CHAPTER

7

THE CONSONANCE MODEL OF DISSONANCE REDUCTION

Thomas R. Shultz
McGill University

Mark R. Lepper
Stanford University

THE SEARCH FOR CONSISTENCY

The pursuit of consistency among one's beliefs and attitudes has long been taken as a sign of human rationality (Abelson, 1971), and consistency-seeking has held a prominent place in social-psychological theorizing for over 50 years (e.g., Festinger, 1957; Heider, 1946, 1958; McGuire, 1960; Newcomb, 1953; Osgood & Tannenbaum, 1955). This widespread concern with issues of cognitive consistency culminated in the late 1960s in the publication of an 84-chapter tome, *Theories of Cognitive Consistency*, edited by six of the major social-psychological theorists of the day (Abelson et al., 1968). Theoretical ideas about cognitive consistency, particularly those involving cognitive dissonance (Aronson, 1969; Festinger, 1957; Wicklund & Brehm, 1976) and cognitive balance (Heider, 1946, 1958; Rosenberg & Abelson, 1960), have been thoroughly investigated and largely supported by systematic experimentation. Indeed, the extensive literature on cognitive consistency continues to constitute one major part of the foundation of contemporary social psychology.

However, the study of cognitive consistency seems to have fallen out of favor (Aronson, 1989; Berkowitz & Devine, 1989). With a few prominent exceptions (e.g., Cooper & Fazio, 1984; Steele, 1988; Thibodeau & Aronson, 1992), little empirical work has been done in this area over the past 20 years. Perhaps this decline was due to the fact that so many aspects of the various consistency theories were already so extensively explored; perhaps it was

PREFACE

Neural network models, also called connectionist or parallel distributed processing models, seem to represent a major paradigm shift in cognitive psychology, cognitive science, and artificial intelligence. Such models move us away from the idea of mind as computer, and instead promise the possibility of brain style models of the mind, admitting the possibility that models of high level cognitive processing can be built from simple neuron-like units. That is, we can build computational models of the mind composed of units functionally similar to the physical units that compose a real brain. This approach has led to some fundamental new insights about the way the mind might work and the way it might interact with the environment.

Surprisingly, given the importance of these models, until recently social psychologists had paid little attention to them. Yet, these models directly address several fundamental characteristics of social perception and social interaction: the simultaneous integration of multiple pieces of information and the quite short time frame within which such integration occurs. Any mundane act of social perception (and any resulting behavior) results from the simultaneous integration of multiple pieces of information, such that the meaning of each piece of information mutually influences and constrains the meaning of each other piece. Thus, social perception can be viewed as the solution of simultaneous mutually interacting constraints. Moreover, this integration typically takes place in a very short time frame, much shorter than would be possible for any kind of reasonable serial integration process. Thus, much of social perception must occur in parallel. Both of due to an inability to penetrate further the reasoning mechanisms that underlie the search for cognitive consistency.

Recently, we proposed that the striving for cognitive consistency can be understood in terms of constraint satisfaction, a process of simultaneous adjustment of beliefs and attitudes to satisfy as many internal and external constraints as possible (Shultz & Lepper, 1992, 1996, 1997). We illustrated this theoretical reinterpretation with an artificial neural network computer model—the consonance model—that captures many of the findings in the major paradigms of cognitive dissonance theory.

The present chapter summarizes and evaluates this enterprise. The potential payoff for this work would be to offer a novel theoretical interpretation of some of the most basic phenomena in social psychology at an abstract, yet mathematically specified, level that potentially can be unified with many other constraint satisfaction phenomena in psychology. Such a model may also lead to the generation of new predictions for empirical research on important consistency phenomena.

CLASSICAL DISSONANCE THEORY

Classical cognitive dissonance theory postulates that dissonance is a psychological state of tension that people are motivated to reduce (Festinger, 1957). Any two cognitions are said to be dissonant when, considered by themselves, one of them follows from the obverse of the other. The amount of dissonance is defined by the ratio of dissonant to total relevant (i.e., dissonant plus consonant) relations, with each relation being weighted for its importance to the person. Cognitive dissonance can, in principle, be reduced by decreasing the number and/or the importance of the dissonant relations and by increasing the number and/or the importance of consonant relations. The manner in which dissonance will actually get reduced in a particular situation is hypothesized to depend on the resistance to change of the various relevant cognitions, with less resistant cognitions being more likely to change. Resistance, in turn, derives from the extent to which a cognitive change would be likely to produce new dissonance, the degree to which a cognition is firmly anchored in reality, and the difficulty of changing those aspects of reality.

Drawing on these principles, researchers sought to create experimental situations in which dissonance could be produced and, in turn, its reduction channeled into particular predictable changes in attitudes or beliefs. This was often accomplished by the creation of dissonant relationships between people's attitudes or beliefs and their overt actions, producing pressure on them to alter their attitudes or beliefs to justify or fit their presumably irrevocable actions. These studies showed that, under the right circumstances, the result

of supposedly rational consistency seeking could be demonstrably irrational behavior. Ironically, as Aronson (1969) put it, dissonance theory portrayed people more as "rationalizing" than as "rational" creatures.

Cognitive dissonance theory proved to be enormously successful. Within five years of the publication of Festinger's original book (1957), the theory had already become the most influential of the consistency models. By now, dissonance theory has generated more than 1,000 published studies covering a great variety of content domains including attitude change, decision rationalization, and responses to belief disconfirmation (Cooper & Fazio, 1984; Thibodeau & Aronson, 1992). The general success of the theory, and its specific ability to encompass both apparently rational and apparently irrational responses, coupled with the possibility that the underlying mechanisms were based on constraint satisfaction, seemed to us to recommend it highly as a candidate for our modeling efforts.

THE CONSONANCE MODEL

Our consonance model is based on the idea that dissonance reduction can be viewed as a constraint satisfaction problem. That is, the motivation to seek cognitive consistency, which is postulated by dissonance and related theories, can be viewed as imposing constraints on the beliefs and attitudes that an individual holds at a given moment (Abelson et al., 1968; Abelson & Rosenberg, 1958; Feldman, 1966). Such consistency problems can be resolved by the satisfaction of a number of soft constraints that can vary in their relative importance. Soft constraints are those that are desirable, but not essential, to satisfy. The fact that constraints can conflict with each other favors softness rather than hardness of constraints and suggests that it is unlikely that inconsistencies can be fully eliminated from most complex belief systems.

Consonance networks correspond to a person's representation of the situation created by the experimental setting for each condition in a particular cognitive dissonance experiment. Units in a network can be variously active, corresponding loosely to the firing rate of neurons. Activations of units represent the direction and strength of the person's attitudes and beliefs. Units can also differ in their resistance to change, reflecting differences in the extent to which particular cognitions may be supported by other cognitions or anchored in reality. Connection weights between cognitions represent psychological implications among a person's beliefs and attitudes. The connections between any two units can be excitatory (+), inhibitory (–), or nonexistent (0). Both unit activations and connection weights can vary across the different conditions of a particular experiment.

Consonance is the degree to which similarly evaluated units are linked by excitatory weights and oppositely valued units are linked by inhibitory weights. Activations change over time cycles in order to satisfy the various constraints and increase consonance. More formally, the consonance contributed by a particular unit i is

$$consonance_i = \sum_j w_{ij} a_i a_j. \tag{1}$$

where w_{ij} is the weight between units i and j, a_i is the activation of the receiving unit i, and a_j is the activation of the sending unit j.

Consonance over the whole network is the sum of the values given by Equation 1 over all receiving units in the network.

$$consonance_n = \sum_i \sum_j w_{ij} a_i a_j. \tag{2}$$

Activation spreads around the network over time cycles in conformity with two update rules:

$$a_i(t+1) = a_i(t) + net_i[ceiling - a_i(t)], \text{ when } net_i \geq 0. \tag{3}$$

$$a_i(t+1) = a_i(t) + net_i[a_i(t) - floor], \text{ when } net_i < 0. \tag{4}$$

where $a_i(t+1)$ is the activation of unit i at time $t+1$, $a_i(t)$ is the activation of unit i at time t, *ceiling* is the maximal level of unit activation, *floor* is the minimal level of unit activation, and net_i is the net input to unit i, defined as

$$net_i = resist_i \sum_j w_{ij} a_j. \tag{5}$$

The parameter $resist_i$ indicates the resistance of receiving unit i to having its activation changed. Smaller values of this parameter indicate greater resistance because smaller values mean less impact of the net input.

At each time cycle, n units are randomly selected and updated according to Equations 3–5. Typically, n is the number of units in the network. The update rules in Equations 3–5 ensure that consonance increases or stays the same across cycles. When consonance reaches asymptote, the updating process is typically stopped.

MAPPING THE CONSONANCE MODEL TO DISSONANCE PHENOMENA

A generic consonance network is used to instantiate any particular dissonance experiment. To date, our simulations used between two and four cognitions per network, but in principle, there is no limit to the number of

cognitions that could be included. Each cognition falls into one of three categories: behaviors, justifications, or evaluations. As noted later, these three categories are differentially resistant to change during dissonance reduction. Our consonance computer program, written in the Common Lisp language, enables us to specify a network including the relevant cognitions, their types and initial activations, and the relations among the cognitions. A concrete example of the start of a network run is given in the next section in which our simulations are presented. At this point, let us turn to a presentation of the set of six theoretical principles that map dissonance theory to the consonance model.

Representation of Cognitions

Principle 1 specifies that a cognition is implemented by the net activation of a pair of negatively connected units, one of which represents the positive pole and the other represents the negative pole. Net activation for the cognition is the difference between activation of the positive unit and activation of the negative unit. Activations range from a floor to a ceiling. In our simulations, the floor parameter is 0 by default. The default ceiling parameter for positive poles is 1 and for negative poles is 0.5. Use of two different default ceilings is based partly on neurological and computational considerations (Anderson, 1995) and partly on the fact that it works well in the domain of cognitive consistency (Shultz & Lepper, 1996). This bipolar representation scheme allows for some degree of ambivalence in cognitions, although the inhibitory connections between the two poles do tend to discourage ambivalence.

Relationships Among Cognitions

Mapping Principle 2 specifies that cognitions are connected to each other based on their causal implications. If quantitative increases in cognition 1 would cause quantitative increases in cognition 2, then cognition 1 is said to be a positive cause of cognition 2. If quantitative increases in cognition 1 would cause quantitative decreases in cognition 2, then cognition 1 is said to be a negative cause of cognition 2. Connection weights range from −1 to 1, with 0 representing a lack of causal relation. The connection scheme for two generic cognitions is illustrated in Fig. 7.1. When two cognitions are positively related, their positive poles are connected with excitatory weights, as are their negative poles; inhibitory weights connect the positive pole of one cognition with the negative pole of the other cognition (Fig. 7.1a). These connections are reversed for cognitions that are negatively related (Fig. 7.1b). Each unit has an inhibitory self-connection specified by the *cap* parameter, and all connection weights are bi-directional. Connection weights

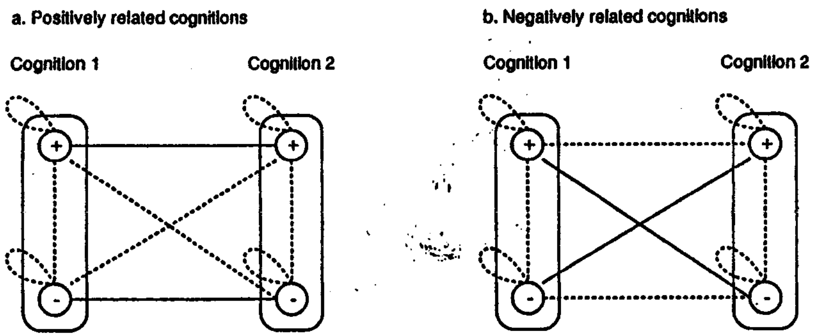

FIG. 7.1. In the consonance model, any two cognitions can be positively (a) or negatively (b) related. Excitatory (positive) weights are portrayed by solid lines, inhibitory (negative) weights by dashed lines. Each cognition is portrayed by a rounded rectangle drawn around the positive and negative poles of the cognition. Each pole is represented as a unit in the network, drawn here as a circle. All connection weights are bi-directional.

have a default value of 0.5, representing a strong or high connection, but we occasionally use a weak or low connection, with a default value of 0.1.

Magnitude of Dissonance

Mapping Principle 3 specifies that total dissonance is the negative of total consonance divided by r, the number of nonzero inter-cognition relations:

$$dissonance = \frac{-consonance_{ll}}{r}. \quad (6)$$

Dividing by r standardizes dissonance across networks by controlling for the number of relevant relations. Self-connections, w_{ii}, are excluded from this computation of dissonance so that dissonance is not an artifact of amount of activation. This definition of dissonance differs from Festinger's (1957) because it is formalized, it measures the amount of dissonance in each inter-cognition relation, it includes within-cognition ambivalence, and it can still vary when all relations are dissonant or all relations are consonant. Our definition of dissonance is analogous to Hopfield's (1982, 1984) notion of energy, and our definition of consonance is analogous to Rumelhart, Smolensky, McClelland, and Hinton's (1986) notion of goodness. Energy (or dissonance) decreases as goodness (or consonance) increases.

Reduction of Dissonance

Principle 4 specifies that networks tend to settle into more stable, less dissonant states as unit activations are updated according to Equations 3, 4, and 5. There are two parameters that affect the dissonance reduction

process, *cap* and *rand%*. A *cap* parameter with a default of –0.5, corresponding to the value of the connection between each unit and itself, w_{ii}, prevents activations from growing to their ceiling. This sort of activation limitation seems appropriate for most dissonance experiments, which do not typically deal with life and death situations. At the start of each run of a network, connection weights, resistances, caps, and initial activations are all randomized by adding or subtracting a random proportion of their initial amounts. The *rand%* parameter specifies the proportion range in which additions or subtractions are randomly selected under a uniform distribution. Typically, we use small (.1), medium (.5), and large (1.0) levels of *rand%*. This increases psychological realism in the sense that not everyone can be expected to share precisely the same parameter values. The randomization of weight values also violates connection weight symmetry such that $w_{ij} \neq w_{ji}$ and thus increases the instability of network solutions. Perhaps most importantly, comparisons of the solutions obtained from networks at various levels of randomization provide a clear indication of the robustness of the results across parameter variations.

Changes in Cognitions

Principle 5 specifies that cognition unit activations, but not connection weights, are allowed to change, and that some cognitions are more resistant to change than others, as implemented in Equation 5. Beliefs (including behaviors and justifications) are more resistant to change than are evaluations. Although participants in dissonance experiments are likely to be well aware of what just happened to them and what they just did, they may not be so sure of how they feel about aspects of the somewhat novel situation they are in. The resist parameter has default values of 0.5 for low and 0.01 for high resistance. As specified in Equation 5, the larger the resistance multiplier, the more readily the unit changes its activation. Additional details about the consonance model and discussion of its various assumptions are presented in Shultz and Lepper (1996).

Importance of Dissonance

Recently, we added a sixth mapping principle to account for a variety of arousal and self-concept phenomena in the dissonance literature (Shultz & Lepper, 1997). This new Principle 6 concerns the psychological importance of the dissonant situation. An importance scalar parameter, with default values of 0.5, 1.0, or 1.5, multiplies all connection weights and unit activations at the start of each run, before the initial randomizations referred to under Principle 4. An importance parameter value of 1.0 is used in control conditions; a value of 0.5 is used for conditions that lessen the importance of a dissonant situation; and a value of 1.5 is used for conditions that enhance the importance of a dissonant situation.

At this point, we move into a review of some of our simulations of cognitive dissonance phenomena.

INSUFFICIENT JUSTIFICATION PARADIGMS

Perhaps the most widely researched and cited experimental paradigm in cognitive dissonance concerns situations that involve psychologically insufficient justification. The "insufficient justification" paradigm deals with situations in which participants are led to engage in some counter-attitudinal action with either rather little or considerably greater justification for that action. Classical dissonance theory predicts that the less the justification for the behavior, the greater the dissonance and, at least when it is difficult to retract one's action, the more people will be motivated to change their attitudes so as to provide additional justification for their actions. Within insufficient justification research, there have been three major sub-paradigms (Lepper, 1983); prohibition, initiation, and forced compliance. Within each of these sub-paradigms, our simulations focused on second-generation experiments that avoided both the plethora of alternative explanations that plagued the original classic experiments and the more narrow and highly qualified interests of most contemporary studies.

Insufficient Justification via Prohibition

The classic study on insufficient justification in a prohibition situation was done by Aronson and Carlsmith (1963) who forbade nursery schoolers from playing with an attractive toy under either mild or severe threat of punishment. In later ratings, children devalued the forbidden toy more in the mild than in the severe threat condition, presumably as a result of greater dissonance under mild threat. In order to rule out a number of alternative explanations of these initial results, Freedman (1965a) included surveillance conditions in this experiment in addition to the original non-surveillance conditions. In the surveillance conditions, the experimenter stayed in the room while the child played with the alternative, less desirable toys. Presumably, the presence of the experimenter lessened dissonance by increasing the potency of the threats. Several weeks later, in another setting, only the children in the mild threat, non-surveillance condition showed significant devaluation of the forbidden toy.

Network specifications for our simulations (Shultz & Lepper, 1996) of the non-surveillance and surveillance conditions of Freedman's experiment are shown in Figs. 7.2a and 7.2b, respectively. There were three cognitions:

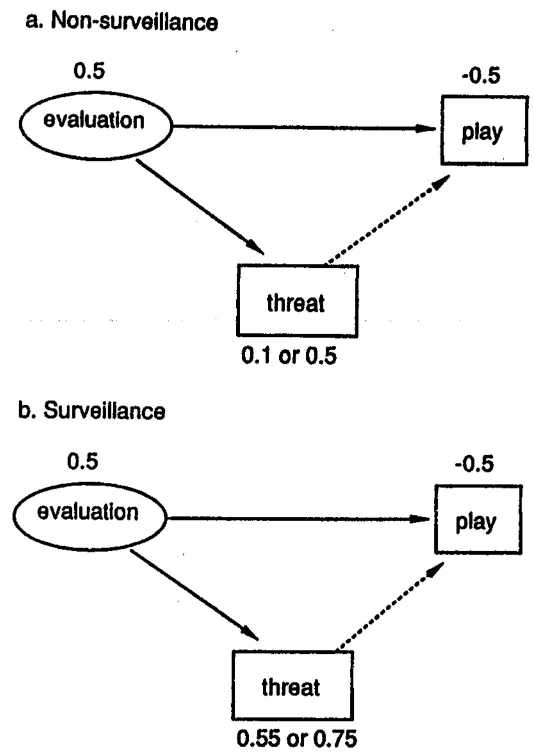

FIG. 7.2. Network specifications for simulation of the Freedman (1965a) experiment. Cognitions with high resistance to change are drawn as rectangles; cognitions with low resistance to change are drawn as ellipses. Positive causal implications are drawn with solid arrows; negative causal implications are drawn with dashed arrows. Each arrow is drawn from cause to effect. Initial activation values are shown next to each cognition. Some details are suppressed in the networks displayed in Figs. 7.2–7.9, in that each cognition is actually implemented with a pair of units and that causal implications among cognitions are implemented with a set of bi-directional connection weights as illustrated in Fig. 7.1.

evaluation of the toy, play with the toy, and threat. Initially, the toy was given a high positive evaluation, play was given a high negative evaluation because it was not done, and threat was either low or high, depending on condition. In the surveillance conditions, the value of these threats was scaled up using a multiplier of 0.5 plugged into Equation 3. The new threat was computed as the old threat plus one half of the difference between 1 and the old threat. Relations among the three cognitions reflected assumed causal relations. The better liked the toy is, the more it would be played with; the better liked the toy, the more threat would be required to prevent play; and the larger the threat, the less the play.

Net evaluation of the toy was computed as the difference between activation of the positive pole and the negative pole of the toy cognition. As shown in Fig. 7.3, after unit activation changes and dissonance reduction reached asymptotic values, the networks showed more devaluation of the forbidden toy under mild than under severe threat, but only in the non-surveillance conditions, thus mirroring the children's data. Figure 7.3 (as well as other figures presenting simulation results in this chapter) shows results for small levels of parameter randomization (*rand%* = .1).

Examples of actual activation updates are shown in Table 7.1 over the first two update cycles from one network run in Freedman's mild threat condition. (For purposes of illustration, the cycles presented in Table 7.1 do not involve any randomization of initial values.) There were six updates per cycle, the number of units in the network. Within each update cycle, the units to be updated were randomly selected. Unit 3, the positive pole of the play cognition, happened to have been selected for the first update. The initial activation for this unit was 0.000, because the toy was not played with. The net input to this unit was –0.050, as specified in Equation 5, before being scaled by the resistance parameter. Net input to a unit is computed as the sum of products of activations on sending units and the connection weights. In this case, there were only three such nonzero products: from unit 1 (activation of 0.5 × weight of 0.5), from unit 4 (activation of 0.5 × weight of –0.5), and from unit 5 (activation of 0.1 × weight of –0.5). Summing these three products yields an unscaled net input of –0.050. Multiplying this net

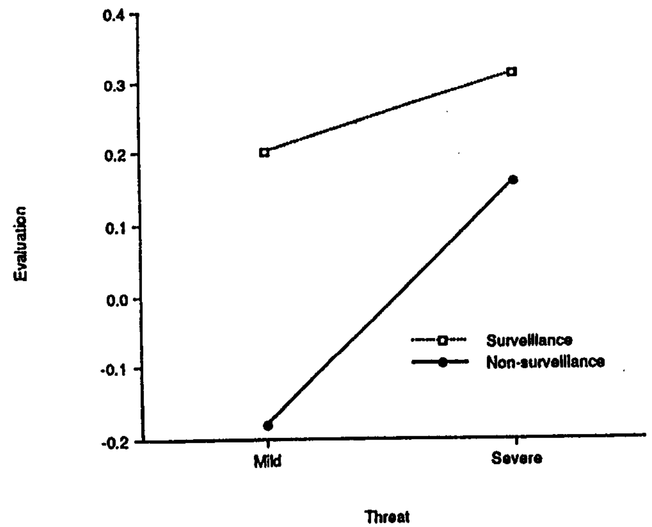

FIG. 7.3. Mean evaluation of the forbidden toy in the simulation of the Freedman (1965a) experiment. Adapted from Shultz and Lepper (1996).

TABLE 7.1
Activation Updates Over the First Two Update Cycles for a Network in Freedman's Mild Threat, Non-surveillance Condition as Designed in Fig. 7.2a

Unit number (randomly selected)	*Unit name*	*Current unit activation*	*Net input to unit*	*Scaled by unit resistance*	*Distance to floor or ceiling*	*Updated unit activation*
3	play+	0.000	–0.050	0.000	0.000	0.000
6	threat–	0.000	–0.550	–0.006	0.000	0.000
4	play–	0.500	–0.450	–0.005	0.500	0.498
1	evaluation+	0.500	–0.449	–0.224	0.500	0.388
4	play–	0.498	–0.393	–0.004	0.498	0.496
6	threat–	0.000	–0.492	–0.005	0.000	0.000
1	evaluation+	0.388	–0.392	–0.196	0.388	0.312
3	play+	0.000	–0.142	–0.001	0.000	0.000
4	play–	0.496	–0.354	–0.004	0.496	0.494
6	threat–	0.000	–0.453	–0.005	0.000	0.000
5	threat+	0.100	0.353	0.004	0.900	0.103
5	threat+	0.103	0.351	0.004	0.897	0.106

input by the resistance scalar of 0.01 yields –0.0005, which in this Table with only three decimal places rounds to 0.000. The updated activation of this unit was thus still 0.000. The actual running of these networks carried many more decimal places, so this example should be considered as only a rough approximation of what really happens. The interested reader can follow through the next few updates to gain a sense of how the program works. At each update, scaled net input is multiplied by the distance to the floor (in the case of negative net input) or the distance to the ceiling (in the case of positive net input) in the last column of Table 7.1, in conformity with Equations 5 and 6.

Insufficient Justification via Initiation

A second insufficient justification sub-paradigm generated by dissonance theory concerns the consequences of having to suffer an initiation in order to join a group. The initial example of an initiation experiment found that people initiated into a boring group liked the group better after undergoing a severe than after a mild initiation (Aronson & Mills, 1959). The classical dissonance explanation was that the greater dissonance created by a painful initiation could be reduced by increasing one's evaluation of the group for which one had suffered. A follow-up experiment by Gerard and Mathewson (1966) eliminated various alternative explanations for the results of the first experiment, in part by adding non-initiation conditions with the same severe and mild levels of unpleasantness. For this purpose, they used electric shock at two different levels, administered either as part of an initiation or as part of an unrelated experiment. Following an overheard boring discussion by a group they had volunteered to join, participants rated the group. As predicted by dissonance theory, people liked the group better after a severe than after a mild initiation. However, in the non-initiation condition, participants liked the group better after receiving mild shock than after receiving severe shock, a finding not predicted by dissonance theory, which speaks only to the initiation conditions.

Network specifications for simulation of the initiation and non-initiation conditions of the Gerard and Mathewson experiment (Shultz & Lepper, 1996) are presented in Figs. 7.4a and 7.4b, respectively. There were three cognitions: evaluation of the group, the behavior of joining the group, and the level of shock received. Joining received an initially positive evaluation because participants did volunteer to join. Evaluation of the group received an initially negative value reflecting the boring discussion that was over-

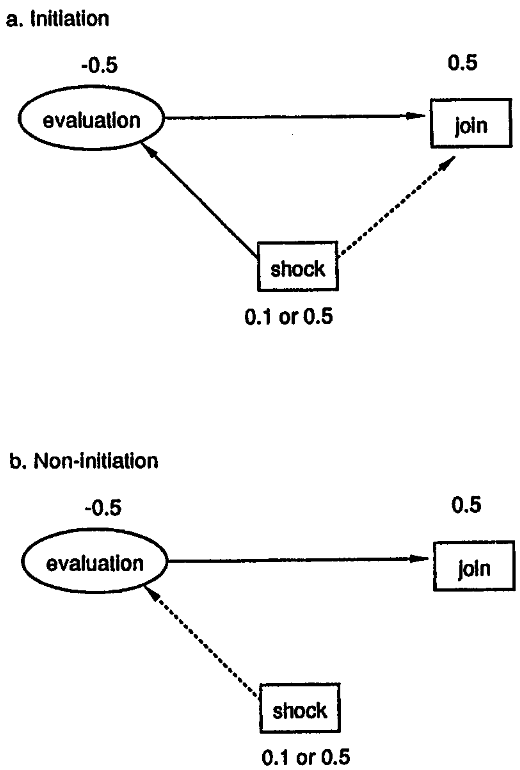

FIG. 7.4. Network specifications for simulation of the Gerard and Mathewson (1966) experiment.

heard. Shock levels were high or low depending on condition. Relations among the cognitions reflected causal implications: The better you like a group, the more likely you are to join it; the more you pay for something, the more you get; and the more you have to pay, the less likely you are to join.

Network design for the non-initiation conditions was similar except that relations between shock and joining were cut to 0 because, without an initiation, there was no longer a causal relation between them, and relations between shock and evaluation were changed from positive to negative because, with shock no longer paying for joining, the negative experience of being shocked would likely affect how one felt about the whole experimental session.

Net liking for the group was computed as the difference between the positive and negative poles of the evaluation cognition. As shown in Fig. 7.5, when dissonance reduction reached asymptote, attitude results showed the same interaction Gerard and Mathewson found with human subjects: a dissonance reduction effect under initiation conditions, with more liking for the group after severe than after mild initiation; and the reverse, annoyance effect without any initiation. Again, this interaction held up over a wide range of parameter randomization. Thus, the consonance model accounts for the full interaction, whereas classical dissonance theory accounts only for the effects of shock under initiation conditions.

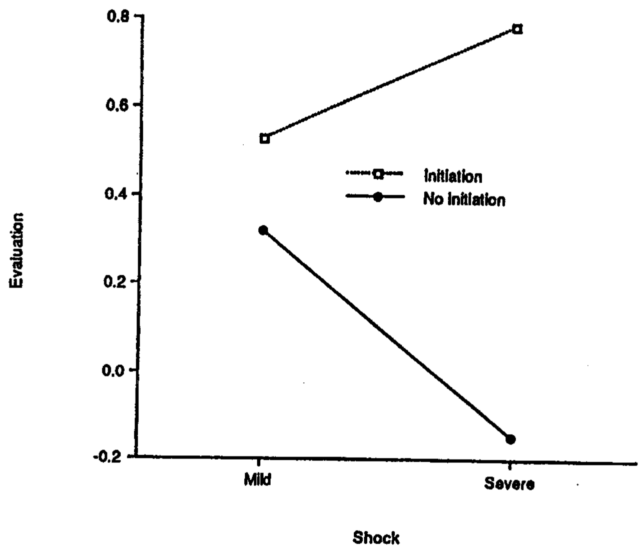

FIG. 7.5. Mean evaluation of the group in the simulation of the Gerard and Mathewson (1966) experiment. Adapted from Shultz and Lepper (1996).

Insufficient Justification via Forced Compliance

The third, and perhaps most famous, insufficient justification sub-paradigm involves what has been called "forced compliance." In the first experiment of this type, Festinger and Carlsmith (1959) found that smaller inducements to voice a belief against one's own attitude led to more change in the direction of the statement than did large inducements. Being paid $1 to lie about how interesting a truly boring task was led to higher subsequent evaluations of the task than did being paid $20 to tell the same lie. Because the receipt of a payment is consonant with telling the lie, it should be more dissonant to lie for a small reward than for a large reward. To rule out a number of alternative explanations, a subsequent study by Linder, Cooper, and Jones (1967) added conditions in which participants were not given a choice about writing a counter-attitudinal essay. They found that the dissonance effect held only when participants had a choice about whether to write the counter-attitudinal essay; without a choice, the opposite effect was found, with higher payment leading to more attitude change in the direction of the view supported in the essay.

Network specifications for this simulation (Shultz & Lepper, 1996) are shown in Fig. 7.6. Again, there were three relevant cognitions: the attitude, writing the essay, and the payment. Initial attitude was high negative, reflecting these liberal college students' negative reactions to banning controversial speakers on campus. Writing the essay was set to high positive because the essay had indeed been written. Payment for writing was either high or low, depending on condition.

Relations among the three cognitions reflected assumed causal implications. In the choice condition, the more one supports the position in the essay, the more likely one would write the essay; the more one is paid, the more one would agree to write the essay; and the more favorable one's attitude, the less one would need to be paid to write an essay at some particular level of support. Relations among cognitions were the same in the no-choice condition, except that the relation between attitude and essay was 0 because, without a choice, there is no causal connection between the two, and the relation between attitude and payment was positive, reflecting a better mood with higher pay.

Attitude toward banning controversial speakers was computed as the difference between activations on the positive and negative poles of the attitude cognition. After dissonance reduction was completed, the networks revealed the same crossover interaction found with Linder et al.'s (1967) college students: a dissonance effect under choice, with more attitude change after low payment than after high payment; and the reverse, mood effect under no choice, with more attitude change after high payment than after low payment (Fig. 7.7). As with the other insufficient justification simu-

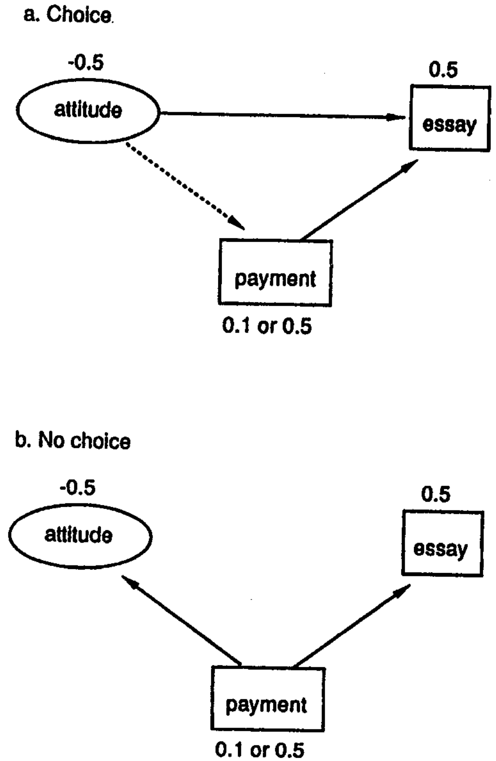

FIG. 7.6. Network specifications for simulation of the Linder et al. (1967) experiment.

lations, the results held up at every level of parameter randomization, but were cleaner at small and medium levels than at high levels. Once again, the consonance model covers the full interaction, whereas classical dissonance theory covers only the dissonance effect under choice.

THE FREE-CHOICE PARADIGM

The second major dissonance paradigm, the free-choice situation, focuses on the effects of people making a choice between alternatives. Making such a decision, it was argued, should create dissonance, due to the fact that the chosen alternative is never perfect and the rejected alternative often has desirable aspects that are necessarily foregone when an irreversible choice is made. Once an irrevocable choice has been made, dissonance can be reduced by viewing the chosen object as more desirable and by viewing the rejected object as less desirable. Such dissonance reduction further sepa-

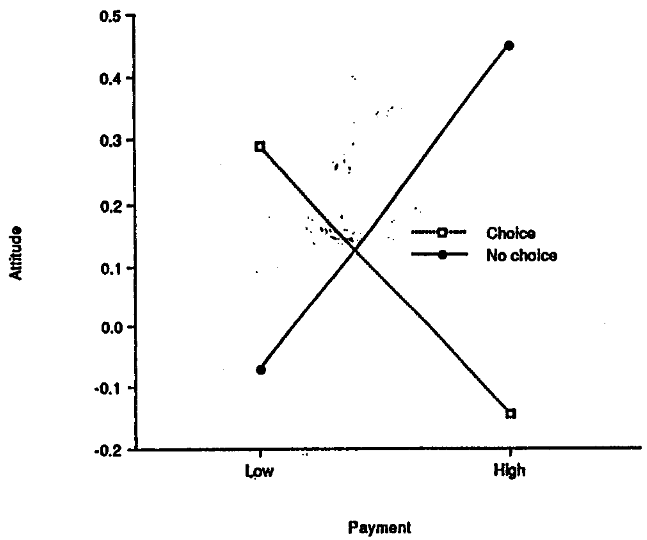

FIG. 7.7. Mean attitude toward the view expressed in the essay in the simulation of the Linder et al. (1967) experiment. Adapted from Shultz and Lepper (1996).

rates the alternative choices in terms of their desirability. Moreover, the magnitude of dissonance should be greater the closer the alternatives are in desirability before the choice is made. The closer the alternatives are in their initial desirability, the more difficult an exclusive choice between them.

In the classic free-choice experiment, Brehm (1956) found that the greater the dissonance created by a choice, the more the increase in separation between the alternatives after the choice was made. There was more separation between alternatives after a difficult choice between two highly preferred alternatives than after an easy choice between a favored and a disliked alternative. Unlike the more counter-intuitive insufficient justification studies, Brehm's free-choice results did not attract many alternative explanations.

However, in planning our simulations, we realized that a third interesting condition could be added, in which participants would be offered a difficult choice, but between two disliked alternatives (Shultz & Lepper, 1996). We referred to this as the difficult/low condition because evaluation of both alternatives starts out relatively low, in contrast to the difficult condition of Brehm (1956) that offered a choice between top-rated alternatives. We referred to that condition as difficult/high. Network specifications for these free-choice simulations are shown in Fig. 7.8. There were three cognitions, two of them evaluations of the alternative objects, and one of them the decision itself. The decision had an initially high value, and the initial values of the alternatives depended on condition. There was a positive relation

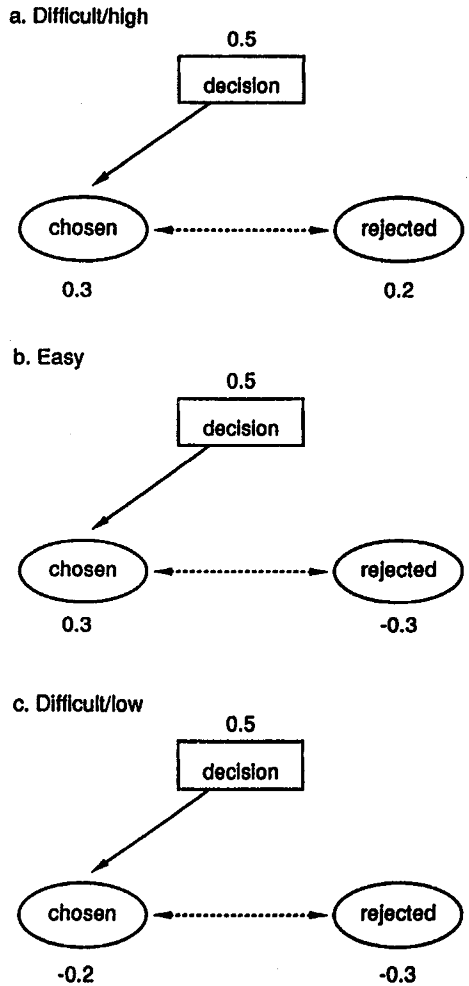

FIG. 7.8. Network specifications for simulation of the Brehm (1956) and Shultz et al. (1996) experiments.

between the decision and the chosen object, reflecting the choice, and a negative relation between the two alternatives reflecting the fact that they were competing for an exclusive choice.

Following Brehm (1956), we computed evaluation differences as final evaluation minus initial evaluation for each alternative. When dissonance reduction in the simulations reached asymptote, evaluation of the chosen object had increased, but more so in the difficult/low condition, and evaluation of the rejected object had decreased, but more so in the difficult/high condition (Fig. 7.9). As parameter randomization increased, the interaction weakened statistically but the pattern of evaluation change remained fairly

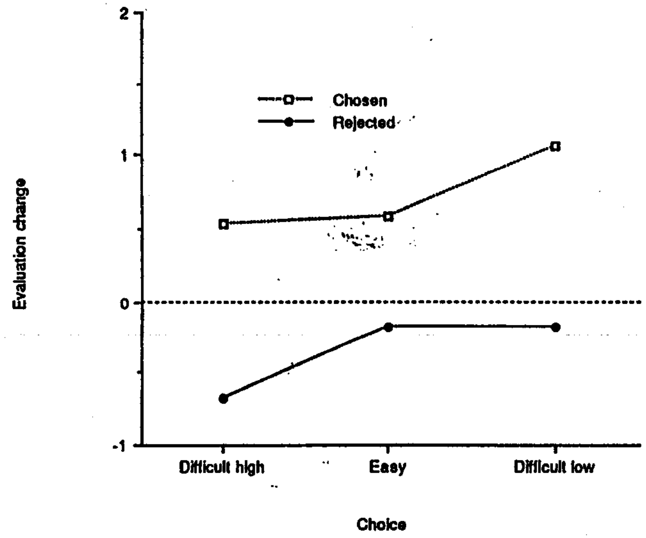

FIG. 7.9. Mean evaluation change in the simulation of the Brehm (1956) and Shultz et al. (1996) experiments. Adapted from Shultz and Lepper (1996).

constant. A new experiment with 13 year olds choosing among posters under these three conditions replicated Brehm's results in the difficult/high and easy choice conditions, and confirmed the predictions of the consonance networks for the difficult/low choice condition (Shultz, Léveillé, & Lepper, 1996).

It is worth noting that the consonance model simulations describe Brehm's (1956) data better than did classical dissonance theory, which only predicts more separation following difficult than following easy choices. Brehm found that most of the action was produced by devaluation of the rejected object in the difficult choice condition, and that was closely simulated by our consonance networks.

AROUSAL AND FORCED COMPLIANCE

More recently, considerable attention has been paid to the study of the arousal properties of cognitive dissonance (e.g., Cooper & Fazio, 1984; Zanna & Cooper, 1974). An important central finding of this work was that dissonance arousal can be externally modulated by administration of a drug, such as an amphetamine, tranquilizer, or placebo. Cooper, Zanna, and Taves (1978), for example, asked university students to write counter-attitudinal essays under either high- or low-choice conditions. These students had taken a pill that they had been led to believe was a placebo. In different conditions,

however, the pill actually contained either phenobarbital, amphetamine, or a placebo. The placebo condition produced the usual dissonance effect: more attitude change in the direction of the essay under high choice than under low choice. In the tranquilizer condition, this dissonance effect was eliminated. In contrast, the ingestion of amphetamine enhanced attitude change under both high- and low-choice conditions. This last finding represented the first time that insufficient justification via forced compliance had produced dissonance effects under low-choice conditions.

Network specifications for our simulation (Shultz & Lepper, 1996) of the Cooper et al. (1978) experiment are shown in Fig. 7.10. In a sense, this study was a scaled-down version of Linder et al.'s (1967) forced compliance experiment, without the payment cognition, and with low, rather than no, choice about whether to write the essay. As in the simulation of Linder et al.'s results, the relation between attitude and essay was positive. Rather than cutting this link to 0, as we did to implement the no choice condition of Linder et al., we used a low value for low choice in contrast to a high value for high choice. To instantiate the arousal manipulations, an importance parameter scaled the initial activations and connection weights before they were randomized. This importance scalar was 1.0 for the placebo condition, 0.5 for what we called the downer condition, and 1.5 for what we called the upper condition. This corresponds very roughly to evidence that drugs such as phenobarbital depress neural firing rates and synaptic trans-

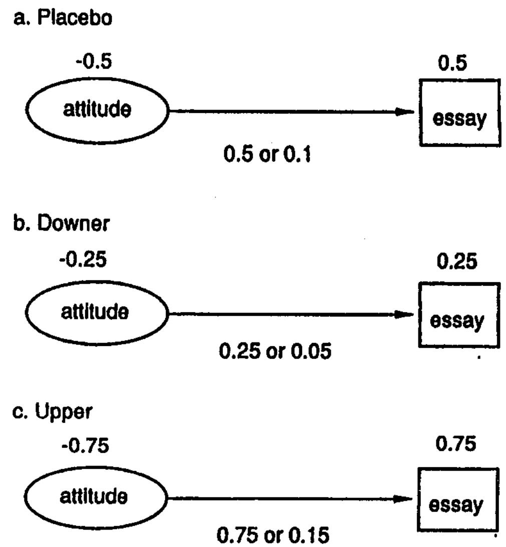

FIG. 7.10. Network specifications for simulation of the three conditions of the Cooper et al. (1978) experiment. Implicational links between attitude and essay are higher for high choice than for low choice.

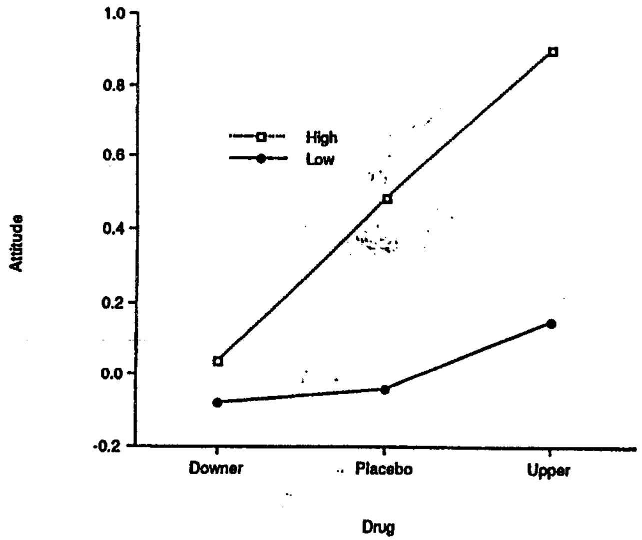

FIG. 7.11. Mean attitude toward the view expressed in the essay in the simulation of the Cooper et al. (1978) experiment. Adapted from Shultz and Lepper (1997).

mission whereas drugs such as amphetamine enhance neural firing rates and synaptic transmission (Quastel, 1975).

Attitude toward the view described in the essay was computed as the difference between activations on the positive and negative poles of the attitude cognition. After the networks had stabilized, the drug by choice interaction effects on attitude mirrored those displayed by Cooper et al.'s university students. There was a dissonance effect in the placebo condition, with more attitude change under high than under low choice; no effect of choice and little attitude change in the downer condition; and enhanced attitude change with the dissonance choice effect in the upper condition (Fig. 7.11). The interaction held for all three levels of parameter randomization.

THE ROLE OF SELF-CONCEPT

There is a contemporary focus on the importance of the self-concept in the arousal of dissonance (Steele, 1988; Thibodeau & Aronson, 1992). The basic idea is that dissonance occurs when behavior is inconsistent with the person's self-concept. Because most people have a positive self-concept, such behaviors as lying or arguing for a position that is contrary to one's own beliefs will arouse dissonance. However, people with a negative self-concept might not experience dissonance from engaging in such behaviors.

Self-Concept in the Forced Compliance Paradigm

Early support for this view can be seen in a study by Epstein (1969), which examined potential differences in the arousal of cognitive dissonance among people scoring high versus low on the trait of Machiavellianism. People scoring low on this trait, it was argued, should experience dissonance in a situation of insufficient justification via forced compliance, but those scoring high on the trait should not. For high Machiavellians, lying or writing a counter-attitudinal essay would not be inconsistent with their self-concept; instead, these sorts of actions would be seen as legitimate tactics of effective social interaction.

Epstein's undergraduate participants all initially supported the fluoridation of water supplies, but those in a dissonance group were induced to give a speech against it. These participants read some anti-fluoridation arguments and were paid $2 to give an anti-fluoridation speech. Participants in a control condition gave no speech, but read the arguments against fluoridation and were paid the $2 anyway.

Consistent with the predictions of classical dissonance theory, Epstein found that attitude change toward an anti-fluoridation position was higher for low Machiavellians who gave a speech than for low Machiavellians who did not give a speech. High scorers on Machiavellianism showed the reverse trend, which Epstein explained by citing evidence that these people are more susceptible to factual arguments than are people scoring low on Machiavellianism.

Network specifications for our simulation (Shultz & Lepper, 1997) of Epstein's experiment are shown in Fig. 7.12. These specifications are identical to those used in our simulations of Linder et al.'s (1967) forced compliance experiment, except that (a) the relation between attitude and speech was low for high Machiavellians (because, for high Machiavellians, there is not such a strong relation between true attitude and public statement), (b) there was an extra cognition for anti-fluoridation arguments (initialized to low) that had a positive relation to attitude (low for low Machiavellians, and high for high Machiavellians), (c) there was no speech cognition in the no-speech condition, and (d) pay was uniformly low.

Attitudes after asymptotes were reached mirrored the crossover interaction found with Epstein's undergraduates. Networks with low Machiavellian parameters revealed a dissonance effect, that is, more attitude change with a speech than without a speech, whereas networks built with high Machiavellian parameters exhibited the reverse effect (Fig. 7.13). The reversal effect for Machiavellians cannot be predicted or explained within classical dissonance theory, but did occur with the various constraints that can be built into consonance networks.

Most recently, even more subtle effects of self-concept on dissonance in a forced compliance situation were found by Steele (1988) in his studies of

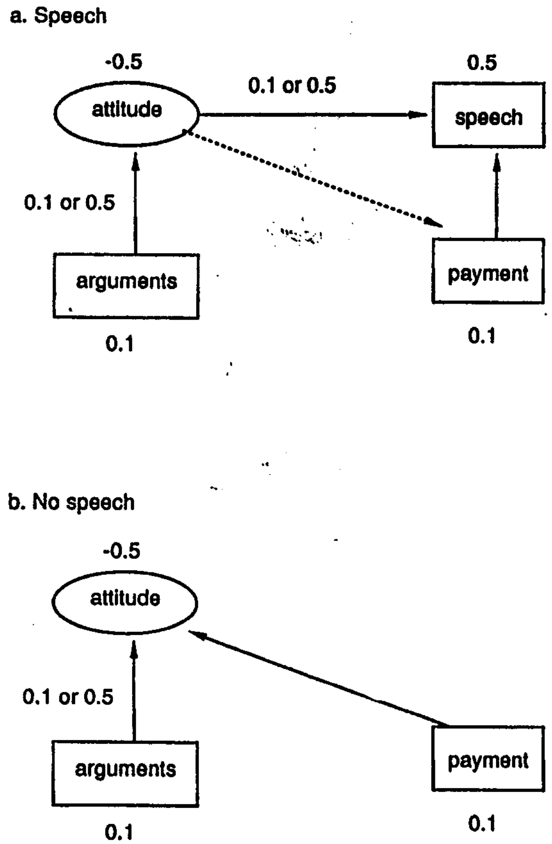

FIG. 7.12. Network specifications for simulation of the Epstein (1969) experiment. For high Machiavellians, the implication between attitude and speech is 0.1 and that between arguments and attitude is 0.5. For low Machiavellians, the implication between attitude and speech is 0.5 and that between arguments and attitude is 0.1.

self-affirmation processes. If dissonance arises when people's actions threaten their self-concepts, then anything that affirms an important aspect of the self-concept, even if it was completely irrelevant to the source of any experimentally induced inconsistency, should minimize the need for people to reduce dissonance via attitude change. Steele's college students were selected for their strong opposition to a tuition hike. These students were then persuaded to write essays in support of a substantial tuition hike, under either high choice or low choice, without any payment for doing so. Some of the students were also previously assessed as having a strong economic-political value orientation. For them, completing a political value scale would

presumably affirm a valued part of their self-concept; for others without this value orientation, this task should have little impact. Indeed, Steele found the familiar dissonance effect of more attitude change under high choice than under low choice conditions, but he also found that self-affirmation eliminated attitude change, even under high choice conditions.

Network specifications for our simulations of Steele's experiment (Shultz & Lepper, 1997) are given in Fig. 7.14. As with the experiment of Cooper et al. (1978), this was a scaled-down version of Linder et al.'s (1967) forced compliance experiment. The key cognitions were attitude and essay with a positive relation between them. As with the Cooper et al. experiment, this positive relation was high under high choice and low under low choice. Initial values were high negative for attitude and high positive for essay. Here, the self-affirmation manipulation, designed to minimize the importance of the dissonance produced by the counter-attitudinal behavior, was implemented with an importance scalar of 0.5, as in the downer condition of the Cooper et al. simulation. This dampened all of the unit activations and connection weights prior to the randomization of parameter values.

After these consonance networks stabilized, they showed the same results as Steele's undergraduate participants: more attitude change in the standard high-choice condition than in either the low-choice or the high-choice/self-affirmation conditions (Fig. 7.15).

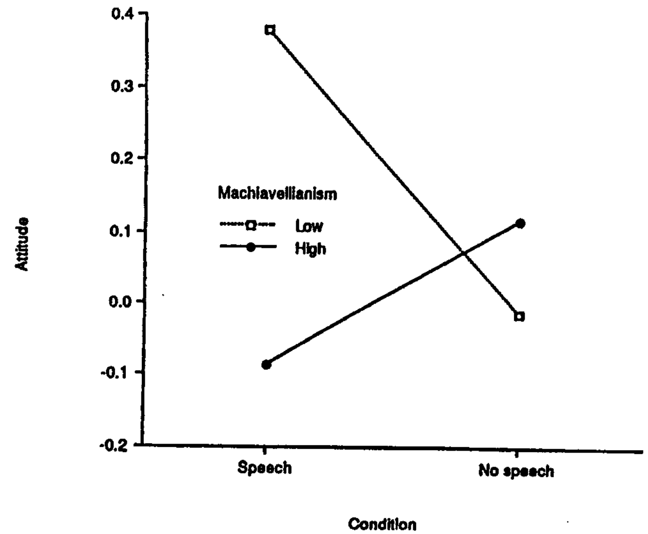

FIG. 7.13. Mean attitude toward the view expressed in the speech in the simulation of the Epstein (1969) experiment. Adapted from Shultz and Lepper (1996).

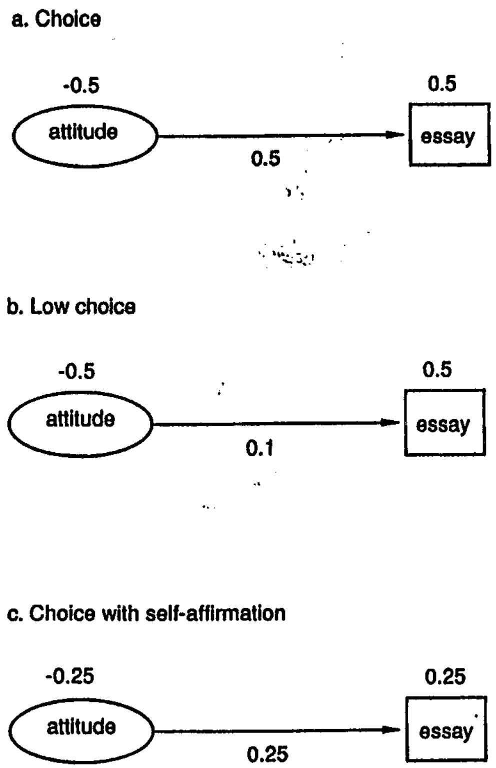

FIG. 7.14. Network specifications for simulation of Steele's (1988) forced compliance experiment. An importance scalar of 0.5 was used in the choice with self-affirmation condition.

Self-Concept in the Free-Choice Paradigm

In a similar vein, Steele (1988) also reported a free-choice experiment that showed self-affirmation effects. In this study, participants rated and ranked 10 record albums and then were given a choice of keeping either their fifth- or sixth-ranked albums. Some of the participants were selected for having a strong scientific value orientation and for having indicated that a lab coat symbolized for them important values and goals; others were selected for whom science was not an important personal value. Within each of these two selected groups, half the participants were asked to wear a lab coat for the rest of the experiment, whereas the others were not. Later, all participants rated the 10 albums once again.

The spreading apart of alternatives after the choice was computed by adding the increase in the value of the chosen item and the decrease in the value of the rejected item. Spread of alternatives was found to be lower in

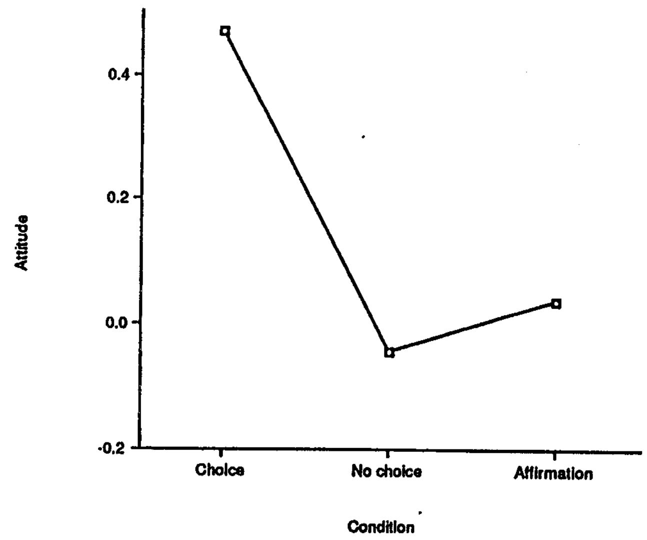

FIG. 7.15. Mean attitude toward the view expressed in the essay in the simulation of Steele's (1988) forced compliance experiment. Adapted from Shultz and Lepper (1996).

the self-affirmation condition, in which scientifically oriented students wore lab coats, than in the other conditions. This result occurred even though the affirmation procedure was irrelevant to the choice that initially produced dissonance.

Network specifications for the standard conditions of this second Steele study are shown in Fig. 7.16 (Shultz & Lepper, 1997). As with the simulations of Brehm's free-choice experiment, there were three principal cognitions: evaluations of each of the two alternatives, and the decision itself. Only a difficult choice was simulated, but otherwise the initial activations and relations were the same as in our simulation of Brehm's experiment (cf. Fig. 7.8). Here, once again, an importance scalar of 0.5 was used to diminish all unit activations and connection weights in the self-affirmation condition.

After stabilization of these consonance networks, spread of evaluation between the two alternative objects was computed as in Steele (1988). Change in the value of one object was the difference between its initial value and its value after stabilization. Spread of evaluation was the sum of the increase in the value of the chosen alternative and the decrease in the value of the rejected alternative. As with Steele's college students, at each level of parameter randomization, there was a smaller spread of the alternatives in the self-affirmation condition than in any of the other three conditions (Fig. 7.17). All of these results on self-concept held at all three levels of parameter randomization.

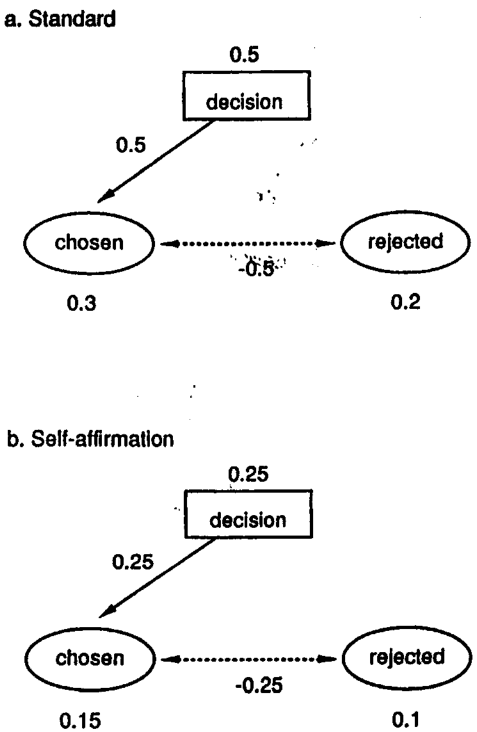

FIG. 7.16. Network specifications for simulation of the standard and self-affirmation conditions of Steele's (1988) free-choice experiment.

THE SELECTIVE EXPOSURE PARADIGM

A third basic paradigm featured in the initial presentation of dissonance theory (Festinger, 1957) involved the phenomenon of selective exposure, referring to biases in the manner in which people seek or avoid additional information that is relevant to a choice they made. In order to reduce cognitive dissonance, people were supposed to prefer information supporting their choices and to avoid information contradicting their choices. Compared to the more successful insufficient justification and free-choice paradigms, however, selective exposure generated much more long-term controversy. This is because of many results that either failed to support the selective exposure predictions or that directly contradicted them by finding a relative preference for dissonant information. Indeed, early reviews of the selective exposure literature concluded that dissonance theory was inapplicable to selective exposure phenomena (Bem, 1967; Freedman & Sears, 1965).

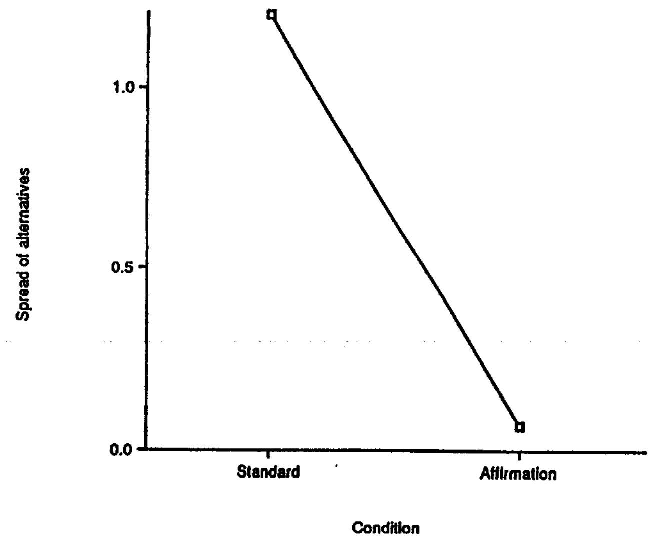

FIG. 7.17. Mean spread of alternatives in the simulation of Steele's (1988) free-choice experiment. Adapted from Shultz and Lepper (1997).

However, subsequent research (reviewed in Frey, 1986) inspired by theoretical reformulations (Festinger, 1964) permitted a somewhat more optimistic appraisal of the ability of dissonance theory to deal with selective exposure effects. These theoretical revisions emphasized, in particular, that there should be a relative preference for dissonant information when this information is perceived to be easily refutable or useful for future decision making (Festinger, 1964).

Replicable Selective Exposure Effects

Nonetheless, because of the controversial nature and relative fragility of selective exposure effects, it is probably wise to consider simulating only results in these areas that have at least one successful reported replication. Two of the phenomena in the selective exposure literature that would seem to meet this requirement concern the effects of variations in the magnitude of dissonance and those of a variety of moderating variables, including the utility, refutability, and credibility of the new consonant or dissonant information.

Selective Exposure and Magnitude of Dissonance. Selective exposure effects are predicted to increase with the magnitude of dissonance, at least in the case of irreversible decisions (Festinger, 1957, 1964). When a decision can be reversed, however, the relation between dissonance and selective exposure is presumably curvilinear, in an inverse U shape. This reflects the idea that, as dissonance approaches a maximum, dissonant information is particularly useful in guiding a decision reversal (Frey, 1986).

One of the first confirmations of the positive relationship between dissonance and selective exposure for irreversible decisions was provided by Mills (1965). Starting with procedures like those used in standard free-choice experiments, Mills asked college women to rank 20 personal care products on desirability and then to make a choice between two of them. In a high dissonance condition, the choice was between two objects that had been ranked very high; in a low dissonance condition, the choice was between one object that had been ranked high and one that had been ranked low. Then, as a measure of selective exposure, participants were asked to rate their interest in reading advertisements for each of the products involved in the choice. As predicted, interest in reading advertisements for the chosen product (i.e., information consonant with the choice) was greater when dissonance was high than when dissonance was low. Contrary to dissonance theory predictions, however, there was no effect of magnitude of dissonance on interest in advertisements for the rejected product (i.e., information dissonant with the choice).

This positive relation between dissonance and preference for consonant information was replicated by Frey (1981). His high school participants ranked 14 books and then were allowed to choose between two of them. Three levels of dissonance were created. High dissonance involved a choice between the second- and third-ranked books, medium dissonance a choice between the second- and seventh-ranked books, and low dissonance a choice between the second- and 13th-ranked books. After the choice was made, participants were allowed to select three commentaries from among six consonant and six dissonant commentaries on their chosen book. The mean number of consonant commentaries chosen was 2.00 in the high dissonance condition, 1.50 for medium dissonance, and 1.41 for low dissonance. The linear trend indicating more preference for consonant information with higher dissonance was statistically significant. As in the Mills (1965) experiment, there was no such effect of magnitude of dissonance on avoidance of dissonant commentaries, contrary to dissonance theory predictions.

Selective Exposure and Utility of Dissonant Information. The direction and strength of selective exposure effects have also been demonstrated to be influenced by a variety of extraneous variables including the utility, refutability, and credibility of the new information (see Frey, 1986, for a review of these studies). In general, dissonant information might actually be preferred to the extent that it is expected to prove useful in future tasks, to be easy to refute, or to originate from a low-credibility source. Some studies

of utility have also included a condition in which consonant information is made useful, with the result that the selective exposure effect is enhanced rather than reversed. Potentially, each of these three moderating variables could be simulated in our consonance networks in the same way, by an extra cognition positively connected to the dissonant information cognition (or to the consonant information cognition). Consequently, and because published experiments tend to manipulate only one of these variables at a time, it is probably necessary to only include one of these variables in simulations. Utility of the extra information might be a good choice because its effects are particularly well documented.

One of the first studies of utility in selective exposure was reported by Canon (1964). Male college students selected one of two possible solutions to a business problem. Then they were told that they would take part in either a debate or a presentation of their point of view concerning this decision. To prepare for this session, participants were given the opportunity to read some number of five articles relevant to the case they had decided upon. Two of these articles supported their decision, two contradicted their decision, and one was neutral. Participants were asked to rate their interest in reading each of these articles on a 100-point scale. Preparing for a debate would make the dissonant information more useful, whereas preparing for a presentation of one's own views would render the consonant information more useful. Participants were more interested in consonant information when it was useful and somewhat more interested in dissonant information when it was useful. There was also an orthogonal manipulation of the subject's confidence, with the result that more confidence led to more interest in dissonant information. However, because this effect was not replicated (Freedman, 1965b), it may not be wise to try to capture it in simulations.

Interest in whatever information was more useful was replicated by Freedman (1965b) who also added an equal utility condition in which participants were expecting to merely react to other evaluations of the case. Again, information preference corresponded to the manipulation of information utility. When only one type of information was useful, that type was preferred.

Difficulty of Simulating Selective Exposure Phenomena

Although it might well be worth trying to simulate these selective exposure phenomena, we have not so far had success doing so with our consonance network models. For us, as for many of the early dissonance theorists, selective exposure proved to be a far more elusive target than the other central paradigms presented in Festinger's original book (1957).

Our simulation attempts underscored, though, one feature of selective exposure that was somewhat obscure in the psychological literature. This is the idea that selective exposure deals principally with any residual dissonance that may be present after evaluative change has occurred. That is, only if the relatively more direct process of evaluative change fails to eliminate dissonance would selective exposure operate.

Although this view became clear to us only through attempting to simulate selective exposure effects, it is consistent with both Festinger's (1957) original emphasis on dissonance as a pre-condition of selective exposure and the fundamental distinction between dissonance reduction via evaluative change and dissonance reduction via selective exposure that was drawn in more recent literature. Frey (1986) pointed out that evaluative change does not require any new behavior on the part of the person; an attitude is simply adjusted. In contrast, selective exposure does require additional behavior in terms of further interaction with the environment. Frey's notion that selective exposure is oriented toward the future is consistent with our view that it is primarily useful in dealing with any residual dissonance after dissonance reduction via evaluative change is complete.

Thus, it would appear that at least some of the fragility of selective exposure effects in the psychological literature may have been due to possible inadvertent dissonance reduction via evaluative change. It would seem that substantial dissonance reduction by evaluative change would interfere with selective exposure effects.

DISCUSSION

Using constraint satisfaction neural networks, our consonance model of dissonance reduction successfully captured all of the basics and many of the subtleties of both insufficient justification and free-choice phenomena (Shultz & Lepper, 1996, 1997). These captured effects include the devaluation of a forbidden toy under mild, but not under severe threat, and under non-surveillance, but not under surveillance conditions; greater liking of a group after severe than after mild initiation into the group; more attitude change following small inducements to make counter-attitudinal statements than following larger inducements; and separation of alternatives after a free-choice. Drawing upon later dissonance literature, they also include enhancement of attitude change in forced compliance experiments under stimulating drugs, and elimination of attitude change in forced compliance experiments under calming drugs; elimination of dissonance effects in forced compliance for participants with either Machiavellian personalities or recent self-affirming experiences; and elimination of dissonance effects in free-choice experiments for participants who have just experienced self-affirmation.

In several of these paradigms, our consonance model fit the psychological data better than did classical cognitive dissonance theory. Examples of superior fits to the human data include mood or annoyance effects in

insufficient justification via initiation or forced compliance, locus of evaluation change in the free-choice paradigm, and greater attitude change among Machiavellians after not giving a speech than after giving a counter-attitudinal speech. These superior data fits were due to the inclusion of constraints that were not part of classical dissonance theory and to the increased precision inherent to this constraint satisfaction approach.

Consonance constraint satisfaction networks thus seem both more general and more precise in their theoretical coverage than is classical cognitive dissonance theory. In addition, novel predictions generated by the consonance network model for a free-choice between undesirable alternatives were confirmed in a new psychological experiment (Shultz et al., 1996). It is, of course, particularly challenging to generate predictions for new research in an area, such as cognitive dissonance, that has been worked over so thoroughly for so many years. Nevertheless, we hope that this new view of dissonance theory will stimulate further fresh investigations of it and other cognitive consistency phenomena. A particularly fruitful line of predictions might involve more complicated scenarios with more cognitions and relations than found in experiments inspired by dissonance theory. Being a verbally formulated theory, dissonance theory is limited in terms of its ability to deal with complex scenarios.

Perhaps the largest potential payoff for this simulation work, however, is the theoretical unification that can follow successful simulations (Smith, 1996). Constraint satisfaction models have been successfully applied to a wide variety of psychological processes, including belief revision, explanation, comprehension, schema completion, analogical retrieval and mapping, and content-addressable memory storage and retrieval (Holyoak & Thagard, 1989; Kintsch, 1988; Rumelhart et al., 1986; Sloman, 1990; Thagard, 1989). A number of phenomena in social psychology have also been successfully modeled with constraint satisfaction networks, including attitude change, impression formation, and cognitive balance (Kunda & Thagard, 1996; Read & Miller, 1993, 1994; Spellman & Holyoak, 1992; Spellman, Ullman, & Holyoak, 1993). Indeed, in our own present work, we are extending our consonance model to a more complete analysis of the case of cognitive balance theory (Heider, 1946, 1958; Read & Miller, 1994; Rosenberg & Abelson, 1960).

The success of constraint satisfaction models across these varied research areas suggests that cognitive dissonance may not be as unique and exotic as it has often appeared to be. Partly because dissonance researchers chose to focus upon those counter-intuitive dissonance phenomena that could not be easily explained by other models, dissonance theory was historically viewed as distinct from basic and perhaps more mundane psychological processes. Once dissonance and its reduction are viewed as another instance of constraint satisfaction, much of its mystery disappears. Viewing dissonance theory in this new light makes it more understandable within the general scheme of psychological explanation.

Indeed, the use of these sorts of models may help us to remember that dissonance reduction, and the related consonance-seeking processes postulated by other consistency theories, are neither inherently rational nor inherently irrational. Whether the outcome of such a process is one or the other in any given case will depend upon the task and the situation. The same willingness to disregard dissonant feedback that allowed the Wright brothers to succeed in pioneering manned flight also led a number of their predecessors to leap off cliffs while flapping the artificial wings attached to their arms. Perhaps our aspirations to rationality and our propensity for rationalization, these models suggest, are but two sides of the same coin.

ACKNOWLEDGMENTS

This research was supported by a grant to the first author from the Social Sciences and Humanities Research Council of Canada, and by grant MH-44321 to the second author from the U.S. National Institute of Mental Health.

Thomas R. Shultz, Department of Psychology, McGill University, 1205 Penfield Avenue, Montreal, Quebec, Canada H3A 1B1. E-mail: shultz@psych.mcgill.ca. Mark R. Lepper, Department of Psychology, Stanford University, Jordan Hall, Building 420, Stanford, CA 94305-2130. E-mail: lepper@psych.stanford.edu.

REFERENCES

Abelson, R. P. (1971). Are attitudes necessary? In B. T. King & E. McGinnies (Eds.), *Attitudes, conflict, and social change* (pp. 19–32). New York: American Marketing Association.

Abelson, R. P., Aronson, E., McGuire, W. J., Newcomb, T. M., Rosenberg, M. J., & Tannenbaum, P. H. (Eds.). (1968). *Theories of cognitive consistency: A sourcebook*. Chicago: Rand McNally.

Abelson, R. P., & Rosenberg, M. J. (1958). Symbolic psycho-logic: A model of attitudinal cognition. *Behavioral Science, 3*, 1–13.

Anderson, J. A. (1995). *An introduction to neural networks*. Cambridge, MA: MIT Press.

Aronson, E. (1969). The theory of cognitive dissonance: A current perspective. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 4, pp. 1–34). Orlando, FL: Academic Press.

Aronson, E. (1989). Analysis, synthesis, and treasuring the old. *Personality and Social Psychology Bulletin, 15*, 508–512.

Aronson, E., & Carlsmith, J. M. (1963). Effect of severity of threat on the devaluation of forbidden behavior. *Journal of Abnormal and Social Psychology, 66*, 584–588.

Aronson, E., & Mills, J. (1959). The effect of severity of initiation on liking for a group. *Journal of Abnormal and Social Psychology, 59*, 177–181.

Bem, D. J. (1967). Self perception: An alternative interpretation of cognitive dissonance phenomena. *Psychological Review, 74*, 183–200.

Berkowitz, L., & Devine, P. G. (1989). Research traditions, analysis, and synthesis in social psychology: The case of dissonance theory. *Personality and Social Psychology, 15*, 493–507.

Brehm, J. W. (1956). Post-decision changes in the desirability of choice alternatives. *Journal of Abnormal and Social Psychology, 52*, 384–389.

Canon, L. K. (1964). Self-confidence and selective exposure to information. In L. Festinger (Ed.), *Conflict, decision, and dissonance* (pp. 83–95). Stanford, CA: Stanford University Press.

Cooper, J., & Fazio, R. H. (1984). A new look at dissonance theory. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 17, pp. 229–266). New York: Academic Press.

Cooper, J., Zanna, M. P., & Taves, P. A. (1978). Arousal as a necessary condition for attitude change following forced compliance. *Journal of Personality and Social Psychology, 36*, 1101–1106.

Epstein, G. F. (1969). Machiavelli and the devil's advocate. *Journal of Personality and Social Psychology, 11*, 38–41.

Feldman, S. (Ed.). (1966). *Cognitive consistency*. New York: Academic Press.

Festinger, L. (1957). *A theory of cognitive dissonance*. Evanston, IL: Row, Peterson.

Festinger, L. (1964). *Conflict, decision, and dissonance*. Stanford, CA: Stanford University Press.

Festinger, L., & Carlsmith, J. M. (1959). Cognitive consequences of forced compliance. *Journal of Abnormal and Social Psychology, 58*, 203–210.

Freedman, J. L. (1965a). Long-term behavioral effects of cognitive dissonance. *Journal of Experimental Social Psychology, 1*, 145–155.

Freedman, J. L. (1965b). Confidence, utility, and selective exposure: A partial replication. *Journal of Personality and Social Psychology, 2*, 778–780.

Freedman, J. L., & Sears, D. O. (1965). Selective exposure. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 2, pp. 57–97). New York: Academic Press.

Frey, D. (1981). Reversible and irreversible decisions: Preference for consonant information as a function of attractiveness of decision alternatives. *Personality and Social Psychology Bulletin, 7*, 621–626.

Frey, D. (1986). Recent research on selective exposure to information. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 19, pp. 41–80). New York: Academic Press.

Gerard, H. B., & Mathewson, G. C. (1966). The effects of severity of initiation on liking for a group: A replication. *Journal of Experimental Social Psychology, 2*, 278–287.

Heider, F. (1946). Attitudes and cognitive organization. *Journal of Personality, 21*, 107–112.

Heider, F. (1958). *The psychology of interpersonal relations*. New York: Wiley.

Holyoak, K. J., & Thagard, P. (1989). Analogical mapping by constraint satisfaction. *Cognitive Science, 13*, 295–355.

Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences, USA, 79*, 2554–2558.

Hopfield, J. J. (1984). Neurons with graded responses have collective computational properties like those of two-state neurons. *Proceedings of the National Academy of Sciences, USA, 81*, 3008–3092.

Kintsch, W. (1988). The role of knowledge in discourse comprehension: A construction-integration model. *Psychological Review, 95*, 163–182.

Kunda, Z., & Thagard, P. (1996). Forming impressions from stereotypes, traits, and behaviors: A parallel-constraint-satisfaction theory. *Psychological Review, 103*, 284–308.

Lepper, M. R. (1983). Social-control processes and the internalization of values: An attributional perspective. In E. T. Higgins, D. N. Ruble, & W. W. Hartup (Eds.), *Social cognition and social development* (pp. 294–330). New York: Cambridge University Press.

Linder, D. E., Cooper, J., & Jones, E. E. (1967). Decision freedom as a determinant of the role of incentive magnitude in attitude change. *Journal of Personality and Social Psychology, 6*, 245–254.

McGuire, W. J. (1960). A syllogistic analysis of cognitive relationships. In C. I. Holland & M. J. Rosenberg (Eds.), *Attitude organization and change* (pp. 65–111). New Haven, CT: Yale University Press.

Mills, J. (1965). Effect of certainty about a decision upon postdecision exposure to consonant and dissonant information. *Journal of Personality and Social Psychology, 2*, 749–752.

Newcomb, T. M. (1953). An approach to the study of communicative acts. *Psychological Review, 60*, 393–404.
Osgood, C. E., & Tannenbaum, P. H. (1955). The principle of congruity in the prediction of attitude change. *Psychological Review, 62*, 42–55.
Quastel, J. H. (1975). Effect of drugs on energy metabolism of the brain and on cerebral transport. In L. L. Iversen, S. D. Iversen, & S. H. Snyder (Eds.), *Handbook of psychopharmacology: Vol. 5. Synaptic modulators* (pp. 1–46). New York: Plenum.
Read, S. J., & Miller, L. C. (1993). Rapist or "regular guy"?: Explanatory coherence in the construction of mental models of others. *Personality and Social Psychology Bulletin, 19*, 526–541.
Read, S. J., & Miller, L. C. (1994). Dissonance and balance in belief systems: The promise of parallel constraint satisfaction processes and connectionist modeling approaches. In R. C. Schank & E. Langer (Eds.), *Beliefs, reasoning, and decision making: Psycho-logic in honor of Bob Abelson* (pp. 209–235). Hillsdale, NJ: Lawrence Erlbaum Associates.
Rosenberg, M. J., & Abelson, R. P. (1960). An analysis of cognitive balancing. In M. J. Rosenberg, C. I. Hovland, W. J. McGuire, R. P. Abelson, & J. W. Brehm (Eds.), *Attitude organization and change* (pp. 112–163). New Haven, CT: Yale University Press.
Rumelhart, D. E., Smolensky, P., McClelland, J. L., & Hinton, G. (1986). Schemata and sequential thought processes in PDP models. In D. E. Rumelhart & J. L. McClelland (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition* (Vol. 2, pp. 7–57). Cambridge, MA: MIT Press.
Shultz, T. R., & Lepper, M. R. (1992). A constraint satisfaction model of cognitive dissonance phenomena. *Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society* (pp. 462–467). Hillsdale, NJ: Lawrence Erlbaum Associates.
Shultz, T. R., & Lepper, M. R. (1996). Cognitive dissonance reduction as constraint satisfaction. *Psychological Review, 103*, 219–240.
Shultz, T. R., & Lepper, M. R. (1997). *Consonance network simulations of arousal and self-concept phenomena in cognitive dissonance: The importance of psychological importance*. Manuscript submitted for publication.
Shultz, T. R., Léveillé, E., & Lepper, M. R. (1996). *Free choice and cognitive dissonance revisited*. Manuscript submitted for publication.
Sloman, S. (1990). *Persistence in memory and judgment: Part-set inhibition and primacy*. Unpublished doctoral dissertation, Stanford University.
Smith, E. R. (1996). What do connectionism and social psychology offer each other? *Journal of Personality and Social Psychology, 70*, 893–912.
Spellman, B. A., & Holyoak, K. J. (1992). If Saddam is Hitler, then who is George Bush? Analogical mapping between systems of social roles. *Journal of Personality and Social Psychology, 62*, 913–933.
Spellman, B. A., Ullman, J. B., & Holyoak, K. J. (1993). A coherence model of cognitive consistency: Dynamics of attitude change during the Persian Gulf war. *Journal of Social Issues, 49*, 147–165.
Steele, C. M. (1988). The psychology of self-affirmation: Sustaining the integrity of the self. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 21, pp. 261–302). New York: Academic Press.
Thagard, P. (1989). Explanatory coherence. *Behavioral and Brain Sciences, 12*, 435–502.
Thibodeau, R., & Aronson, E. (1992). Taking a closer look: Reasserting the role of the self-concept in dissonance theory. *Personality and Social Psychology Bulletin, 18*, 591–602.
Wicklund, R. A., & Brehm, J. (1976). *Perspectives on cognitive dissonance*. Hillsdale, NJ: Lawrence Erlbaum Associates.
Zanna, M. P., & Cooper, J. (1974). Dissonance and the pill: An attribution approach to studying the arousal properties of dissonance. *Journal of Personality and Social Psychology, 29*, 703–709.

CHAPTER

8

Toward an Integration of the Social and the Scientific: Observing, Modeling, and Promoting the Explanatory Coherence of Reasoning

Michael Ranney
University of California, Berkeley

Patricia Schank
SRI International

It may seem odd for two cognitive scientists, each with little specific expertise in social psychology, to present a chapter that focuses on social cognition. Indeed, our past work may seem much more in the realm of scientific reasoning than in that of social reasoning. But one question that we have been asking, both of ourselves and of our colleagues, is, "What is the difference between 'scientific reasoning' and plain old 'reasoning'?" Generally, people hem and haw when confronted with this question, then speak of the latter as if it were social reasoning—and quite often, they mention socially based ruminations that involve suboptimal decisions, faulty heuristics, and inappropriately biased values, goals, and the like (see Gigerenzer, 1991; Tversky & Kahneman, 1974, and many others). Useful follow-up questions to such respondents include, "Well, is the difference between these two sorts of reasoning qualitative or quantitative?" Put another way (as many—including Einstein, 1950—seem to have occasionally wondered), "Is scientific reasoning just (a) more likely to employ formal tools (like deduction or mathematics) and/or (b) more likely to involve the vigilant search for disconfirmation—something that 'just plain folks' (Lave, 1988, p. 4) do, but less frequently?"

Put rather bluntly, we have not been able to reject the hypothesis that the word "scientific" in "scientific reasoning" is superfluous. In an era of