European Journal of Social Psychology
Eur. J. Soc. Psychol. **33**, 515–533 (2003)
DOI: 10.1002/ejsp.160

Acquisition of dispositional attributions: effects of sample size and covariation

FRANK VAN OVERWALLE*
Vrije Universiteit Brussel, Belgium

Abstract

Two experiments examined whether dispositional attributions are sensitive to the sample size of the evidence indicating a given level of covariation between person and behavior. Participants were given high or low levels of covariation (i.e. consensus and distinctiveness), and the acquisition of dispositional attributions was monitored by requesting dispositional trait ratings at fixed intervals. The results showed that dispositional attributions were sensitive to sample size, and increased given more evidence on high person-behavior covariation while they decreased given more evidence on low person-behavior covariation. Additional analyses suggested that in making dispositional inferences (e.g. about the actor), there was a slight preference for agreement information (e.g. low distinctiveness) over difference information (e.g. low consensus). The effects of sample size are inconsistent with current statistical or probabilistic models of covariation, but are in line with connectionist networks using an error-correcting learning algorithm. Copyright © 2003 John Wiley & Sons, Ltd.

When Sally, your new neighbor, talks a lot with your close friend Bart at her house-warming party, you might start to believe that she is rather talkative and sociable. When you notice her talking again with some of your friends on several occasions, your judgment about her sociability will probably become more firm and pronounced. The idea that we make more extreme dispositional attributions when we possess more supportive evidence is captured by the *law of large numbers* or *sample size effect*, which says that our judgments 'should be more confident when they are based on a larger number of instances' (Nisbett, Krantz, Jepson, & Kunda, 1993, p. 339).

However, when you observe that Sally rarely talks with anyone except with your close friend Bart, your attribution will probably shift and you might start wondering whether she considers Bart to be a quite attractive and interesting person. This change in dispositional attribution is captured by the *principle of covariation*, which predicts that attributions are made to a condition that covaries uniquely with the effect, that is, 'that condition which is present when the effect is present and which is absent when the effect is absent' (Kelley, 1967, p. 194). Since Sally seems to talk only with Bart, this implies that her behavior covaries uniquely with Bart, suggesting that his attractive personality is the cause of her attention. The combined effect of covariation information and sample size in making dispositional attributions is the focus of the present article.

---

*Correspondence to: Frank Van Overwalle, Department of Psychology, Vrije Universiteit Brussel, Pleinlaan 2, B-1050, Brussel, Belgium. E-mail: Frank.VanOverwalle@vub.ac.be

Received 9 March 2002
Accepted 7 December 2002

SAMPLE SIZE AND COVARIATION

Evidence from different research areas in social cognition indicates that people have an intuitive appreciation of the law of large numbers and seem to recognize that when making inferential judgments, more evidence is better than less. For instance, when receiving more supportive information, people make more extreme causal attributions (Van Overwalle & Van Rooy, 2001a), hold more extreme impressions about other persons (Anderson, 1967, 1981), make more polarized group decisions (Fiedler, 1996; Ebbesen & Bowers, 1974), endorse more firmly an hypothesis (Fiedler, Walther, & Nickel, 1999), make more extreme predictions (Manis, Dovalina, Avis, & Cardoze, 1980) and agree more with persuasive messages (Eagly & Chaiken, 1993).

Although attributions to dispositions share similarities with causal attributions (e.g. Kelley, 1967) and person impressions (e.g. Anderson, 1981), they do not involve the same type of judgment. In *causal attribution*, people seek the cause of a behavior (e.g. something about Sally), while in *dispositional attribution*, people make a specific trait inference about an actor (e.g. Sally is talkative). In addition, in research on *person impression*, people infer dispositions from traits or behavioral descriptions that refer uniquely and directly to the actor, while in dispositional attribution research, people have to infer the actor's causal role from covariation information involving the actor and other information concerning other persons or the situation, before making a correspondent inference to the actor's dispositions.

Although it seems quite reasonable to assume that attributions to dispositions would also be susceptible to the amount of evidence just like attributions to causes or trait impressions as shown in previous research (Anderson, 1967, 1981; Van Overwalle & Van Rooy, 2001a), to date no such study has been carried out. Therefore, the aim of the present studies is to test the hypothesis that sample size can influence dispositional attributions, irrespective of the level of covariation. To test this, I manipulated sample size by simply repeating the information given while keeping the covariation between the actor and his or her behavioral outcomes constant.

Why do I test this sample size hypothesis, which seems so intuitively plausible and which seems to be already captured by Kelley's (1967) consistency dimension (i.e. the extent to which the outcome is constant over time)? A first reason is that the consistency dimension is rather inflexible. It only predicts that consistent behavior, such as Paul's repeated failure at his driver's license test, should result in strong attributions to the actor's disposition (e.g. Paul is clumsy), while inconsistent behavior should result in weak attributions. Other interesting issues were not addressed in his covariation model. For instance, what about mixed consistency? The most likely prediction from Kelley's covariation model is that this would result in a moderate dispositional attribution, while the present approach is more flexible in that it predicts a wider range of weak to strong dispositional attributions depending on when the inconsistent information is provided in the stream of consistent information, and the behavioral consistency towards other persons. This issue is further elaborated in the General Discussion.

A second reason is that, surprisingly, all prominent attribution models in social psychology that use a statistical formulation—often to further articulate Kelley's original ideas—actually contradict that sample size plays a independent role on dispositional ratings (e.g. Anderson & Sheu, 1995; Busemeyer, 1991; Cheng & Novick, 1990; Försterling, 1989; Hogarth & Einhorn, 1992). As I will argue shortly, only connectionist models make the prediction that variations in sample size lead to variations in dispositional attributions (Read & Montoya, 1999; Van Overwalle, 1998; Van Overwalle & Van Rooy, 1998). This has far-reaching implications on how we view the dispositional attribution process. In contrast to statistical models that describe attribution as a laborious application of statistical rules, connectionist models assume that judgments of causality are often an implicit on-line adjustment process outside awareness, and that only the outcome of that process is open to consciousness (Smith & DeCoster, 1998). This latter view seems to describe much better how causes and dispositions are learned and understood in the hustle of everyday social life.

Statistical Models

Why do statistical models fail to predict changes in dispositional judgments when causal evidence is repeated? The key limitation of all these models is that they are 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 because simply repeating evidence does not change the relevant proportions, these models predict the same constant dispositional judgments (see Van Overwalle & Van Rooy, 2001a, for the same logic given causal attributions). Thus, although statistical models may adequately describe dispositional 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.

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 which makes them sensitive to sample size (the step-by-step belief-adjustment model of Hogarth & Einhorn, 1992; the serial averaging strategy of Busemeyer, 1991). 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 additional causes. As I will discuss later, this restriction is untenable for dispositional attributions as they typically involve comparisons between at least two types of causes involving the person (i.e. consensus) and the stimulus (i.e. distinctiveness). In addition, it has been shown that these models are mathematically identical to a simplified, single-cause version of the delta algorithm applied in connectionist models (Wasserman, Kao, Van Hamme, Katagiri, & Young, 1996; Appendix D). Thus, a straightforward manner to extend these models to allow for multiple causes would actually render them similar to connectionist models. For all these reasons, I will ignore these models.

Connectionist Models

The connectionist framework presents an entirely different perspective on causal judgments that is inspired by basic neurological properties of the human brain (Read & Montoya, 1999; Van Overwalle, 1998; see also Allan, 1993; McClelland & Rumelhart, 1988; Shanks, 1995). Unlike earlier statistical models, adaptive connectionist models with a learning algorithm predict differences in dispositional judgments given changes in sample size.

To explain the properties of adaptive networks, I focus on one of the simplest architectures, the feedforward model (McClelland & Rumelhart, 1988, pp. 83ff.; Van Overwalle, 1998). 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). The activation of the output node reflects the effect predicted by the network.

A key feature of adaptive connectionist networks is that the connection weights are adapted in response to information on new co-occurences between causes and effect, using a learning algorithm. The learning algorithm I focus on here is the *delta algorithm* (McClelland & Rumelhart, 1988). When the occurrence of the effect is underestimated by the network, the weights are adjusted upward; when the occurrence of the effect is overestimated, adjustments are made downward. This makes the network sensitive to sample size, as weights typically start at low levels and gradually increase or decrease depending on the presence or absence respectively of the effect. This process goes on until the predicted output of the network matches the actual effect, that is, until the mismatch or error between one's mental causal representation and reality is reduced, and learning stops. Although this mechanism does not tally frequencies nor computes statistical probabilities, it forces the weights to converge to the probabilistic norm of causality (Cheng & Novick, 1990) after a sufficient number of observations (Chapman & Robbins, 1990; Van Overwalle, 1996). Thus, the delta learning algorithm respects the statistical principles of both sample size and covariation.

USING DIFFERENT TYPES OF COVARIATION INFORMATION

A secondary aim of this research is to explore how different types of covariation information contribute to dispositional attributions. Research by Hilton, Smith, and Kim (1995) and Van Overwalle (1997) seems to suggest that participants use covariation evidence in a significantly different way for making dispositional attributions than for making causal attributions. In general, two types of covariation can be applied, the *method of difference* and the *method of agreement* (J. S. Mill, 1872/1973).

The method of difference focuses on differences between causal conditions and designates the cause to that condition that differs from other conditions. Consider, for example, Sally who is the only person that talks a great deal at parties while others talk considerably less (i.e. low consensus of the actor). Because her behavior differs from that of others, this designates Sally as the cause of her behavior. In contrast, the method of agreement focuses on generalities in someone's behavior so that the cause is designated to that condition that agrees in all instances. For example, if Sally talks a great deal to every guest (i.e. low distinctiveness), this implies generalization across the recipients or stimuli of her behavior, and hence attributions are again made to Sally.

The same logic also applies to dispositional attributions of the stimulus person. Consider, for example, the method of difference. When Sally talks a great deal only to Bart (i.e. high distinctiveness of the stimulus), this implies that Bart differs from the other stimuli so that attributions are made to the stimulus person, Bart. In contrast, the method of agreement is involved when all guests at a party talk a lot to Bart (i.e. high consensus of the actors). This again designates Bart as the cause of their behavior, because Bart is the only factor that covaries with the behavior of all guests (i.e. on which they all agree in their behavior).

In sum, when making covariation-based dispositional judgments, perceivers may focus on differences or generalities in behaviors. To make a dispositional attribution to the actor, information should indicate low consensus (method of difference), low distinctiveness (method of agreement), or both. To make an attribution to the stimulus person, information should reflect high distinctiveness (method of difference), high consensus (method of agreement), or both (see also Kelley, 1967).

However, there is some uncertainty as to how much each type of information implied by the two methods is used in making dispositional attributions. Initial findings by Hilton et al. (1995) suggested that dispositional attributions privilege the use of Mill's method of agreement, whereas causal explanations privilege the use of Mill's method of difference. According to Hilton et al., this is due to the fact that for making dispositional attributions, people have to learn more about generalizations in characteristics or dispositions of the actor. Thus, for dispositional attributions, people seem to prefer low distinctiveness (of stimuli) for inferring actor dispositions and high consensus (of actors) for inferring stimulus dispositions. However, in a subsequent study, Van Overwalle (1997) suggested that dispositional attributions depend on the joint and equal use of difference and agreement information. He argued that the comparison information provided by Hilton et al. (1995) was not complete, so that participants could not use the method of difference fully because they lacked the necessary information to compare and differentiate all relevant cases. When all possible comparison information was provided, Van Overwalle (1997) found that dispositional attributions were guided by an equal amount of the methods of difference and agreement. The present research was designed to provide additional evidence on this issue.

Given that the connectionist delta algorithm converges to the predictions of the covariation model as mentioned earlier, one would expect that it also incorporates Mill's methods of difference and agreement. And, indeed, it does. It incorporates the method of difference, because the more the target and the less other comparison factors are paired together with a behavior (e.g. Sally talks a lot but nobody else), the stronger the connection with the target factor becomes. Conversely, it incorporates the method of agreement because the more the target and other comparison factors are paired with a behavior (e.g. Sally and everyone else talks a lot), the weaker the connection with the target factor becomes (due to the property of competition, see Van Overwalle & Van Rooy, 1998, 2001b).

EXPERIMENT 1

In the first experiment, participants read several stories depicting an actor engaging in a behavior directed towards another stimulus person or object. Covariation of the behavior with the actor and with the stimulus was either high or low. To examine the effect of sample size given confirming evidence, this covariation pattern was repeated across eight blocks of trials. In contrast, to examine the effect of disconfirming information, in another condition, this covariation pattern was reversed halfway the blocks. At regular intervals (i.e. after each block), participants judged a relevant dispositional characteristic of the actor and the stimulus.

As predicted by the sample size effect, when covariation is repeated over all eight blocks, the differences between the covariation levels should gradually grow stronger. At the final judgment, I expect the highest dispositional attributions when there is high covariation with the target factor (actor or stimulus), and the lowest dispositional judgments when there is high covariation with the other factor (i.e. the stimulus when the actor is the target; the actor when the stimulus is the target). For instance, when only Sally talks a lot to all other people (target covariation) I expect the highest talkative disposition for Sally, whereas if everyone talks to Bart and nobody else (other covariation) I expect the lowest talkative disposition for Sally. When covariation with the actor and stimulus are both high or both low, dispositional judgments should end at an intermediate level. For instance, when only Sally talks a lot to Bart and nobody else (both target and other covariation) I expect only a moderate talkative disposition for Sally because, in fact, Bart's attractiveness partly instigated her behavior. Thus, the greater attribution to Bart draws away talkative dispositions for Sally. Finally, when everyone talked a lot to everyone else (no target or other covariation), this seems to be the norm in this group so that talkative dispositions for individual members like Sally should again be moderate.

The intermediate covariation levels allow exploring the relative contribution of the method of difference versus agreement. For instance, if covariation with Sally and Bart is both high, this implies that Sally and Bart are different from comparison actors and stimuli, suggesting the sole application of the method of difference. If participants prefer the method of difference, then they should give relatively higher dispositional ratings in this condition. Conversely, if covariation with Sally and Bart is both low, this implies an identical outcome across all actors and stimuli, suggesting solely the method of agreement. If participants prefer this method as claimed by Hilton et al. (1995), they should give somewhat higher ratings in this condition. Finally, if both methods are applied equally strongly, ratings in these two conditions should be equally high. Other comparisons are also informative, but have less generality because they involve a combination of Mill's methods of difference and agreement, and thus lend themselves less easily to estimate the relative weight of each method.

In addition, when the covariation levels are reversed after the first half of blocks, I expect a gradual reversal of dispositional judgments so that at the final judgments, there are little differences left, because the information in the second half of the event descriptions completely overrides the information in the first half of the descriptions. Van Overwalle and Van Rooy (2001a) reported this pattern of results with causal attributions.

Method

Participants

Participants were 67 male and female freshmen from the Dutch-speaking Vrije Universiteit Brussel, who participated for a partial requirement of an introductory psychology course. They were tested individually.

Materials

Stories All instructions, materials and questions appeared on an IBM-compatible PC screen, and the whole experiment was monitored by MEL software. The participants read eight short event descriptions that were shown in eight consecutive blocks of four trials. The four trials in an event/block each depicted the outcome of a focal actor, a focal stimulus (person or object), a comparison actor and a comparison stimulus. For instance, consider an event involving Jasmine as focal actor and Corinne as focal stimulus. The first block contained all four combinations of the focal and comparison actors with the focal and comparison stimuli. Hence, these four trials read as follows (the focal actor and stimulus are in italics):

- *Jasmine* cheated on her friend, *Corinne*.
- *Jasmine* cheated on her friend, Karen.
- Sigrid cheated on her friend, *Corinne*.
- Sigrid cheated on her friend, Karen.

Each trial was displayed consecutively on a separate screen in a random order for each participant. In the remaining blocks, the focal actor and stimulus remained the same, whereas the comparison actors and stimuli changed at each block. Each block was described as a consecutive occasion in time. After each block, the participants rated the dispositional characteristics of the focal actor and stimulus.

Covariation High covariation of the actor was created by having only the focal actor be followed by the outcome (e.g. Jasmine cheated on her friend, Corinne) while the comparison actor attained the opposite outcome (e.g. Sigrid did NOT cheat on her friend, Corinne). In contrast, low covariation was created by having the comparison actors be followed by the same outcome as the focal actor. As can be seen from the example, the opposite outcome was created by negating the focal outcome, and the negative wording of an outcome was always given in capitals to ensure correct encoding of the outcome information. The same logic was used for high and low covariation of the stimulus.

Table 1. Design of Experiment 1 illustrated for the event 'Jasmine cheated on her friend, Corinne' with Jasmine as target actor and Corinne as target stimulus

	Repeated covariation			
Trial type	Target (actor) (LL)	None (HL)	Both (LH)	Other (stimulus) (HH)
Jasmine + *Corinne*	*Cheated*	*Cheated*	*Cheated*	*Cheated*
Jasmine + Karen[a]	*Cheated*	*Cheated*	Did not cheat	Did not cheat
Sigrid[a] + *Corinne*	Did not cheat	*Cheated*	Did not cheat	*Cheated*
Sigrid[a] + Karen[a]	Did not cheat	*Cheated*	Did not cheat	Did not cheat

	Reversed covariation			
Trial type	Other (stimulus) (HH)	Both (LH)	None (HL)	Target (actor) (LL)
Jasmine + *Corinne*	*Cheated*	*Cheated*	*Cheated*	*Cheated*
Jasmine + Veronique[a]	Did not cheat	Did not cheat	*Cheated*	*Cheated*
Gwen[a] + *Corinne*	*Cheated*	Did not cheat	*Cheated*	Did not cheat
Gwen[a] + Veronique[a]	Did not cheat	Did not cheat	*Cheated*	Did not cheat

Note. The focal actor, focal stimulus, and focal outcome of the story are given in italics. In this example, the target is the focal actor 'Jasmine', and the other causal factor is the focal stimulus 'Corinne', and the focal outcome is 'cheated'. Between parentheses is indicated whether, for this example, consensus and distinctiveness were H = High or L = low respectively. In the repeated condition, each covariation condition in the top panel was repeated over eight consecutive blocks; in the reversed condition, each covariation condition in the top panel was reversed after the fourth block as shown in the same column of the bottom panel.

[a] There were eight different comparison targets and stimuli (one for each block) in Experiment 1, and there were six in Experiment 2.

Taken together, I created four possible combinations of covariation as illustrated in Table 1. *Target*-covariation denotes that covariation with the target was high (and low for the other factor). In the example, target covariation is indicated by the fact that Jasmine cheated on her friends and that no one else cheated. *Other*-covariation denotes that covariation with the other factor was high (and low for the target); *both*-covariation indicates that covariation with both the target and the other factor was high, while *none*-covariation denotes that covariation was low for both. Table 1 illustrates how these covariation levels were created by changing the focal outcome and also demonstrates how they relate to high or low consensus and distinctiveness.

Repeated versus Reversed In the repeated condition, the four covariation patterns remained identical throughout all blocks of an event. In the reversed condition, these covariation patterns were shown in the first half of block, and then reversed in the second half of blocks. That is, the target- and other-covariation conditions were switched, as well as the both- and none-covariation conditions. As illustrated in Table 1, participants went through one of the columns (or conditions) of the table. In the repeated conditions, they went only through one of the top columns for all eight blocks. In the reversed condition, they went through the same column of the top and bottom panel for the first four and last four blocks respectively.

To control for potential effects of implicit causality, the number of action versus state verbs was counterbalanced across all events (Rudolph & Försterling, 1997). The computer randomly determined for each subject the selection of the events in each covariation condition, the order in which the events appeared, as well as the order of the trials within each block.

Procedure

The instructions appeared on the screen and the use of the rating scale was practiced. For each event, the participants then read eight blocks of four trials. After each block, they rated the disposition of the actor and the stimulus on two 11-point scales. Following Van Overwalle (1997), the particular questions and anchors were specifically fitted to each story. For instance, in the prior example, the participants had to judge Jasmine and Corinne by rating the following dispositions: 'To what extent is Jasmine a cheat' (0 = *not at all a cheat* to 10 = *very much a cheat*) and 'To what extent is Corinne naive' (0 = *not at all naive* to 10 = *very much naive*). The participants indicated their answer by moving through the scale points in steps of 1, using the left and right arrow keys. The order in which the dispositions were rated was randomized for each participant and each event.

Results

Because I made the same predictions for actor and stimulus dispositions, these two ratings were collapsed together (with the provision that they held the same role of *target* and *other* factor in each covariation condition). This was justified, as a preliminary analyses of variance (ANOVA) with Measure (actor or stimulus), Target Covariation (high or low), Other Covariation (high or low) and Block (1 to 8) as within-subject factors revealed no significant interactions with Measure in both the repeated and the reversed conditions.

Given that I made different predictions for repeated versus reversed covariation patterns, I discuss the results separately.

Repeated Condition

The top panel of Figure 1 depicts the mean disposition ratings of the repeated condition. The ratings were analyzed using an ANOVA with Target Covariation (high or low), Other Covariation (high or low), and Block (first vs. last) as within-subjects factors. The main effect of Target Covariation was significant, $F(1, 66) = 4.72$, $p = 0.033$, as well as the main effect of Other Covariation, $F(1, 66) = 192.74$, $p < 0.0001$. This indicated that participants were sensitive to the different levels of covariation presented to them. The main effect of Block did not reach significance, $F < 1$. Of most interest was the expected interaction between Covariation and Block, which indicates whether the disposition ratings are adjusted after each block as a consequence of the increasing sample size. Consistent with my prediction, the interaction was highly significant for both Target Covariation, $F(1, 66) = 20.14$, $p < 0.0001$, and Other Covariation, $F(1, 66) = 82.03$, $p < 0.0001$.

I explored this interaction further by conducting trend analyses in each covariation conditions. As can be seen in Figure 1, in line with my prediction, there was a significant linear increase over blocks in the target-condition, $F_{\text{lin}}(1, 66) = 61.30$, $p < 0.0001$, and a linear decrease in the other-condition, $F_{\text{lin}}(1, 66) = 40.86$, $p < 0.0001$. There was also a linear increase when none of the two factors covaried with the outcome, $F_{\text{lin}}(1, 66) = 21.52$, $p < 0.0001$, while there was a linear decrease when they both covaried, $F_{\text{lin}}(1, 66) = 9.82$, $p < 0.003$.[1]

---

[1] The delta learning algorithm would predict a linear increase that is negatively accelerating, that is, that gradually flattens towards asymptote when multiple trials with the same (consistent) information have been processed. However, given that, at most, eight trials repeated the information, this negatively accelerating effect might still be quite weak in the present data set. This was indeed the case, as a significant quadratic effect was found only in the target-condition, $F(1, 66) = 20.52$, $p < 0.0001$, but not in the other covariation conditions. In Experiment 2, none of the covariation conditions showed a significant quadratic effect.

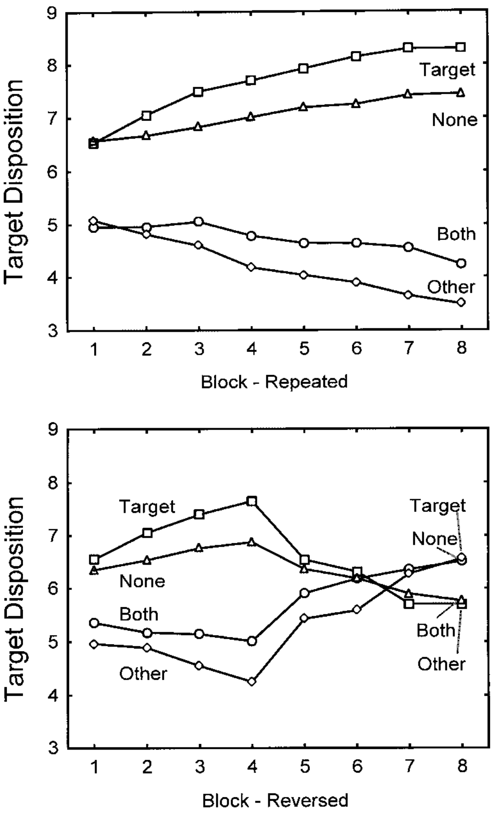

Figure 1. Dispositional ratings of the target in function of covariation pattern and block, for the repeated condition (top panel) and reversed condition (bottom panel): Experiment 1

Because none-covariation implies the method of agreement over all factors (e.g. all actors and stimuli elicited the same outcome) while both-covariation implies the method of difference over all factors (e.g. the focal actor and stimulus together elicited a different and unique outcome), the higher ratings for the none-covariation condition suggests that, in general, the method of agreement produced stronger dispositional attributions than the method of difference. However, the method of difference still had some effect, as can be seen by comparing the target-condition with the none-condition, and by comparing the both-condition with the other-condition, because for each comparison the target covaried in the first condition but not in the second, while covariation by the other factor was constant. For instance, the method of agreement would predict that target-covariation (HH) and none-covariation (HL) should both lead to dispositional attributions to the target actor, yet the ratings were somewhat stronger for the target-covariation (HH) where there is low consensus, suggesting some contribution for the method of difference. Planned contrasts confirmed that at the last block, all adjacent covariation patterns differed reliably from each other, $F\text{s}(1, 66) = 6.82\text{–}103.00$, $p\text{s} < 0.011$, indicating that both methods contributed to the dispositional ratings.

Reversed Condition

The bottom panel of Figure 1 depicts the mean disposition ratings of the reversed condition. The ratings show a similar pattern as the repeated condition for the first four blocks, and a reversal of this pattern in the last four blocks, as expected. Like before, the none-condition led to higher ratings than the both-condition, which suggests that the method of agreement was stronger than the method of difference. The ratings between the first and fourth blocks were subjected to the same ANOVAs as described above in the repeated condition (the ratings of the fourth to eight blocks were not analyzed as they confound change in sample size with change in covariation). These analyses gave very similar results and are therefore not reported.

More importantly, because of the reversal of the covariation patterns, I predicted that there should be a null-effect at the last block, that is, no significant differences between the covariation conditions. However, contrary to this prediction, it can be seen that there is a crossover (or recency effect), which suggests that participants gave more weight to the more recent (disconfirming) covariation information. Planned comparisons confirmed that the ratings of the target- and none-conditions were reliably higher than the ratings of the both- and other-conditions in the seventh block, $F(1, 66) = 7.19$, $p < 0.01$, and in the last block, $F(1, 66) = 15.92$, $p < 0.001$.

Discussion

The findings in the repeated condition demonstrate that participants tended to make progressively more extreme dispositions as more covariation evidence was provided that confirmed their initial judgments. The findings in the reversed condition further indicate that when initial evidence was disconfirmed, the participants converged to the novel level of covariation with increasing observations. These findings indicate that sample size and covariation both contribute to the dispositional attributions we make. Contrary to my expectations, when initial evidence was disconfirmed in the reversed condition (with an equal amount of disconfirming and confirming information), the disposition ratings were not restored to a neutral scale midpoint, but were somewhat more influenced by the recent disconfirmatory information.

EXPERIMENT 2

The results of the first experiment indicated that although both Mill's method contributed to dispositional ratings, the method of agreement was more strongly applied than the method of difference. However, this result may be due to a methodological limitation in that information was not provided on all possible combinations of comparison actors and stimuli, but only on the comparison actor and stimulus of the current block. Thus, in total, only eight out of 64 comparison combinations of comparison actors and comparison stimuli were filled with relevant information. As suggested by Van Overwalle (1997), when comparison information is not complete, participants cannot use the method of difference fully because they lack the necessary information to compare and differentiate all relevant cases. To rule out this potential methodological explanation and to determine whether the method of agreement is indeed preferred in making dispositional attributions, I carried out a second experiment in which complete information on all comparison combinations was provided.

Method

Participants

Participants were 71 male and female freshmen from the Dutch-speaking Vrije Universiteit Brussel, who participated for a partial requirement of an introductory psychology course.

Materials and Procedure

All instructions, materials, and questions were similar as in the previous experiment, with the following modification. To provide complete information on all combinations of comparison cases, for each block, in addition to information on the current comparison actor and comparison stimulus, information was also provided on all new combinations of the current comparison cases with comparison actors and stimuli from all previous blocks. As a result, the number of trails increased by two in each block (e.g. their number was 4, 6, 8 and so on). To keep the overall number of trials reasonable, however, only six blocks of trials were provided in this experiment.

Results

Again, the actor and stimulus ratings were collapsed and analyzed together, as a preliminary ANOVA with Measure (actor or stimulus), Target Covariation (high or low), Other Covariation (high or low) and Block (1 to 6) as within-subjects factors revealed no significant interactions with Measure.

Repeated Condition

The top panel of Figure 2 depicts the mean disposition ratings of the repeated condition. The results are very much like those of the prior experiment, save for the difference between the none- and

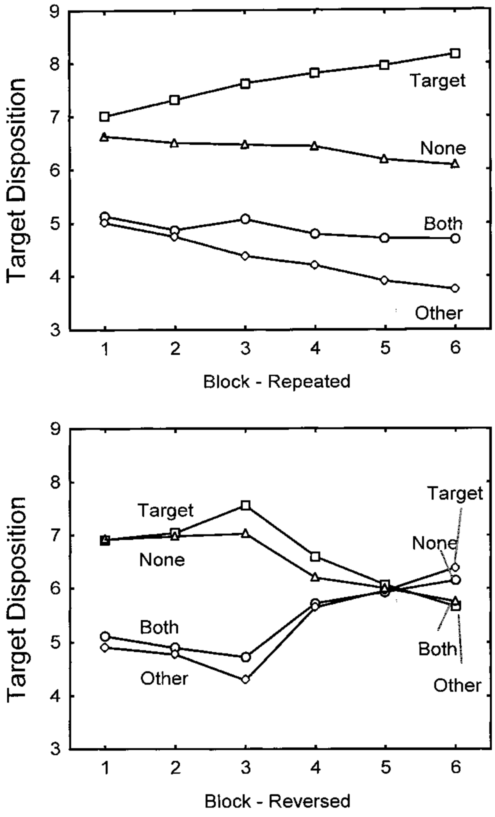

Figure 2. Dispositional ratings of the target in function of covariation pattern and block, for the repeated condition (top panel) and reversed condition (bottom panel): Experiment 2.

both-conditions that was much more attenuated, as I had hoped for, although it did not disappear entirely. However, if anything, this difference tended to grow a little smaller over blocks rather than increase.

To verify these observations, the ratings were analyzed using an ANOVA with Target Covariation (high or low), Other Covariation (high or low), and Block (first versus last) as within-subjects factors. The main effects of Target Covariation and Other Covariation were significant, $F\text{s}(1, 70) = 51.39-207.69$, $p\text{s} < 0.0001$, suggesting that participants differentiated between the levels of covariation. The main effect of Block also reached significance, although this effect was rather weak, $F(1, 70) = 6.31$, $p < 0.02$. Of major interest was the interaction between Covariation and Block. Consistent with my sample size prediction, the interaction was highly significant for both Target Covariation, $F(1, 70) = 33.99$, $p < 0.0001$, and Other Covariation, $F(1, 70) = 16.20$, $p < 0.0001$.

I explored the interaction further by conducting trend analyses in each covariation condition. The results confirmed that, as predicted, there was a significant linear increase in the target-condition, $F_{\text{lin}}(1, 70) = 34.76$, $p < 0.0001$, and a linear decrease in the other-condition, $F_{\text{lin}}(1, 70) = 46.59$, $p < 0.0001$. Contrary to the first experiment, the ratings in the remaining conditions now showed only very mild linear trends, which was marginal for the none-condition, $F_{\text{lin}}(1, 70) = 3.73$, $p = 0.057$ and just reached significance for the both-condition, $F_{\text{lin}}(1, 70) = 4.30$, $p = 0.042$. As in the first experiment, however, the difference between all four covariation patterns was significant at the final judgment, as planned contrasts revealed that all adjacent covariation patterns differed reliably from each other at the last block, $F\text{s}(1, 70) = 17.27-77.68$, $p\text{s} < 0.0001$. The higher dispositional ratings of the none-condition (implying the method of agreement over all factors) compared to the both-condition (implying the method of difference over all factors) suggests that the method of agreement—although not growing substantially over sample size—was still more strongly applied generally than the method of difference. Nevertheless, as I had hoped, the method of difference had a more substantial impact than in the first experiment, as can be seen from the greater differences in the ratings between target-covariation and none-covariation, and between both-covariation and other-covariation.

Reversed Condition

The bottom panel of Figure 1 depicts the mean ratings of the reversed condition. As in the first experiment, ANOVAs on the ratings between the first and third block gave very similar results as in the repeated condition, and are therefore not reported. The most interesting result is that the recency effect became unreliable in this experiment, as planned comparisons revealed that none of the covariation conditions differed significantly from each other at the last block, $F\text{s}(1, 70) < 2.38$, $p\text{s} > 0.12$. This is surprising, because in this experiment the number of trials in the second half of blocks was actually larger (due to the repetition of the earlier comparison factors) than in the first half of blocks.

Discussion

This experiment replicated the main results of the previous experiment, and demonstrated again that dispositional ratings were adjusted given more covariation evidence, although the level of covariation was kept constant. Moreover, a number of unexpected results in Experiment 1 were attenuated or disappeared by providing all potential covariation information on all current and previous comparison cases in this experiment. First, the greater weight of the method of agreement was attenuated and did not increase further over trials—although it remained reliable. Second, the recency effect disappeared and became unreliable even with the greater amount of disconfirming information in this experiment.

MODEL SIMULATIONS

To test my prediction that sample size effects are best described by adaptive connectionist models that incorporate the delta learning algorithm, I ran simulations of the manipulations of covariation and size with the feedforward model described earlier (Van Overwalle, 1998; Van Overwalle & Van Rooy, 1998, 2001a). As a way of comparison, I also ran simulations of two prominent statistical models of causality (Cheng & Novick, 1990; Försterling, 1989). However, because these models cannot predict sample size changes as demonstrated earlier, I extended these models with additional weighting parameters that reflect the number of observations (see also Lober & Shanks, 2000; Van Overwalle & Van Rooy, 2001a). To test the fit of the models, I compared the simulated values with the dispositional ratings obtained from the repeated condition in the two studies, as this condition allows examining the effect of sample size independently from covariation changes.

Method

To capture the idea of the method of difference and agreement by comparison actors or stimuli, I borrowed Weiner's (1985) notion of causal dimensions and referred to comparison actors and stimuli as the *external* or *general* context respectively (see also Van Overwalle, 1996, 1997). Thus, the architecture of all models simulated consisted of causal variables reflecting an actor and an external context (to capture consensus), a stimulus and a general context (to capture distinctiveness) and all possible combinations between these two sets. The contextual factors were assumed to be always present. As consensus and distinctiveness reflect two independent covariation dimensions, trial input in the feedforward model was presented in three separate time slices: one for each dimension and one for their interaction (for more mathematical details, see Van Overwalle, 1996). As such, each dimension and their interaction can accrue connection weights independently. In the statistical models, this independence is built-in.

For testing dispositional attributions (as opposed to causal attributions), in all models, a disposition was coded as the outcome predicted given the combined presence of a single target factor and the context of the other factor. For instance, an actor disposition was defined as the outcome of a single *actor* in a *general* stimulus context (i.e. irrespective of a particular stimulus) or a LL covariation pattern. Similarly, a stimulus disposition was defined as the outcome of a single *stimulus* given the presence of an *external* context (i.e. irrespective of a particular actor) or a HH covariation pattern. This coding scheme was deemed most appropriate for dispositions, as the first aspect of the covariation pattern reflects the method of difference (e.g. low consensus of the actor), while the second aspect reflects the method of agreement (e.g. low distinctiveness of the actor).

For running the simulations, in the probabilistic (Cheng & Novick, 1990) and statistical models (Försterling, 1989), the causal variables were entered in the appropriate equations together with the outcome, whereas in the feedforward connectionist model, the causal variables were represented by input nodes and the outcome was represented by an output node (see Van Overwalle, 1996). The fit of the simulated models with the observed data in the experiments was measured using a correlation coefficient, after searching for the highest correlation given all admissible values of two free parameters. For the statistical and probabilistic models, I estimated two different weights parameters (see also Van Overwalle & Van Rooy, 2001a) reflecting the frequencies of cases that involve either the presence or absence of the target cause (ε_t and ε_x respectively, see the Appendix for more details). For the feedforward model, there were two different learning rate parameters for target and context causes (ε_t and ε_x respectively); the initial weights were set to 0.20. It should be noted that small changes in the parameter values did not alter the results meaningfully for all the models tested.

Table 2. Fits of the models to the data in the repeated condition

	Experiment	
Model	1	2
Weighted ANOVA	0.25 ($\omega_t = 0.20$, $\omega_x = 0.70$)	0.24 ($\omega_t = 0.30$, $\omega_x = 0.50$)
Weighted probabilistic	0.25 ($\omega_t = 0.60$, $\omega_x = 0.02$)	0.37 ($\omega_t = 0.81$, $\omega_x = 0.01$)
Feedforward connectionist	0.96 ($\varepsilon_t = 0.10$, $\varepsilon_x = 0.03$)	0.91 ($\varepsilon_t = 0.20$, $\varepsilon_x = 0.07$)

Note. Cell entries are correlations over all random runs within each covariation condition, and then averaged; ω_t = best-fitting weight for target; ω_x = for context; ε_t = best-fitting learning rate for target; ε_x = for context.

All the models were run using the same order of blocks as in the experiments. The weights in the feedforward model were updated after each trial. Because trial order in the experiments was randomized within each block, I ran 20 simulations with a random trial order in each block, and computed an average of the simulated values across all runs (this was not necessary for the statistical models as they are independent of trial order).

Results and Discussion

After the simulations were run, I computed the correlation between the simulated values and observed data points in each of the four covariation patterns in the repeated condition to explore how well the simulations reproduced the sample size effect, and averaged them. These averaged correlations are shown in Table 2.

As can be seen, the fit of the extended statistical models—the weighted ANOVA and weighted probabilistic model—was insufficient, just as they were for causal attributions in the earlier study by Van Overwalle and Van Rooy (2001a). The fit of the feedforward connectionist model was much more adequate ($r > 0.90$), and the correlation exceeded 0.80 in all individual covariation patterns. These results confirm that the acquisition pattern of dispositional judgments can be more adequately accounted for by an adaptive connectionist approach using the error-reducing delta algorithm.

GENERAL DISCUSSION

The present experiments show that our attributions of the traits and dispositions of people are shaped by the covariation with their behaviors or outcomes, and by the amount of evidence we have to support these attributions. We give increasing dispositional estimates when there is more evidence on the high covariation between a person and his or her behavior, and decreasing dispositional estimates when there is more evidence on the low covariation between person and behavior. The present results parallel those of previous research on causal attributions, which were also found to be influenced by sample size (Van Overwalle & Van Rooy, 2001a).

I also documented that the error-reducing delta algorithm incorporated in many adaptive connectionist models (Van Overwalle, 1998; Read & Montoya, 1999) nicely captures this on-line adjustment process given increasing evidence. In contrast, many earlier statistical models of attribution that attracted a lot of attention and excitement in social psychological research, were unable to account for this simple and obvious effect (e.g. Anderson & Sheu, 1995; Busemeyer, 1991; Cheng & Holyoak, 1995; Försterling, 1989; Fales & Wasserman, 1992; Hogarth & Einhorn, 1992).

More germane to dispositional attributions, is the finding that in detecting covariation, people apply both Mill's methods of difference and agreement consistent with Van Overwalle's prediction (1997). However, as noted by Hilton et al. (1995), there is some small but consistent preference for the method of agreement although it does not grow over increasing evidence (when given sufficient information on all comparison cases). Hilton et al. (1995) argued that in order to learn more about some general dispositions of a person or entity, that is, to make dispositional inferences, the method of agreement is preferentially used. This means, for example, that for inferring dispositions about actors, we are somewhat more influenced by agreement information indicating that the actor behaved in similar ways in diverse situations (low distinctiveness) than by difference information indicating that the actor was among few who behaved this way (low consensus). Similarly, for making dispositional inferences about stimuli, we are somewhat more influenced by the agreement information that everyone acted the same way in response to this stimulus (high consensus) than by difference information suggesting that this was the only stimulus the actor responded to in this manner (high distinctiveness).

Limitations

A limitation of this study is that I did not explicitly include consistency information, or the extent to which the outcome is constant over time (Kelley, 1967). However, repeating information about a target's outcome may have increased the *perceived* (rather than *communicated*) amount of consistency, and may so explain the present sample size effects. However, as noted in the introduction, Kelley's consistency dimension is rather inflexible in that it does not give precise or correct predictions for more complex situations such as mixed consistency. For instance, the covariation model would predict that the HH covariation pattern (consistency of the actor on 1 trial only, see Table 1) should result in only weakly positive disposition ratings (e.g. Jasmine cheated only once). In contrast, the findings show instead strongly negative dispositions because the HH pattern also involves high distinctiveness (e.g. everyone cheated once, but always on Corinne, see Table 1), thus focusing attributions to the stimulus and away from the actor. As another example, the covariation model would predict that the combination of the LL pattern followed by the reversed HH pattern (high consistency on all trials followed by consistency on one trial only) should result in moderate dispositional ratings (e.g. Jasmine cheated always in the beginning, but later on only once), but actually the ratings were lowest compared to all other reversed conditions, for the same reason as in the previous example. The present findings are consistent with the domain of causal attributions, where it was also observed that consistency could not explain all the effects of sample size (Van Overwalle & Van Rooy, 2001b).

Another limitation is that the dispositions were rated repeatedly after each block of trials. How can we be sure that this repeated expression of dispositional judgments is not responsible for the present sample size results? In attitude research, for instance, it has been shown that the repeated expression of attitudes increases the attitude extremity (e.g. Downing, Judd, & Brauer, 1992). In person impression research, it has been documented that recency effects tend to disappear when trait ratings are requested at the end rather than intermittently during a series of trait adjectives describing the person (e.g. Dreben, Fiske, & Hastie, 1979). In order to rule out this possibility, one needs a condition where dispositional attributions are requested at the end of the behavioral information, and verify whether size manipulations still influence the dispositional ratings (holding constant the level of covariation). In an as yet unpublished experiment by Van Overwalle (2003), this manipulation was part of a larger study on discounting in attributional dispositions. The results showed that dispositions where stronger after behavioral descriptions (very similar to the present ones) that were repeated five times rather than one time. This is entirely consistent with the present results. However, it is possible that the impact of repeated information may be somewhat stronger given concurrent as opposed to final dispositional ratings, if only because repeated ratings are an additional source that repeats and reactivates the same information.

Implications

An interesting finding is that the present pattern of results is very much like the person impression paradigm, where trait inferences are made after receiving a series of trait adjectives about a person (Anderson, 1981). When participants receive adjectives implying a high level of a trait, their trait ratings gradually increase over trials, whereas they decrease when the adjectives imply a low level of the trait. In addition, just like in the reversed condition, when the high information is followed by low information, the pattern of results show a similar gradual increase and decrease of trait impression (Stewart, 1965). Although causal responsibility is taken for granted by providing an explicit trait adjective, in contrast to my dispositional attribution design where the target's causal status has to be deduced from the covariation information provided, these strong similarities point to very similar, if not identical underlying cognitive processes.

One of the most important implications of connectionist models is that they provide a radically different perspective on the process by which dispositional attributions are inferred. Whereas earlier statistical models assumed that causal perceivers tally frequencies of all possible actor-outcome occurrences and then enter these frequencies into algebraic equations to arrive at a dispositional judgment, the connectionist approach sees causal inference as an effortless, almost unconscious process of which only the outcome enters consciousness. That is, our brain—as developed during evolution to deal increasingly better with understanding and predicting its environment—takes care of most of the inference process and communicates only the outcome of this process to awareness. Such efficient and automatic process is a natural property of connectionist models because the computations required in the learning algorithm operate in parallel at a lower autonomous level that involves only the input and output nodes activated (i.e. present) at a given trial. Algebraic models lend themselves less easily to such efficient and autonomous computation because they require memorizing and computing all relevant information of all trials. The picture of causal reasoning as an predominantly autonomic and parallel process seems much closer to our intuitions on causal and dispositional attribution, as most people seem to spent little effort at all in making adequate judgments. Evidence that covariation is automatically integrated while making spontaneous trait inferences from behavioral descriptions was provided in a study by Van Overwalle, Drenth, and Marsman (1999). They found that when covariation information indicated the actor as the cause, then trait inferences about the actor where spontaneously made, but when covariation information indicated the stimulus as the cause, there were no such spontaneous inferences.

In a recent theoretical paper, Van Overwalle and Labiouse (2003) further developed the implications of a connectionist approach in explaining other related findings on trait inferences including primacy and recency effects in impression formation (Dreben et al. 1979), the asymmetric diagnosticity of ability and morality related behaviors in predicting implied traits (Skowronski & Carlston, 1987), the increased recall for trait-inconsistent information (Hastie, 1980), assimilation and contrast in priming of traits and trait-implying exemplars (Stapel, Koomen, & van der Pligt, 1997), and parallel discounting of trait inferences by situational information (Trope & Gaunt, 2000). Moreover, they demonstrated mathematically that the connectionist approach was consistent with Anderson's (1981) prominent algebraic models of impression formation. In conclusion, the connectionist approach presented here accounts impressively well for dispositional attribution and person impression data, and offers an alternative interpretation to earlier theoretical distinctions and dilemmas in these areas.

APPENDIX

This appendix describes how the statistical models were extended with weighting parameters to allow these models to account for sample size.

Weighted Probabilistic Model

To make the probabilistic model (Cheng & Novick, 1990) sensitive to the number of observations, I weighted the conditional probabilities P 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 target factor T is present and if ω_x denotes a proportion of the frequencies when factor T is absent, then the strength of a target factor T may be formalized as:

$$\Delta P_T = \omega_t P(O \mid T) - \omega_x P(O \mid \sim T) \tag{A1}$$

and, likewise, the strength of a context factor X:

$$\Delta P_X = \omega_x P(O \mid \sim T) \tag{A2}$$

Weighted ANOVA Model

Försterling's (1989) ANOVA formulation defines causal strength as an analog to the effect size η^2 of a standard ANOVA, which is given by $\eta^2_T = SS_{between}/SS_{total}$ for a target factor and $\eta^2_X = SS_{within}/SS_{total}$ for a context factor. To make Försterling's (1989) model sensitive to the number of observations, I weighted η^2 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 factor T this becomes:

$$\eta^2_T = [\omega_t \, SS_{between}]/[(\omega_t + \omega_x)SS_{total}] \tag{A3}$$

and, similarly, for the context factor X:

$$\eta^2_X = [\omega_x SS_{within}]/[(\omega_t + \omega_x)SS_{total}] \tag{A4}$$

ACKNOWLEDGEMENTS

I wish to thank Dirk Van Rooy, Karen Jordens, Bert Timmermans and Tim Vanhoomissen for their comments on an earlier version of the manuscript, and Ingrid Ponjaert for allowing us to use her subject pool.

REFERENCES

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

Anderson, N. H. (1967). Averaging model analysis of set-size effects in impression formation. *Journal of Experimental Psychology, 75*, 158–165.

Anderson, N. H. (1981). *Foundations of information integration theory*. New York: Academic Press.

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

Busemeyer, J. R. (1991). Intuitive statistical estimation. In N. Anderson (Ed.), *Contributions to information integration theory, Vol. 1: Cognition*. Hillsdale, NJ: Erlbaum.

Chapman, G. B., & Robbins, S. J. (1990). Cue interaction in human contingency judgment. *Memory and Cognition, 18*, 537–545.

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

Downing, J. W., Judd, C. M., & Brauer, M. (1992). Effects of repeated expression on attitude extremity. *Journal of Personality and Social Psychology, 63*, 17–29.

Dreben, E. K., Fiske, S. T., & Hastie, R. (1979). The independence of evaluative and item information: Impression and recall order effects in behavior-based impression formation. *Journal of Personality and Social Psychology, 37*, 1758–1768.

Eagly, A. H., & Chaiken, S. (1993). *The psychology of attitudes*. San Diego, CA: Harcourt Brace.

Ebbesen, E. B., & Bowers, R. J. (1974). Proportion of risky to conservative arguments in a group discussion and choice shifts. *Journal of Personality and Social Psychology, 29*, 316–327.

Fales, E., & Wasserman, E. A. (1992). Causal knowledge: What can Psychology teach philosophers? *Journal of Mind and Behavior, 13*, 1–28.

Fiedler, K. (1996). Explaining and simulating judgment biases as an aggregation phenomenon in probabilistic, multiple-cue environment. *Journal of Personality and Social Psychology, 103*, 193–214.

Fiedler, K., Walther, E., & Nickel, S. (1999). The auto-verification of social hypotheses: Stereotyping and the power of sample size. *Journal of Personality and Social Psychology, 77*, 5–18.

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.

Hastie, R. (1980). Memory for behavioral information that confirms or contradicts a personality impression. In R. Hastie, T. M. Ostrom, E. B. Ebbesen, R. S. Wyer, D. L. Hamilton, & D. E. Carlston (Eds.), *Person memory: The cognitive basis of social perception* (pp. 155–177). Hillsdale, NJ: Erlbaum.

Hilton, D. J., Smith, R. H., & Kim, S. H. (1995). The process of causal explanation and dispositional attribution. *Journal of Personality and Social Psychology, 68*, 377–387.

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

Kelley, H. H. (1967). Attribution in social psychology. *Nebraska Symposium on Motivation, 15*, 192–238.

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

Manis, M., Dovalina, I., Avis, N. E., & Cardoze, S. (1980). Base rates can affect individual predictions. *Journal of Personality and Social Psychology, 38*, 231–248.

McClelland, J. M., & Rumelhart, D. E. (1988). *Explorations in parallel distributed processing: A handbook of models, programs and exercises*. Cambridge, MA: Bradford.

Mill, J. S. (1872/1973). *System of logic* (8th ed.). In J. M. Robson (Ed.), *Collected works of John Stuart Mill* (Vols. 7 and 8). Toronto: University of Toronto Press.

Nisbett, R. E., Krantz, D. H., Jepson, C., & Kind, Z. (1983). The use of statistical heuristics in everyday inductive reasoning. *Psychological Review, 90*, 339–363.

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.

Rudolph, U., & Försterling, F. (1997). The psychological causality implicit in verbs: A review. *Psychological Bulletin, 121*, 192–218.

Shanks, D. R. (1995). Is human learning rational? *Quarterly Journal of Experimental Psychology, 48a*, 257–279.

Skowronski, J. J., & Carlston, D. E. (1987). Social judgment and social memory: The role of cue diagnosticity in negativity, positivity and extremity biases. *Journal of Personality and Social Psychology, 52*, 689–699.

Smith, E. R., & DeCoster, J. (1998). Knowledge acquisition, accessibility, and use in person perception and stereotyping: Simulation with a recurrent connectionist network. *Journal of Personality and Social Psychology, 74*, 21–35.

Stapel, D. A., Koomen, W., & J. van der Pligt (1997). Categories of category accessibility: The impact of trait concept versus exemplar priming on person judgments. *Journal of Experimental Social Psychology, 33*, 47–76.

Stewart, R. H. (1965). Effect of continuous responding on the order effect in personality impression formation. *Journal of Personality and Social Psychology, 1*, 161–165.

Trope, Y., & Gaunt, R. (2000). Processing alternative explanations of behavior: Correction or integration? *Journal of Personality and Social Psychology, 79*, 344–354.

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). Dispositional attributions require the joint application of the methods of difference and agreement. *Personality and Social Psychology Bulletin, 23*, 974–980.

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. (2003). A recurrent connectionist model of person impression formation. *Personality and Social Psychology Review*, in press.

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.* New York: Erlbaum.

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

Van Overwalle, F., & Van Rooy, D. (2001b). How one cause discounts or augments another: A connectionist account of causal competition. *Personality and Social Psychology Bulletin, 12*, 1613–1626.

Van Overwalle, F., Drenth, T., & Marsman, G. (1999). Spontaneous trait inferences: Are they linked to the actor or to the action? *Personality and Social Psychology Bulletin, 25*, 450–462.

Weiner, B. (1985). An attributional theory of achievement motivation and emotion. *Psychological Review, 92*, 548–573.

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.