist approach to causal attribution. In S. J. Read & L. C. Miller (Eds.), *Connectionist models of social reasoning and social behavior* (pp. 142-171). Hillsdale, NJ: Lawrence Erlbaum.
- Van Overwalle, F., & Van Rooy, D. (2000). *When more observations are better than less: A connectionist account of the acquisition of causal strength*. Manuscript submitted for publication.
- Wasserman, E. A., Kao, S.-F., Van Hamme, L., Katagiri, M., & Young, M. E. (1996). Causation and association. *The Psychology of Learning and Motivation, 34*, 207-264.
- Weiner, B. (1985). An attributional theory of achievement motivation and emotion. *Psychological Review, 92*, 548-573.
- Wells, G. L., & Ronis, D. L. (1982). Discounting and augmentation: Is there something special about the number of causes? *Personality and Social Psychology Bulletin, 8*, 566-572.
- Williams, D. A., & Docking, G. L. (1995). Associative and normative accounts of negative transfer. *Quarterly Journal of Experimental Psychology, 48a*, 976-998.
- 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.

Received July 24, 2000

Revision accepted January 23, 2001