

**Memories, Thoughts,
and Emotions:**
*Essays in Honor of
George Mandler*

Edited by

William Kessen

Yale University

Andrew Ortony

Northwestern University

Fergus Craik

University of Toronto



1991

LAWRENCE ERLBAUM ASSOCIATES, PUBLISHERS
Hillsdale, New Jersey Hove and London

22

Value and Emotion

Andrew Ortony
Northwestern University

In some remote causal way, I am writing this piece as a result of the fact that my first serious introduction to the psychology of emotion was George Mandler's book *Mind and Emotion* (Mandler, 1975). My initial reaction upon reading it was, I recall, rather mixed. I felt it was interesting and largely believable, but I also found it frustrating and incomplete. I began corresponding with George on a number of issues raised by *Mind and Emotion*, and soon, with considerable encouragement from him, I got going on the study of emotion. I started doing what George felt somebody, but not he, should be doing, namely, I started thinking about the cognitive antecedents of the different emotions. My colleagues and I worked on this problem for nearly ten years, finally publishing our efforts in a book entitled *The Cognitive Structure of Emotions* (Ortony, Clore, & Collins, 1988). One of the main themes of that work was the idea that there are three classes of emotions. This first class consists of emotions resulting from appraisals rooted in goals, the second from appraisals rooted in standards and norms, and the third from appraisals grounded in tastes and attitudes.

In the early 1980s, George began to devote serious attention to the notion of value. This interest was primarily stimulated by his belief that value plays a central role in emotion, an opinion with which I am in strong agreement. Indeed, at least the last two of the three classes of emotions that my colleagues and I proposed clearly implicate what George thinks of as values. But the view that values are central to emotions does not appear to be one that is widely held by emotion theorists—certainly, it is difficult to find books on emotion with index entries for “value,” and coherent discussions of the psychological, especially cognitive, foundations of value are few and far between. George was disturbed by this neglect.

In taking on the problem of value, George was grappling with an immensely difficult and understudied topic. In *Approaches to a Psychology of Value* (Mandler, 1989) he bemoans the fact that psychologists have paid so little attention to the issue. He wrote, referring to his earlier efforts to bring some ideas about it into the public forum: "The purpose [of Mandler (1982)] was in part to expose some budding ideas to public inspection, but also to invite and generate some discussion on the problem itself. The invitation was apparently declined. Psychology in general is still valueless (though, of course, *not* value free). I shall now try again."

In this chapter, I shall present a few of my own preliminary thoughts on the issue of value. In doing so, perhaps I can contribute in some small way to George's hope that the field begin to take the question of value more seriously. I shall take George's concluding remarks in his aforementioned work at face value: "I started by reasserting my attempt to move psychologists to discuss seriously and in a principled way the sources and conditions of value. I certainly do not claim any kind of completeness for this account, nor any certainty for the propositions advanced here. I do hope that they will be the beginning of a study of value—even if it begins by questioning what I have tried to do here" (Mandler, 1989, p. 22). George knows me well enough to know that it is inevitable that I should begin by questioning what he has tried to do!

George's primary focus is on values as preferences, as likes and dislikes, rather than as moral beliefs and opinions. In this sense, values need (and often can be given) no justification. I cannot (and normally am not expected to) justify the fact that I like baseball or that I like strawberries, although sometimes it is possible to explain such an attitude. In general, however, I just like what I like, and dislike what I dislike. Of course, my liking of strawberries need not be innate—perhaps I did not like them as a child. But acquired or not, my liking of strawberries is a matter of taste (no pun intended), not of adopted choice. Such values, George wants to argue, contribute to emotional experience. On the other hand, and this is important in the present context, values do not, he asserts, arise out of emotion.

If values do not arise from emotions, where do they come from? What leads people to evaluate objects and events positively or negatively? George's basic answer to this question is in terms of schema theory—in terms of schema conformity and discrepancy (see Figure 22.1). With some caveats, George's proposal is that there are three unambiguous sources of positive affect.

First, as a limiting case, we have the situation in which some object or event (which I shall call an experience, for simplicity) is congruous with some existing schema. In this case we have positive value with, curiously, zero intensity. The next, slightly more intense case of positive value arises when the experience is slightly incongruous but, because it is only slightly so, it is readily assimilated to an existing schema. The third unambiguously positive case arises when there is severe incongruity between an input and an active schema that can be resolved by

assimilating it to an alternate schema. There remain two other cases. The first of these involves successful accommodation. That is, the experience can be assimilated only to a *changed* schema. This can give rise to either positive or negative value. In either case, George maintains that the resulting intensity is quite high. However, the most intense case is, according to George, the negative case in which a severely incongruous experience cannot be assimilated even after accommodation.

There is much more to George's approach to value than I have described; however, this structural account lies at the heart of his view, and has done so since he started seriously thinking about the issue. In the rest of this chapter I want to focus on three questions. The first is whether or not the general account that George proposes can do what he intends, the second concerns George's supposition that the relation between values and emotions is unidirectional—from values to emotions, and the third is a proposal for a model that might move us toward an account of the relation between emotions and beliefs, values, and experience.

IS SCHEMA CONFORMITY AND DISCREPANCY ENOUGH?

I agree with George that there are many occasions upon which positive and negative reactions arise as a result of the kind of schema matches and mismatches that he describes. The question that I want to address first is whether all such reactions result from the mechanisms that George proposes and whether all such mechanisms lead to such reactions. I think that the answer to both may be "no." My reason is based on the idea, which I shall elaborate in the next section, that representations often, indeed usually, include affective information (Fiske, 1982). For example, my ARMED-BURGLAR schema surely includes my distaste for such people. It is not, nor can it be, a representation that incorporates only stereotypic factual (i.e., affect-free, non-evaluative) information about things like what it is that armed burglars do and how and why they do it. The most compelling real-world reason for believing that this must be the case is the existence of prejudice.

To see the implications of affective information in schemas, it may be easier to operationalize schema conformity and discrepancy in terms of realized and violated expectations. I suspect that George would be willing to accept this, at least as an analogy. So, imagine you are lying in bed at night. You know that there have been a number of occasionally violent burglaries in your neighborhood recently. You think you hear the sound of an intruder in your house and become increasingly frightened. You listen, motionless, ever more convinced that you have the right explanation for what you hear. Suddenly, a change in the way the moonlight is illuminating the door of your bedroom leads you to realize that the




FIG. 22.1. Several possible outcomes of schema congruity and incongruity in terms of both values and affective intensity. The resultant value is shown as either positive (POS) or negative (NEG). Degree of affective intensity is shown to vary from zero to +++ (from Mander, 1982).

door is being slowly opened. You quietly pull the bedcovers over your head while peeking out to see who is coming in. As you peek, you are overcome with the fear that if you can see the intruder, he can see you. When you finally see the figure of a person wearing a ski mask with a gun in hand your suspicions that you have an armed burglar in your room are unambiguously confirmed. Your heart is pounding—you think you are going to die.

Meanwhile, back at the (cognitive) ranch, you have a perfect schema fit. Total congruity. You have no difficulty assimilating what you see to your active ARMED-BURGLAR schema. If I understand George's view correctly, this means that there ought to be positive affect resulting from this confirmed expectation. But such a conclusion does not seem credible. Surely, the only outcome that could possibly produce positive affect in such a situation would be a violation of the expectation. But if I am right that successful assimilation by the active ARMED-BURGLAR schema would lead to negative rather than positive affect, it follows that not all cases of successful assimilation to an active schema lead to affectively positive reactions. Although I shall not pursue this now, it is presumably possible to find comparable cases in which the schema congruity and incongruity conditions that George describes in fact generate affective reactions that differ in valence relative to the predictions of the model. The part of George's model that is, I think, correct, is the quantitative part concerning the relative intensity of the associated reactions, but it seems to me that we need to rethink the qualitative aspects that concern the origin of positive and negative valence. We will have to acknowledge that representations can come with affect already in them, in which case we are left with a new problem, namely, that of explaining the origin of value in schemas, and that is what this chapter is all about.

EMOTIONS, VALUES, AND TRAITS

I mentioned earlier that George believes the relation between values and emotions is unidirectional—from values to emotions. He writes (Mandler, 1989): "I have become convinced that the problem of value is at the heart of human emotional experience. I do not assert the converse, that emotion is at the heart of the problem of value. Values, even simple preferences, do not arise out of emotion; rather, they contribute to emotional experiences" (p. 4). Unfortunately, George does not explain why values do not arise out of emotions, and because his intuitions and mine diverge on this point, I shall try to make a case for the claim that in addition to values being an important source of emotions, emotions are an important source of values and, more specifically, that they can be the source of value in schemas.

Imagine that you are introduced to a man whom you have no particular desire to meet. As you learn his name and what he does, you go through the motions of

being interested in meeting him. But your initial reaction is in fact one of indifference—you neither like him nor dislike him. Maybe you build a skeletal schema for him, but because he appears to have no interesting or salient characteristics, no such characteristics are incorporated into your representation of him. Now suppose that the person starts talking and behaving in a way that you find offensive. He persists, and you become irritated, and as he continues you find yourself increasingly disliking him. Let us now compare your original attitude of indifference to your new attitude of dislike. Such a change in attitude must surely have concomitants in your representation of him. Your schema for him must contain the material necessary for this new, negative attitude—it did not have it before, but it has it now. This material might include the trait-based information that the person is very unfriendly and very aggressive, as well as the emotional information that he caused you to be upset.

The first point that I want to make is that values assigned to trait variables are direct manifestations of our values. If I do not value friendliness as a characteristic of people, then even if I can scale people (or their behaviors) on a friendly–unfriendly dimension, such a scale value cannot contribute to my dispositional liking of them. The second point is that insofar as trait representations do in fact represent values, representations of individuals in terms of traits are representations that incorporate values. The third point is that the particular constellation of traits (and the values on them) that are represented in the schema for some particular target can result from emotional interactions with the target. And finally, I want to suggest that one's dispositional liking for an individual is normally determined by one's emotional interactions with that person, independent of trait representations. In other words, if somebody makes me angry, and if a consequence of this can be that it changes my dispositional liking for that person, then assuming that dispositional liking is part of my value system, we have to conclude that emotions can influence values.

Not surprisingly, our discussion of value has already brought us into the realm of traits. This suggests that an obvious place to look for some serious discussion of value would be in the literature of social cognition, especially the part that has to do with impression formation. Unfortunately, this turns out not to be the case. The level of analysis that we need lies below that on which this literature usually focuses. Research on impression formation, person perception, and person memory, insofar as it deals with the issue of where attitudes come from, usually comes to rest on ideas such as that evaluative judgments are determined by the affective value of (currently) salient beliefs (e.g., Fishbein & Ajzen, 1975), or that they are the result of some sort of algebraic combination of weighted scale values on a set of traits (Anderson, 1971). The problem with the first kind of analysis is that it does not offer an account of the ultimate origin of the affective value of a belief. If my attitude toward or evaluation of somebody is negatively affected by learning that he is a Republican, this can only be explained if I have an explanation of why I have a negative attitude toward Republicans. The problem with the

second approach is that values on attributes are simply not sufficient to explain everyday facts. I can have scale values for a set of attributes about a person I have never met, and yet not know whether I like him.

Recent work by Srull and Wyer (e.g., Wyer & Srull, 1986; Srull & Wyer, 1989) perhaps comes closest to dealing with the issue, although they tend to view impression formation as a goal of social information processing rather than as a spontaneous product of it. Their general view is that trait information is assembled at the time that an evaluative judgment is first made by classifying behaviors in terms of the most accessible applicable traits that they exemplify. This leads to the construction of a general evaluative concept, likeableness, that becomes associated with the target and that can be used later to make an evaluative judgment without reference to the specific behaviors from which it was constructed. Srull and Wyer (1989) suggest that the mechanism for generating a general evaluative concept of an individual from trait-behavior clusters is probably similar to that proposed by Anderson (1971, 1981).

Comprehensive as their proposals are, I have a number of problems with the Srull and Wyer model. Many of these have to do with the kind of data that the model seeks to accommodate. The data are generally derived from a research paradigm that lacks ecological validity. In the typical experiment in person perception research, subjects read descriptions of the characteristics and activities of (usually fictitious) individuals and then are required to make a judgment, for example, about how likeable such an individual would be. However, because the targets are rarely real people who are personally significant to the subjects, there is no reason to suppose that subjects care one way or the other about them (Berscheid, 1982); they neither like nor dislike those about whom they are making judgments or recalling behaviors. Maybe this is why subjects are typically asked to give likeableness ratings (rather than liking ratings). Unfortunately, likeableness ratings are usually treated as surrogates for liking ratings (see, for example, Fiske & Pavelchak, 1986; Srull & Wyer, 1989). The problem is that the general evaluative concept of likeableness is nothing more than a very general trait. Although it might on occasion be the underlying representation for dispositional liking, it is not always so, and quite possibly is not even normally so. Insofar as trait information enables likeableness judgments, it enables predictions or expectations about liking. But predicting that one will like somebody is not the same as liking them. Predictions can be wrong. Furthermore, there is nothing contradictory about asserting that, for example, Ronald Reagan is likeable but that one does not like him, and it is certainly not unheard of for people to like individuals while evaluating them negatively on their most salient traits (we might call this the "you're a bastard but I love you" syndrome).

Another problem I have with the Wyer and Srull model (and many other accounts of person perception) is that they seem to presuppose a stability of social judgments that I suspect is not so common in the real social world. Thus, it is possible to interpret their model as proposing that a global evaluative social

judgment (e.g., of likeableness) is a read-out of what appears to be a relatively stable evaluative concept of the target. Now this might be an unfair criticism, but I think Wyer and Srull would acknowledge that their model does not focus on the dynamic aspects of person representation or social judgment. This is something of a problem because, as Berscheid (1982) points out, in moving from the laboratory to naturalistic settings, "the focus radically changes from accounting for stability and constancy in affective behaviors towards [a target] to accounting for instability and inconstancy" (p. 42f.).

My final discomfort with the Srull and Wyer model is that it treats traits as categorical rather than continuous constructs—a target is viewed as either friendly or unfriendly, pleasant or unpleasant. Perhaps they do not believe that traits are categorical, but their model in no way exploits them as continuous. In the model that I am about to sketch, this aspect of traits turns out to be quite important, at least at the level of representation if not at the level of conscious awareness.¹

A MODEL OF PERSON REPRESENTATION

The main purpose of the last section's brief excursion into impression formation was to set the stage for a more detailed examination of the relationship between values and emotions in which we pay more attention to the issue of how mental representations of value-laden objects might be structured. Because I find it an interesting topic in its own right, I shall restrict myself to some proposals about how we might consider individual people to be represented, rather than worrying about the representation issue with respect to objects, events, or categories.

Clearly, because we are able to make evaluative judgments about other people, our representations of those people must contain information that will enable the construction of such evaluations. It may be that in some cases, we really do directly store such an evaluation so that a specific judgment really is a simple read-out of the (current, or most recently computed) value of how much one likes (or dislikes) the target. However, we still need an account of how such stored values change and of where they come from. Central to my proposals is the idea that they do not necessarily always come from the same sort of information.

It seems to me that we need to distinguish three types of information that can, in principle, be part of a representation of a person (and possibly of other things too). One of the three parts of person representations contains the accumulated

¹In fact, the treatment of traits as categorical rather than continuous constructs appears to have been an almost intentional decision by the field. My claim is not that the internal representation of trait information is literally quantitative, but that it qualitatively represents the approximate shape of the distribution. A categorical equivalent of this would be to treat, for example, "extremely aggressive" as a different trait from "somewhat aggressive," and so on.

value-free facts (including observed behaviors) held to be true of the individual. I shall call this aspect of the representation the *fact* component, and I view it as relatively stable, normally changing only by accretion. For example, someone's representation of Margaret Thatcher might initially contain only the beliefs that she is female, in her 60s, married, and the British prime minister. These are all matters of (putative) fact, not matters of value; they are not represented as being good or bad. This is not to say that they cannot have evaluative implications, but only that if they do, these implications are not part of the fact component.

Important evaluative implications that factual knowledge may have are represented in the value-laden component of the representation, a component that I shall call the *value* component. For example, the fact that Margaret Thatcher is the British prime minister might, for some individuals, imply that because she is a politician she is likely to be very unscrupulous and ambitious. These would be inferences to traits inherited from the stereotype of a category to which Thatcher had been assigned (Fiske & Neuberg, 1990). Thus, the resulting representation might contain relatively extreme initial assignments on the evaluative dimensions scrupulous–unscrupulous and ambitious–unambitious. In the case of person representations, I envision the value component as being comprised of a set of evaluative dimensions, which we usually call traits, to which values have been assigned on the basis of inferences from directly or indirectly acquired propositional information or observed behaviors.

Whereas the value component of the representation of some individuals contains trait-based information, I think it may be helpful to suppose that at any point in time only a few trait dimensions are represented here. Specifically, I have in mind that the value component contains information pertaining to only the currently most diagnostic traits for that individual. I take a trait to be diagnostic with respect to a target if there is a history of extreme values, high variability, or if some current observation seems extreme with respect to it. So, if Smith is one of the most selfish people I know, my value component for Smith might include the selfish–generous dimension. And if I view him as sometimes (to some people, in some situations) quite friendly, and sometimes (to the same or other people, in other situations) rather unfriendly, then the friendly–unfriendly dimension might be represented in the value component. Finally, if I view Smith as being neither particularly intelligent nor particularly stupid, but suddenly he does something that causes me to think he must be remarkably stupid, then the intelligent–unintelligent dimension could become (at least temporarily) diagnostic. Such highly diagnostic information is readily accessible, and it is accessible not simply as trait identities that are particularly strongly associated with the target, but also as distributional information. Indeed, it is by virtue of their distributional properties that traits get into the value component in the first place. They get in because the distributions are unimodal but highly skewed, or bimodal or multimodal; none of them is normal (even in the case of a trait inconsistent behavior such as Smith's surprisingly stupid act, there is the

beginning of a bimodal distribution). This is why I think it helpful to view traits as continuous rather than as categorical constructs. Furthermore, their distributional properties will surely play a role in the construction of social judgments (see also Kahneman & Miller, 1986).² What drives the system is deviations from normalcy. In all cases, however, I think it reasonable to suppose that some points on the distribution may be related to specific retrievable entries in the fact component.

The idea that some of the information we have encoded about individuals is fact-based and some trait-based is not particularly new, even though the constraints that I have proposed for trait representations may be. But it is the third aspect of person representations that is the most distinctive feature of the model I am proposing. This aspect is what I call the *affective experience* component. It is curious that social psychologists who deal with attitudes and impression formation have devoted so little attention to the role of affective experience induced by the target (although see Berscheid & Walster, 1978, and Clore, 1975). Perhaps the recent neglect is a consequence of the typical research paradigm that I criticized earlier—fictitious people of no personal significance tend not to induce affective experiences in us. In any event, many of the things with respect to which we have attitudes are things that directly or indirectly give rise to emotional reactions in us. People talk to us, or do things to us or for us, or we observe or otherwise come to know about what they do or did or what happened to them. In some cases, the emotions we have in response to such events are immediate feelings of (momentary) liking or disliking, emotions that Ortony, et al. (1988) call attraction emotions. In other cases, attraction emotions arise indirectly, as a result of an initial experience of a more specific emotion. So, for example, we often find ourselves (momentarily) disliking somebody who is making us angry. In either case, what is important in the current context is the idea that we end up with relatively undifferentiated emotions of momentary liking or disliking. What I am proposing is that the affective experience component contains a summary record of the liking and disliking experiences that the target person has induced in us, again in the form of something like a frequency distribution of positive and

²There is a complex issue lurking here. Imagine you meet somebody who engages in the courageous act of climbing up a dangerous cliff in order to save somebody's life. If, in Kahneman and Miller's (1986) terms, one recruits ordinary people as the norm or contrast set, then courage would certainly be diagnostic for this individual and thus, according to the model, the relevant dimension ought to be incorporated into the value space. But now suppose that you see someone else doing exactly the same thing but you know that he is a stunt man in a movie that is being shot. The same observed behavior evaluated with respect to stunt men could seem quite normal. There are several ways of dealing with this. One, which is not very parsimonious, but which might still be right, is multiple representations—one for each major role the person fills, with the context activating the appropriate one. A simpler solution would be to use an exception marker. This would involve thinking of the dimensions in the value space as being tagged with the names of any category with respect to which the dimension was known *not* to be diagnostic. In this scheme, we would have an extreme value on the courageous-fearful dimension with a tag such as "not as stunt man."

negative feelings of different intensity. This is not a record of the particular emotions induced, but only a record of the magnitude and sign of the undifferentiated affect (see Diener, Larsen, Levine, & Emmons, 1985 for the relevance of frequency of undifferentiated positive and negative affect for subjective well-being). I am assuming that the event itself along with the facts pertaining to its particular emotional impact are represented in and, in principle, recoverable from the fact component.

When we come to ask what is "dispositional" liking, we might imagine that it is primarily determined by a temporally weighted frequency function from the affective experience component, perhaps modulated by some sort of summary assessment of currently salient information in the value component. The value component associated with the representation of a person is influenced by changes in both the fact and the affective experience components. In other words, the attributions that I make with respect to individuals are influenced both by things that I (come to) believe to be true of them, and by the quality of my emotional experiences with respect to them. For example, if I learn that Smith, whom I was previously disposed to like, was once convicted and imprisoned for child abuse, I might introduce some (even many) dimensions into the value component that were previously not explicitly represented—I might now think him very dishonest, untrustworthy, insincere, and cruel whereas previously I just viewed him as normal along these dimensions.

The interesting and more difficult case is the one in which there already is an explicitly represented dimension, let us say one with a positively skewed distribution of trustworthy behaviors, and I then observe a highly inconsistent behavior. In such a situation, it seems reasonable to suppose that the behavior is located in the appropriate place on the dimension in the value component. But how would we respond if asked how trustworthy we thought the target was? Different answers could result from different processes. One answer might be "I used to think he was very trustworthy, but now I'm beginning to have doubts." Another might be "Moderately," and yet another "Very." These are interestingly different. The first is a qualitative response perhaps based on the realization that a bimodal distribution might be emerging. It reflects uncertainty as to how to construct a summary value. The second commits to a summary value, but gives great weight to the inconsistent information.³ The third possibility would be one that essentially discounted the inconsistent behavior as an aberration having no noticeable effect on the central tendency. There is no a priori reason to believe that different individuals treat such cases in the same way, or that the same individual resolves different cases in the same way. This is an empirical question.

³There is evidence that attitudes are strongly governed by primacy (e.g., Anderson, 1965; Asch, 1946; Dreben, Fiske, & Hastie, 1979). I do not know the extent to which this effect interacts with familiarity with the target. It might be, for example, that the effect is strong for relatively unfamiliar targets, but that recency plays a larger role with familiar targets. This is a complex issue beyond the scope of the present chapter.

GLOBAL JUDGMENTS AND PERSON REPRESENTATIONS

In the last section I outlined a tripartite model of person representation. I want now to consider how the kind of representations that I have suggested might be used in making global social judgments and how they change over time.

I shall start by discussing the question of global judgments of liking, because these judgments are presumably based on what George calls "simple preferences." So, imagine you have a colleague, Fred, and you are asked "How much do you like Fred?" The simplest way in which the representation you have of Fred can provide a basis for a response is through the affective experience component. Recall that this component encodes a distribution of the frequency of undifferentiated positive and negative feelings of different intensity. If the current shape of the distribution reflects the fact that Fred has induced mainly negative and moderately intense emotions in you, you could simply retrieve some measure of central tendency and transform it into a response indicating moderate (or considerable, or whatever) dislike. At the same time, however the central tendency is determined, it would seem desirable that it be capable of being sensitive to highly salient recent information.⁴ If Fred's most recent actual or reported behaviors have been drastically out of character and have generated positive feelings in you, one might sometimes expect the judgment to reveal a dispositional liking that is more positive than it used to be.

What I am proposing is that the primary source of information in a person representation for making liking judgments is the affective experience component. In a nutshell, what I am suggesting is that dispositional liking is primarily determined by the history of momentary liking and disliking, and that momentary liking and disliking (themselves emotions) can have their roots in other emotions. The bottom line here is that the primary source of dispositional liking (i.e., simple preferences) is emotion.

Now, even though the primary determinant of how much one likes or dislikes (i.e., one's attitude toward) somebody is the history of affective experience induced by that person, it seems reasonable to suppose that in some evaluative judgments the contribution of the affective experience component is modulated or even eliminated. Of particular interest here is the possibility that it plays a much less central role in judgments of likeableness. These judgments might well be based primarily on information in the value component, by somehow integrating diagnostic trait information. The contribution of information from the affective

⁴A simple way in which recency effects could be modeled would be for the current summary value to be computed by taking an increment away from the previous summary value in the direction of the most recent value. The size of this increment would affect the sensitivity to recent information. I am indebted to Bill Revelle for this observation.

tive experience component might then be modulatory rather than constitutive. It could be eliminated if one made the liking judgment while in the midst of an emotional state, possibly even an emotional state not induced by the target.

I now want to summarize what I think might be the different roles of and interconnections between the three components in a person representation. First, the most important claim is that the affective experience component is the primary source of attitudes. Because the affective experience component is always updated in response to the emotions induced by the target, attitudes are subject to change, sometimes, as in rapidly evolving close relationships, quite quickly and quite dramatically. Second, the value component is the primary source of likeableness judgments, and the primary source of predictions and explanations about the (especially interpersonal) behavior of the target. The contents of this component are subject to change in two main ways. One is through the introduction of new diagnostic traits inferred from newly encountered extreme behaviors (coded in the fact component) that are out of character with respect to the existing trait structure in the value component. The second way in which the value component can change is generally more gradual and involves a change in shape of the distribution on a trait dimension.

My third main point is that observations and beliefs about the target are coded in the fact component and are sometimes relatable to particular points in distributions in both the affective experience and the value components. And finally, the affective experience component can be an information source for the fact component. For example, I can represent the fact that the target made me angry yesterday, or that the target often makes one angry, and this can lead to an evaluative inference (e.g., the target is exceptionally aggressive, annoying, uncooperative, etc.) that is manifested by a change in the value component.

The claim that liking judgments are normally determined by the affective component rather than the value component seems intuitively plausible. It amounts to saying that you like someone, say Mary, because most of your encounters with her have made you feel good rather than because you find her intelligent, attractive, and trustworthy, even though she is not very punctual and not very generous. Furthermore, by eschewing reliance on traits as the basis for liking judgments, we can explain the difficulty of making liking judgments about unknown others. Suppose I tell you that I have a female friend who is novelist. I tell you that she is uncommonly beautiful, intelligent, sensitive, and warm. On the basis of this information (and more like it) you could obviously construct both a fact component and a value component. However, there would be nothing in the affective experience component. In such a case, there is something odd about my asking you how much you like her. Your response would be "She sounds terrific, but I don't know whether I like her. I haven't met her." (Indeed, it is almost a joke to say "I like her already"—before even meeting her.) So, with respect to dispositional liking, we generally have no opinion if we have no

corresponding affective experience. On the other hand, the value component is often the source of *predictions*. You might confidently expect that you would like her when you met her.

One way of viewing the difference between the affective experience component and the value component is to say that the former contains first order and the latter contains second order affective information. The second order information is very important because it allows the incorporation of affective expectations by inference in cases where we have had no direct experience (as well as in cases where we have). Few of us have actually had an encounter with an armed burglar, yet we all know that we do not like the idea. When we think we are encountering a particular armed burglar, our dislike of the idea is presumably based on inferences from a stereotype that lead us to expect the individual to be selfish, unkind, ruthless, dangerous, and so on. We infer (perhaps from expected behaviors, perhaps directly) extreme values on dimensions corresponding to such characteristics, and consequently, we expect to experience negative emotions in an encounter. In the example of the nighttime intruder that I gave earlier, all we need is for the generic ARMED-BURGLAR schema to contain this kind of affective information. From the theoretical perspective, this amounts to saying that we have a value component for social categories as well as for individuals, and in fact, there seems to be no reason why categories should not have something like an affective experience component too.

... There are some empirical data that can be interpreted as providing at least mild support for aspects of the model I am proposing. Abelson, Kinder, Peters, and Fiske (1982) reported the results of some analyses of data from national surveys relating to political person perception. In particular, they looked at the relationship between two kinds of independent variables on summary evaluations of and preferences for well-known political figures. Their independent variables were, first, affective experiences induced by the politicians and, second, trait ratings. In other words, they examined the relationship between aspects of what we would call the affective experience component and (no doubt, at least to some extent) the value component on (summary evaluative) liking judgments.⁵ The results were clear. Judgments based on self reports about whether targets had induced emotions (e.g., fear, anger, pride) in subjects were better predictors of summary evaluative judgments than were the ratings on a set of traits (e.g., dishonest, weak, knowledgeable). Furthermore, there appeared to be no differential effect for individual emotions over and above the fact that they induced positive or negative feelings. All of this is consistent with the idea that emotions induced by a target leave relatively simple affective traces that are the primary source of summary evaluations of liking or disliking. In other words, the study

⁵The reason I hedge a little here on whether this really constitutes pitting the affective experience space against the value space is that for the traits examined to appear in the value space, it would be necessary, on my account, for them to have the right kind of distributional properties. In fact there is no reason to believe that this was generally true.

not only suggests a separable representational difference between what Abelson et al. (1982) called affective and semantic information, but subjects were also making their evaluative judgments by giving more weight to the affective information, that is to the fact that targets had induced emotions (derived from the affective experience component) than to semantic information, that is, their trait representations (derived at least in part from the value component).

Finally, I want to consider briefly the possibility that values in a broader sense than simple preferences might arise from emotions. How, for example, might one explain the fact that most of us value helpfulness so that helpful-unhelpful becomes a dimension that can potentially play a role in the value component? With some caveats having to do with social and cultural sanctions, I think we can answer this by supposing that the affective experience component is developmentally prior to the value component. In other words, children first learn to value something (including people, objects, and types of behavior) because they like it, and they like it because it makes them feel good. Consider a child who notices that every time that he needs help and gets it, he has a positive feeling toward the helper and a positive feeling vis-à-vis his or her improved situation. Let us further suppose that the child realizes that when he or she gives help to someone else, the other behaves in a way that is also suggestive of positive feelings. These positive feelings, perhaps along with those that arise from the general social sanctioning the child receives and sees for helping behavior, might well lead the child to believe that there is something intrinsically good about helping. Presumably, in some such way, the child comes to value the attribute of helpfulness; that is, helpfulness comes to be viewed as a value-laden attribute rather than as a value-free one.

CONCLUSION

I have tried to show that emotions are indeed a source of value. Primarily, I have focused on this in the sense of values as simple preferences, trying to show that it is the affective residue of distinct emotions induced by someone that constitutes the primary source of dispositional liking for that person. I think the same sort of model could be used to explain liking outside the interpersonal domain. I am sure that my proposals are inadequate in many ways. There are many complex factors at work. Nevertheless, I think that there is a fertile research ground here. And if I am right about this, and if we start to cultivate it, no one will be happier than George Mandler.

ACKNOWLEDGMENTS

Preparation of this article was supported in part by grants from the National Science Foundation, BNS 8318077 and BNS 8721853, and in part by Andersen

Consulting through Northwestern University's Institute for the Learning Sciences. I am grateful to Gerald Clore for helpful comments on a draft of this chapter.

REFERENCES

- Abelson, R. P., Kinder, D. R., Peters, M. D., & Fiske, S. T. (1982). Affective and semantic components in political person memory. *Journal of Personality and Social Psychology*, *42*, 619-630.
- Anderson, N. H. (1965). Primacy effects in impression formation using a generalized order effect paradigm. *Journal of Personality and Social Psychology*, *2*, 1-9.
- Anderson, N. H. (1971). Integration theory and attitude change. *Psychological Review*, *78*, 171-206.
- Anderson, N. H. (1981). *Foundations of information integration theory*. New York: Academic Press.
- Asch, S. E. (1946). Forming impressions of personality. *Journal of Abnormal and Social Psychology*, *41*, 258-290.
- Berscheid, E. (1982). Attraction and emotion in interpersonal relations. In M. S. Clark & S. T. Fiske (Eds.), *Affect and cognition* (pp. 37-54). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Berscheid, E., & Walster, E. H. (1978). *Interpersonal attraction* (2nd ed.). Reading, MA: Addison-Wesley.
- Clore, G. L. (1975). *Interpersonal attraction: An overview*. Morristown, NJ: General Learning Press.
- Deiner, E., Larsen, R. J., Levine, S., & Emmons, R. A. (1985). Intensity and frequency: Dimensions underlying positive and negative affect. *Journal of Personality and Social Psychology*, *53*, 767-774.
- Dreben, E. K., Fiske, S. T., & Hastie, R. (1979). The independence of item and evaluation information: Impression and recall order effects in behavior-based impression formation. *Journal of Personality and Social Psychology*, *37*, 1758-1768.
- Fishbein, M., & Ajzen, I. (1975). *Belief, attitude, intention and behavior: An introduction to theory and research*. Reading, MA: Addison-Wesley.
- Fiske, S. T. (1982). Schema-triggered affect: Applications to social perception. In M. S. Clark & S. T. Fiske (Eds.), *Affect and cognition* (pp. 55-78). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Fiske, S. T., & Neuberg, S. L. (1990). A continuum of impression formation, from category-based to individuating processes: Influences of information and motivation on attention and interpretation. In M. P. Zanna (Ed.), *Advances in experimental social psychology* (Vol. 23) (pp. 1-74). New York: Academic Press.
- Fiske, S. T., & Pavelchak, M. A. (1986). Category-based versus piecemeal-based affective responses: Developments in schema-triggered affect. In R. M. Sorrentino & E. T. Higgins (Eds.), *Handbook of motivation and cognition: Foundations of social behavior* (pp. 167-203). New York: Guilford.
- Kahneman, D., & Miller, D. T. (1986). Norm theory: Comparing reality to its alternatives. *Psychological Review*, *93*, 136-153.
- Mandler, G. (1975). *Mind and emotion*. New York: Wiley.
- Mandler, G. (1982). The structure of value. In M. S. Clark & S. T. Fiske (Eds.), *Affect and cognition* (pp. 3-36). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Mandler, G. (1989). *Approaches to a psychology of value*. Unpublished manuscript.

- Ortony, A., Clore, G. L., & Collins, A. (1988). *The cognitive structure of emotions*. New York: Cambridge University Press.
- Srull, T. K., & Wyer, R. S. (1989). Person memory and judgment. *Psychological Review*, 96, 58-83.
- Wyer, R. S., & Srull, T. K. (1986). Human cognition in its social context. *Psychological Review*, 93, 322-359.