Implications for Conceptual Structure

Associationists and rule-based theorists tend to have different views concerning the determinants and extent of conceptual coherence. Associationists tend to believe that our beliefs are usually consistent with each other because they reflect the world and the world is necessarily coherent, for it must obey the laws of nature. People may have contradictory beliefs because different aspects of our experience may provide evidence for opposing views. Our experience in the home may suggest that people tend to be generous but our experience on the highway may suggest that people tend to be selfish. On this view, coherence is a property of concepts by virtue and to the extent that experience in the world is coherent.

Rule-based theorists tend to believe that we possess a more potent urge for coherence. Rules can reflect structure in the world as well as conform to their own syntax and semantics, which may impose further structure. Any formal calculus of belief embodies assumptions about which beliefs are consistent with each other. Thus, rules enforce their own principles of coherence and, accordingly, rule-based theorists tend to believe that people try to conform. Some of them (e.g., Keil, 1989; Murphy, 1993) imply that people try to construct a global rule-based theory that causes them to try to be globally coherent in their everyday lives (and not just when doing philosophy or science).

Allowing humans to be both associationists and rule-governed suggests a way to reconcile these views. People may have an urge for coherence, but that urge is for *local* coherence. Rules are applied in such a way that current explanations, the temporary contents of working memory, are internally consistent and consistent with the long-term knowledge deemed relevant. For the most part, people can rely on the world to maintain coherence across situations (unless our perceptions are terribly distorted). Because they reflect objects and events in the world fairly directly, the associative system can do some of that work.

CONCLUSIONS

People are renowned for their willingness to behave in ways that they cannot justify, let alone explain. Instead of performing a complete analysis of their interests, people vote for a politician because they have always for voted for that person or buy an item because it is associated with an image that they would like to project. However, most people only go so far. They would not do something that they consider irrational if it entailed a real penalty or cost. Fewer people buy an item after it has been linked to cancer. So, on one hand, people "follow their noses" by allowing associations to guide them; on the other hand, they are compelled to behave in a manner more justifiable. The fact that people are pulled in two directions at once suggests two forces pulling.

23. The Affect Heuristic

Paul Slovic, Melissa Finucane, Ellen Peters, and Donald G. MacGregor

This chapter introduces a theoretical framework that describes the importance of affect in guiding judgments and decisions. As used here, *affect* means the specific quality of "goodness" or "badness" (1) experienced as a feeling state (with or without consciousness) and (2) demarcating a positive or negative quality of a stimulus. *Affective responses* occur rapidly and automatically – note how quickly you sense the feelings associated with the stimulus words *treasure* or *hate*. We argue that reliance on such feelings can be characterized as the *affect heuristic*. In this chapter, we trace the development of the affect heuristic across a variety of research paths followed by ourselves and many others. We also discuss some of the important practical implications resulting from ways that this heuristic impacts our daily lives.

BACKGROUND

Although affect has long played a key role in many behavioral theories, it has rarely been recognized as an important component of human judgment and decision making. Perhaps befitting its rationalistic origins, the main focus of descriptive decision research has been cognitive, rather than affective. When principles of utility maximization appeared to be descriptively inadequate, Simon (1956) oriented the field toward problem-solving and information-processing models based on bounded rationality. The work of Tversky and Kahneman (1974) and Kahneman, Slovic, and Tversky (1982) demonstrated how boundedly rational individuals use such heuristics as availability, representativeness, and anchoring and adjustment to make judgments, and how they use simplified strategies such as "elimination by aspects" to make choices (Tversky, 1972). Other investigators elaborated the cognitive strategies underlying judgment and choice through models of constructed preferences (Slovic, 1995; Payne, Bettman, & Johnson, 1993), dominance structuring (Montgomery, 1983), and comparative advantages (Shafir, Osherson, & Smith, 1989). In 1993, the entire volume of the journal *Cognition* was dedicated to the topic Reason-Based Choice, in which it was argued that "Decisions . . . are often reached by focusing

Financial support for the writing of this chapter was provided by the National Science Foundation under grant SES 9876587.

on reasons that justify the selection of one option over another" (Shafir, Simonson, & Tversky, 1993, p. 34). Similarly, a state-of-the-art review by Busemeyer, Hastie, and Medin (1995) was titled "Decision Making from a Cognitive Perspective." In keeping with its title, it contained almost no references to the influence of affect on decisions.

Despite this cognitive emphasis, the importance of affect is being recognized increasingly by decision researchers. A limited role for affect was acknowledged by Shafir et al. (1993), who conceded, "People's choices may *occasionally* stem from affective judgments that preclude a thorough evaluation of the options" (p. 32, emphasis added).

A strong early proponent of the importance of affect in decision making was Zajonc (1980), who argued that affective reactions to stimuli are often the very first reactions, occurring automatically and subsequently guiding information processing and judgment. According to Zajonc, all perceptions contain some affect. "We do not just see 'a house': We see a *handsome* house, an *ugly* house, or a *pretentious* house" (p. 154). He later adds,

> We sometimes delude ourselves that we proceed in a rational manner and weight all the pros and cons of the various alternatives. But this is probably seldom the actual case. Quite often "I decided in favor of X" is no more than "I liked X".... We buy the cars we "like," choose the jobs and houses we find "attractive," and then justify these choices by various reasons. (p. 155)

Affect also plays a central role in what have come to be known as *dual-process theories* of thinking, knowing, and information processing (Chaiken and Trope, 1999; Sloman, 1996). As Epstein (1994) observed,

> There is no dearth of evidence in every day life that people apprehend reality in two fundamentally different ways, one variously labeled intuitive, automatic, natural, non-verbal, narrative, and experiential, and the other analytical, deliberative, verbal, and rational. (p. 710)

One of the characteristics of the experiential system is its affective basis. Although analysis is certainly important in some decision-making circumstances, reliance on affect and emotion is a quicker, easier, and more efficient way to navigate in a complex, uncertain, and sometimes dangerous world. Many theorists have given affect a direct and primary role in motivating behavior. Epstein's (1994) view on this is as follows:

> The experiential system is assumed to be intimately associated with the experience of affect,... which refer[s] to subtle feelings of which people are often unaware. When a person responds to an emotionally significant event,... the experiential system automatically searches its memory banks for related events, including their emotional accompaniments.... If the activated feelings are pleasant, they motivate actions and thoughts anticipated to reproduce the feelings. If the feelings are unpleasant, they motivate actions and thoughts anticipated to avoid the feelings. (p. 716)

Also emphasizing the motivational role of affect, Mowrer (1960a, 1960b) conceptualized conditioned emotional responses to images as prospective gains

and losses that directly "guide and control performance in a generally sensible adaptive manner" (1960a, p. 30). He criticized theorists who postulate purely cognitive variables such as expectancies (probabilities) intervening between stimulus and response, cautioning that we must be careful not to leave the organism at the choice point "lost in thought." Mowrer's solution was to view expectancies more dynamically (as conditioned emotions such as hopes and fears) serving as motivating states leading to action.

One of the most comprehensive and dramatic theoretical accounts of the role of affect in decision making is presented by neurologist Antonio Damasio (1994), in his book *Descartes' Error: Emotion, Reason, and the Human Brain*. Damasio's theory is derived from observations of patients with damage to the ventromedial frontal cortices of the brain that has left their basic intelligence, memory, and capacity for logical thought intact but has impaired their ability to feel – that is, to associate affective feelings and emotions with the anticipated consequences of their actions. Close observation of these patients combined with a number of experimental studies led Damasio to argue that this type of brain damage induces a form of sociopathy (Damasio, Tranel, & Damasio, 1990) that destroys the individual's ability to make rational decisions; that is, decisions that are in his or her best interests. Persons suffering this damage became socially dysfunctional even though they remain intellectually capable of analytical reasoning. Commenting on one particularly significant case, Damasio observes:

> The instruments usually considered necessary and sufficient for rational behavior were intact in him. He had the requisite knowledge, attention, and memory; his language was flawless; he could perform calculations; he could tackle the logic of an abstract problem. There was only one significant accompaniment to his decision-making failure: a marked alteration of the ability to experience feelings. Flawed reason and impaired feelings stood out together as the consequences of a specific brain lesion, and this correlation suggested to me that feeling was an integral component of the machinery of reason. (p. XII)

In seeking to determine "what in the brain allows humans to behave rationally," Damasio argues that thought is made largely from images, broadly construed to include sounds, smells, real or imagined visual impressions, ideas, and words. A lifetime of learning leads these images to become "marked" by positive and negative feelings linked directly or indirectly to somatic or bodily states (Mowrer and other learning theorists would call this *conditioning*): "In short, *somatic markers are ... feelings generated from secondary emotions*. These emotions and feelings *have been connected, by learning, to predicted future outcomes of certain scenarios*" (Damasio, 1994, p. 174). When a negative somatic marker is linked to an image of a future outcome, it sounds an alarm. When a positive marker is associated with the outcome image, it becomes a beacon of incentive. Damasio concludes that somatic markers increase the accuracy and efficiency of the decision process, and their absence degrades decision performance.

Damasio tested the somatic marker hypothesis in a decision-making experiment in which subjects gambled by selecting cards from any of four decks. Turning each card resulted in the gain or loss of a sum of money, as revealed on the back of the card when it was turned. Whereas normal subjects and patients with brain lesions outside the prefrontal sectors learned to avoid decks with attractive large payoffs but occasional catastrophic losses, patients with frontal lobe damage did not, thus losing a great deal of money. Although these patients responded normally to gains and losses when they occurred (as indicated by skin conductance responses immediately after an outcome was experienced), they did not seem to learn to anticipate future outcomes (e.g., they did not produce normal skin conductance responses when contemplating a future choice from a dangerous deck). In other words, they failed to show any proper anticipatory responses, even after numerous opportunities to learn them.

Despite the increasing popularity of affect in research programs and recent attempts to acknowledge the importance of the interplay between affect and cognition, further work is needed to specify the role of affect in judgment and decision making. The ideas articulated here are intended as a step toward encouraging the development of theory about affect and decision making and demonstrating how such a theory can be tested.

The basic tenet of this chapter is that images, marked by positive and negative affective feelings, guide judgment and decision making. Specifically, it is proposed that people use an *affect heuristic* to make judgments; that is, representations of objects and events in people's minds are tagged to varying degrees with affect. In the process of making a judgment or decision, people consult or refer to an "affect pool" containing all the positive and negative tags consciously or unconsciously associated with the representations. Just as imaginability, memorability, and similarity serve as cues for probability judgments (e.g., the availability and representativeness heuristics), affect may serve as a cue for many important judgments. Using an overall, readily available affective impression can be far easier – more efficient – than weighing the pros and cons or retrieving from memory many relevant examples, especially when the required judgment or decision is complex or mental resources are limited. This characterization of a mental short-cut leads to labeling the use of affect a "heuristic."

EMPIRICAL EVIDENCE: MANIPULATING PREFERENCES THROUGH CONTROLLED EXPOSURES

The fundamental nature and importance of affect has been demonstrated repeatedly in a remarkable series of studies by Zajonc and his colleagues (see, e.g., Zajonc, 1968; Zajonc & Markus, 1982). The concept of stimulus exposure is central to all of these studies. The central finding is that, when objects are presented to an individual repeatedly, the "mere exposure" is capable of creating a positive attitude or preference for these objects.

In the typical study, stimuli such as nonsense phrases, or faces, or Chinese ideographs are presented to an individual with varying frequencies. In a later session, the individual judges these stimuli on liking, or familiarity, or both. The more frequent the prior exposure to a stimulus, the more positive the response. A meta-analysis by Bornstein (1989) of mere exposure research published between 1968 and 1987 included more than 200 experiments examining the exposure–affect relationship. Unreinforced exposures were found to reliably enhance affect toward visual, auditory, gustatory, abstract, and social stimuli.

Winkielman, Zajonc, and Schwarz (1997) demonstrated the speed with which affect can influence judgments in studies using a subliminal priming paradigm. Participants were "primed" through exposure to a smiling face, a frowning face, or a neutral polygon presented for 1/250th second, an interval so brief that there is no recognition or recall of the stimulus. Immediately following this exposure, an ideograph was presented for 2 seconds, following which the participant rated the ideograph on a scale of liking. Mean liking ratings were significantly higher for ideographs preceded by smiling faces. This effect was lasting. In a second session, ideographs were primed by the other face: the face not associated with the stimulus in the first session. This second priming was ineffective because the effect of the first priming remained.

It is not just subliminal smiles that affect our judgment. La France and Hect (1995) found that students accused of academic misconduct who were pictured as smiling received less punishment than nonsmiling transgressors. Smiling persons were judged as more trustworthy, good, honest, genuine, obedient, blameless, sincere, and admirable than nonsmiling targets.

The perseverance of induced preferences was tested by Sherman, Kim, and Zajonc (1998), who asked participants to study Chinese characters and their English meanings. Half of the meanings were positive (e.g., beauty), half were negative (e.g., disease). Then participants were given a test of these meanings followed by a task in which they were given pairs of characters and were asked to choose the one they preferred. Participants preferred characters with positive meaning 70% of the time. Next, the characters were presented with neutral meanings (desk, linen) and subjects were told that these were the "true" meanings. The testing procedure was repeated and, despite learning the new meanings, the preferences remained the same. Characters that had been paired initially with positive meanings still tended to be preferred.

These various studies demonstrate that affect is a strong conditioner of preference, regardless of whether the cause of that affect is consciously perceived. They also demonstrate the independence of affect from cognition, indicating that there may be conditions of affective or emotional arousal that do not necessarily require cognitive appraisal. This affective mode of response, unburdened by cognition and therefore much faster, has considerable adaptive value.

EVALUATING GAMBLES

The affect heuristic can explain a finding that has intrigued and perplexed the first author since he first observed it in 1984. Slovic and Amos Tversky were reexamining the early studies by Slovic and Lichtenstein (1968) and Lichtenstein and Slovic (1971, 1973) that pointed to compatibility between stimulus attributes and response scales as an explanation for preference reversals. Such reversals were exhibited when an individual chose Gamble A (with a high probability of winning a modest amount of money) over Gamble B (with a smaller probability of a larger payoff) but assigned a larger monetary value (buying price or selling price) to Gamble B. Presumably, the reversal occurred because the gamble payoffs were given more weight in the pricing response mode than in choice, due to the compatibility between prices and payoffs, both of which were measured in dollars.

Tversky and Slovic decided to replicate the earlier reversal studies with three changes:

1. The complexity of the gamble was minimized by eliminating losses. Each gamble consisted merely of a stated probability of winning a given amount. There was no possible loss of money.
2. Following Goldstein (1984) and Goldstein and Einhorn (1987), who observed reversals with ratings and prices, we included ratings of a gamble's attractiveness along with choices and pricing as methods of eliciting preferences. The attractiveness scale ranged between 0 (not at all attractive) and 20 (very attractive).
3. To ensure the strategic equivalence of our three elicitation procedures, we devised a method for linking preferences to outcomes that was identical across all conditions. Subjects were told that a pair of bets would be selected and the bet that received the higher attractiveness rating (or the higher price, or that was preferred in the choice task) would be the bet they would play. Consequently, the preferences elicited by prices and ratings should not differ from each other or from the preferences elicited by direct choices. Some of the gambles were, in fact, actually played.

Using this design, we observed strong differences between response modes, leading to many preference reversals. Particularly striking was the difference between ratings and prices. Ratings produced an overwhelming dominance of high probability bets over high payoff bets (the bet with higher probability of winning had the higher attractiveness rating 80% to 90% of the time, but was assigned a higher price only 10% to 15% of the time). The mean evaluations of the two bets shown in Table 23.1 were typical.

Seeking to explain the results shown in Table 23.1 in terms of compatibility, we linked the compatibility effect to the ease of mapping the stimulus component of a gamble onto the response scale. The easier it is to execute such a mapping, the greater the weight given the component. In principle, a gamble's payoff is more compatible with a price response than with a rating because prices and payoffs are both expressed in dollars. Therefore, payoffs should get greater weight in pricing

Table 23.1. Mean Evaluations of Two Bets

Bet	Mean Price	Mean Rating (0–20 Scale)
29/36 to win \$2	\$1.25	13.2
7/36 to win \$9	\$2.11	7.5

than in rating. The extremely high weight given probabilities when rating attractiveness may be explained by the fact that the probabilities are more readily coded as attractive or unattractive than are the payoffs. For example, 29 out of 36 chances to win are very attractive odds. However, a \$9 payoff may be harder to map on a rating scale because its attractiveness depends on what other payoffs are available.

According to this explanation, if we could make a gamble's payoff more compatible with the attractiveness rating, we would presumably enhance the weight given to payoff in the rating response mode. We attempted to do this in a new experiment, focusing on the gamble 7/36 to win \$9. To make the payoff more compatible with regard to the scale of attractiveness, we added a very small loss (5¢) to the gamble:

7/36 win \$9
29/36 lose 5¢

Whereas the attractiveness of \$9 might not be readily apparent, we reasoned that a bet offering \$9 to win and only 5¢ to lose should appear to have a very attractive payoff ratio. This led us to predict that one might increase the attractiveness of a gamble (p to win X) by adding a loss component to it.

The results exceeded our expectations. The gamble with no loss had the lower attractiveness rating (mean = 9.4 on the 0–20 scale). Adding a 5¢ loss led to a much higher attractiveness rating (mean = 14.9). Even the bet

7/36 win \$9
29/36 lose 25¢

was judged more attractive (mean = 11.7) than the bet with no loss.

Would adding a small loss to the gamble enhance its attractiveness in choice as it did in rating? We addressed this question by asking 96 University of Oregon students to choose between playing a gamble and receiving a gain of \$2. For half of the students, the gamble was 7/36 win \$9; for the others, the gamble had the 5¢ loss. Whereas only 33.3% chose the \$9 gamble over the \$2, 60.8% chose the (\$9, −5¢) gamble over the \$2. A replication study with \$4 as the alternative to the gamble produced similar results. The enhancement produced by adding a small loss thus holds for choices as well as for rating responses.

The enhanced attractiveness produced by small losses was originally predicted and explained in terms of compatibility, and we now see it also as an example of the affect heuristic. This broader perspective was induced, in part, by results obtained later by Mellers, Richards, and Birnbaum (1992), Hsee (1995, 1996a, 1996b, 1998), and our own subsequent studies of imagery, affect, and decision making. These convergent streams of research are described in the following sections.

IMAGE, AFFECT, AND DECISION MAKING

The early anomalous findings with gambles were laid aside while other means of explaining the differences between ratings, choices, and pricing responses were developed (see Tversky, Slovic, & Kahneman, 1990). At the same time, Slovic and colleagues at Decision Research embarked on a research program designed to test whether introducing a hazardous facility into a region might stigmatize that region and cause people to avoid going there to recreate, retire, or do business (Slovic et al., 1991). Believing self-report to be unreliable ("If they build it, will you not come?"), research on stigmatization was conducted through a number of empirical studies designed to examine the relationship between imagery, affect, and decision making. After conducting these studies, we learned that they fit closely with a large body of existing theory and research, such as the work of Damasio, Mowrer, and Epstein, described earlier.

Several empirical studies have demonstrated a strong relationship between imagery, affect, and decision making. Many of these studies used a word-association technique, which involves presenting subjects with a target stimulus, usually a word or very brief phrase and asking them to provide the first thought or image that comes to mind. The process is then repeated a number of times (e.g., three to six), or until no further associations are generated. Following the elicitation of images, subjects are asked to rate each image they give on a scale ranging from very positive (e.g., +2) to very negative (e.g., −2), with a neutral point in the center. Scoring is done by summing or averaging the ratings to obtain an overall index.

This imagery method has been used successfully to measure the affective meanings that influence people's preferences for different cities and states (Slovic et al., 1991) as well as their support or opposition to technologies such as nuclear power (Peters & Slovic, 1996).

Table 23.2 illustrates the method in a task where one respondent was asked to give associations to each of two cities and, later, to rate each image affectively. The cities in this example show a clear affective preference for San Diego over Denver. Slovic et al. (1991) showed that summed image scores such as these were highly predictive of expressed preferences for living in or visiting cities. In one study, they found that the image score predicted the location of *actual* vacations during the next 18 months.

Subsequent studies have found affect-laden imagery elicited by word associations to be predictive of preferences for investing in new companies on the stock market (MacGregor, Slovic, Dreman, & Berry, 2000) and predictive of adolescents' decisions to take part in health-threatening and health-enhancing behaviors such as smoking and exercise, respectively (Benthin et al., 1995).

EVALUABILITY

The research with images points to the importance of affective impressions in judgments and decisions. However, the impressions themselves may vary

Table 23.2. Images, Ratings, and Summation Scores for One Respondent

Stimulus	Image Number	Image	Image Rating
San Diego	1	Very nice	2
San Diego	2	Good beaches	2
San Diego	3	Zoo	2
San Diego	4	Busy freeway	1
San Diego	5	Easy to find way	1
San Diego	6	Pretty town	2
			10
Denver	1	High	2
Denver	2	Crowded	0
Denver	3	Cool	2
Denver	4	Pretty	1
Denver	5	Busy airport	−2
Denver	6	Busy streets	−2
			1

Note: Based on these summation scores, this person's predicted preference for a vacation site would be San Diego.
Source: Slovic et al. (1991).

not only in their valence, but also in the precision with which they are held. It turns out that the precision of an affective impression substantially impacts judgments.

We refer to the distributional qualities of affective impressions and responses as *affective mappings*. Consider, for example, some questions posed by Mellers et al. (1992): "How much would you like a potential roommate if all you knew about her was that she was said to be intelligent?" Or, "Suppose, instead, all you knew about her was that she was said to be obnoxious?" Intelligence is a favorable trait but it is not very diagnostic (e.g., meaningful) for likeableness, hence its affective map is rather diffuse. In contrast, obnoxiousness will likely produce a more precise and more negative impression.

How much would you like a roommate said to be both intelligent *and* obnoxious? Anderson (1981) has shown that the integration of multiple pieces of information into an impression of this sort can be described well by a weighted average model where separate weights are given to intelligence and obnoxiousness, respectively. Mellers et al. (1992) further showed that the weights in such integrative tasks are inversely proportional to the variance of the impressions. Thus we would expect the impression produced by the combination of these two traits to be closer to the impression formed by obnoxiousness alone, reflecting greater weight given to obnoxiousness due to its smaller variance (more precise affective mapping). The meaning of a stimulus image appears to be reflected in the precision of the affective feelings associated with that image. More precise affective impressions reflect more precise meanings and carry more weight in impression formation, judgment, and decision making.

Table 23.3. Attributes of Two Dictionaries

	Year of Publication	Number of Entries	Defects
Dictionary A	1993	10,000	No, it's like new
Dictionary B	1993	20,000	Yes, the cover is torn; otherwise it's like new

Source: Adapted from Hsee (1998).

Hsee (1996a, 1996b, 1998) developed the notion of *evaluability* to describe the interplay between the precision of an affective impression and its meaning or importance for judgment and decision making. Evaluability is illustrated by an experiment in which Hsee asked people to assume they were music majors looking for a used music dictionary. In a joint-evaluation condition, participants were shown two dictionaries, A and B (Table 23.3), and asked how much they would be willing to pay for each. Willingness-to-pay was far higher for Dictionary B, presumably because of its greater number of entries. However, when one group of participants evaluated only A and another group evaluated only B, the mean willingness to pay was much higher for Dictionary A. Hsee explains this reversal by means of the *evaluability principle*. He argues that, without a direct comparison, the number of entries is hard to evaluate, because the evaluator does not have a precise notion of how good or how bad 10,000 (or 20,000) entries is. However, the defects attribute is evaluable in the sense that it translates easily into a precise good/bad response and thus it carries more weight in the independent evaluation. Most people find a defective dictionary unattractive and a like-new one attractive. Under joint evaluation, the buyer can see that B is far superior on the more important attribute, number of entries. Thus number of entries becomes *evaluable* through the comparison process.

According to the evaluability principle, the weight of a stimulus attribute in an evaluative judgment or choice is proportional to the ease or precision with which the value of that attribute (or a comparison on the attribute across alternatives) can be mapped into an affective impression. In other words, affect bestows meaning on information (cf., Osgood, Suci, & Tannenbaum, 1957; Mowrer, 1960a, 1960b) and the precision of the affective meaning influences our ability to use information in judgment and decision making. *Evaluability* can thus be seen as an extension of the general relationship between the variance of an impression and its weight in an impression-formation task (Mellers et al., 1992).

Hsee's work in evaluability is noteworthy because it shows that even very important attributes may not be used by a judge or decision maker unless they can be translated precisely into an affective frame of reference. As described in the next section, Hsee finds evaluability effects even with familiar attributes such as the amount of ice cream in a cup (Hsee, 1998). We also demonstrate similar effects with other familiar concepts such as amounts of money or human lives.

PROPORTION DOMINANCE

In situations that involve uncertainty about whether we will win or lose or that involve ambiguity about some quantity of something (i.e., how much is enough), there appears to be one information format that is highly evaluable, leading it to carry great weight in many judgment tasks. This is a representation characterizing an attribute as a proportion or percentage of something, or as a probability. At the suggestion of Chris Hsee (personal communication), we refer to the strong effects of this type of representation as *proportion dominance*.

Proportion (or probability) dominance was evident in the studies of gambles described at the beginning of this chapter. Ratings of a gamble's attractiveness tend to be determined far more strongly by the probabilities of winning and losing than by the monetary payoffs. The curious finding that adding a small loss to a gamble increases its rated attractiveness, explained originally as a compatibility effect, can now be seen to fit well with the notions of affective mapping and evaluability.

According to this view, a probability maps relatively precisely onto the attractiveness scale because probability has a lower and upper bound (0 and 1) and a midpoint below which a probability is "poor" or "bad" (i.e., has worse than an even chance) and above which it is "good" (i.e., has a better than even chance). People know where a given value, such as 7/36, falls within the bounds, and exactly what it means – "I'm probably not going to win." In contrast, the mapping of a dollar outcome (e.g., $9) onto the attractiveness scale is diffuse, reflecting a failure to know how good or bad or how attractive or unattractive $9 is. Thus, the impression formed by the gamble offering $9 to win with no losing payoff is dominated by the relatively precise and unattractive impression produced by the 7/36 probability of winning. However, adding a very small loss to the payoff dimension brings the $9 payoff into focus and thus gives it meaning. The combination of a possible $9 gain and a 5¢ loss is a *very attractive* win/loss ratio, leading to a relatively precise mapping onto the upper end of the scale. Whereas the imprecise mapping of the $9 carries little weight in the averaging process, the more precise and now favorable impression of ($9, −5¢) carries more weight, thus leading to an increase in the overall favorability of the gamble.

The effect of adding a small loss to the gamble can also be explained by norm theory (Kahneman & Miller, 1986; Kahneman & Miller, Chapter 20, this volume). However, a norm-theoretical explanation is consistent with an affective account. It asserts that the gamble with no loss is a relatively mediocre representative of the set of all positive gambles, whereas the gamble with a small loss is a relatively attractive member of the class of mixed (win/loss) gambles.

Proportion dominance surfaces in a powerful way in a very different context, the life-saving interventions studied by Fetherstonhaugh, Slovic, Johnson, and Friedrich (1997), Baron (1997), Jenni and Loewenstein (1997), and Friedrich et al.

(1999). For example, Fetherstonhaugh et al. found that people's willingness to intervene to save a stated number of lives was determined more by the proportion of lives saved than by the actual number of lives that would be saved. However, when two or more interventions were directly compared, number of lives saved become more important than proportion saved. Thus, number of lives saved, standing alone, appears to be poorly evaluable, as was the case for number of entries in Hsee's music dictionaries. With a side-by-side comparison, the number of lives became clearly evaluable and important, as also happened with the number of dictionary entries.

Slovic (unpublished), drawing on proportion dominance and the limited evaluability of numbers of lives, predicted (and found) that people, in a between-groups design, would more strongly support an airport-safety measure expected to save 98% of 150 lives at risk than a measure expected to save 150 lives. Saving 150 lives is diffusely good, and therefore only weakly evaluable, whereas saving 98% of something is clearly very good because it is so close to the upper bound on the percentage scale, and hence is readily evaluable and highly weighted in the support judgment. Subsequent reduction of the percentage of 150 lives that would be saved to 95%, 90%, and 85% led to reduced support for the safety measure but each of these percentage conditions still garnered a higher mean level of support than did the Save 150 Lives Condition (Table 23.4).

Turning to a more mundane form of proportion dominance, Hsee (1998) found that an overfilled ice cream container with 7 oz of ice cream was valued more highly (measured by willingness to pay) than an underfilled container with 8 oz of ice cream (Fig. 23.1). This "less is better effect" reversed itself when the options were juxtaposed and evaluated together. Thus, the proportion of the serving cup that was filled appeared to be more evaluable (in separate judgments) than the absolute amount of ice cream.

Table 23.4. Proportion Dominance and Airport Safety: Saving a Percentage of 150 Lives Receives Higher Support Ratings Than Does Saving 150 Lives.

	Potential Benefit				
	Save 150 Lives	Save 98%	Save 95%	Save 90%	Save 85%
Mean support[a]	10.4	13.6	12.9	11.7	10.9
Median[a]	9.8	14.3	14.1	11.3	10.8
Percentage of ratings ≥ 13	37	75	69	35	31

[a] Cell entries in these rows describe mean and median responses to the question, "How much would you support this proposed measure to purchase the new equipment?" (Critics argue that the money spent on this system could be better spent enhancing other aspects of airport safety.) The response scale ranged from 0 (*would not support at all*) to 20 (*very strong support*). An overall ANOVA resulted in $F_{4,200} = 3.36$, $p = .01$. The Save 98% and Save 95% conditions were both significantly different from the Save 150 Lives condition at $p < .05$, Tukey HSD test.

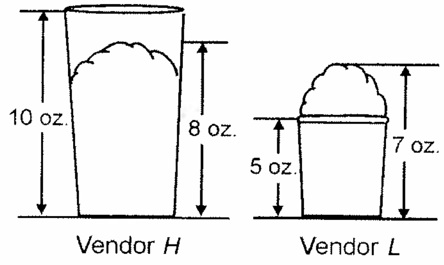

Figure 23.1. Stimuli in ice cream study by Hsee (1998). Participants were given the sizes of the cups and the amounts of ice cream.

INSENSITIVITY TO PROBABILITY

Outcomes are not always affectively as vague as the quantities of money, ice cream, and lives that were dominated by proportion in the above experiments. When consequences carry sharp and strong affective meaning, as is the case with a lottery jackpot or a cancer, the opposite phenomenon occurs – variation in probability often carries too little weight. As Loewenstein, Weber, Hsee, and Welch (2001) observe, one's images and feelings toward winning the lottery are likely to be similar whether the probability of winning is 1 in 10 million or 1 in 10,000. They further note that responses to uncertain situations appear to have an all or none characteristic that is sensitive to the *possibility* rather than the *probability* of strong positive or negative consequences, causing very small probabilities to carry great weight. This they argue, helps explain many paradoxical findings such as the simultaneous prevalence of gambling and the purchasing of insurance. It also explains why societal concerns about hazards such as nuclear power and exposure to extremely small amounts of toxic chemicals fail to recede in response to information about the very small probabilities of the feared consequences from such hazards. Support for these arguments comes from Rottenstreich and Hsee (2001), who show that, if the potential outcome of a gamble is emotionally powerful, its attractiveness or unattractiveness is relatively insensitive to changes in probability as great as from .99 to .01.

MIDCOURSE SUMMARY

We now see that the puzzling finding of increased attractiveness for the gambles to which a loss was appended is part of a larger story that can be summarized as follows:

1. Affect attached to images influences judgments and decisions.
2. The evaluability of a stimulus image is reflected in the precision of the affective feelings associated with that image. More precise affective impressions reflect more precise meanings (i.e., greater evaluability) and carry more weight in impression formation, judgment, and decision making.

3. The anomalous findings from the experiments with gambles, ice-cream preferences, and life-saving interventions, suggest that, without a context to give affective perspective to quantities of dollars, ice cream, and lives, these quantities may convey little meaning. Amounts of anything, no matter how common or familiar or intrinsically important, may in some circumstances not be evaluable.
4. Probabilities or proportions, however, are often highly evaluable, reflecting the ease with which people recognize that a high probability of a desirable outcome is good and a low probability is bad. When the quantities or outcomes to which these probabilities apply are affectively pallid, probabilities carry much more weight in judgments and decisions. Just the opposite occurs when the outcomes have precise and strong affective meanings – variations in probability carry too little weight.

THE AFFECT HEURISTIC IN JUDGMENTS OF RISK AND BENEFIT

Another stream of research that, in conjunction with many of the findings reported previously, led us to propose the affect heuristic had its origin in the early study of risk perception reported by Fischhoff, Slovic, Lichtenstein, Reid, and Coombs (1978). One of the findings in this study and numerous subsequent studies was that perceptions of risk and society's responses to risk were strongly linked to the degree to which a hazard evoked feelings of dread (see also Slovic, 1987). Thus, activities associated with cancer are seen as riskier and more in need of regulation than activities associated with less dreaded forms of illness, injury, and death (e.g., accidents).

A second finding in the study by Fischhoff et al. (1978) was even more instrumental in the study of the affect heuristic: The finding that judgments of risk and benefit are negatively correlated. For many hazards, the greater the perceived benefit, the lower the perceived risk and vice versa. Smoking, alcoholic beverages, and food additives, for example, tend to be seen as very high in risk and relatively low in benefit, whereas vaccines, antibiotics, and X-rays tend to be seen as high in benefit and relatively low in risk. This negative relationship is noteworthy because it occurs even when the nature of the gains or benefits from an activity is distinct, and qualitatively different from the nature of the risks. That the inverse relationship is generated in people's minds is suggested by the fact that risk and benefits generally tend to be positively (if at all) correlated in the world. Activities that bring great benefits may be high or low in risk, but activities that are low in benefit are unlikely to be high in risk (if they were, they would be proscribed).

A study by Alhakami and Slovic (1994) found that the inverse relationship between perceived risk and perceived benefit of an activity (e.g., using pesticides) was linked to the strength of positive or negative affect associated with that activity. This result implies that people base their judgments of an activity or a technology not only on what they *think* about it but also on what they *feel* about it. If they like an activity, they are moved to judge the risks as low and the benefits as high; if they dislike it, they tend to judge the opposite – high risk and low benefit.

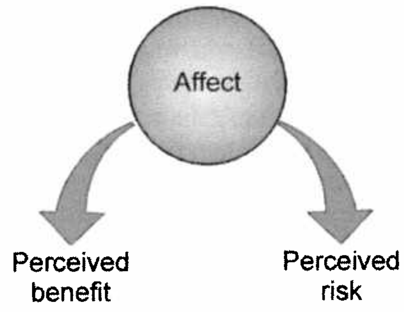

Figure 23.2. A model of the affect heuristic explaining the risk/benefit confounding observed by Alhakami and Slovic (1994). Judgments of risk and benefit are assumed to be derived by reference to an overall affective evaluation of the stimulus item.
Source: Finucane et al. (2000).

Alhakami and Slovic's (1994) findings suggested that use of the affect heuristic guides perceptions of risk and benefit as depicted in Fig. 23.2. If so, providing information about risk should change the perception of benefit and vice-versa (Fig. 23.3). For example, information stating that risk was low for some technology should lead to more positive overall affect that would, in turn,

Figure 23.3. Model showing how information about benefit (A) or information about risk (B) could increase the overall affective evaluation of nuclear power and lead to inferences about risk and benefit that coincide affectively with the information given. Similarly, information could decrease the overall affective evaluation of nuclear power as in C and D. *Source:* Finucane et al. (2000).

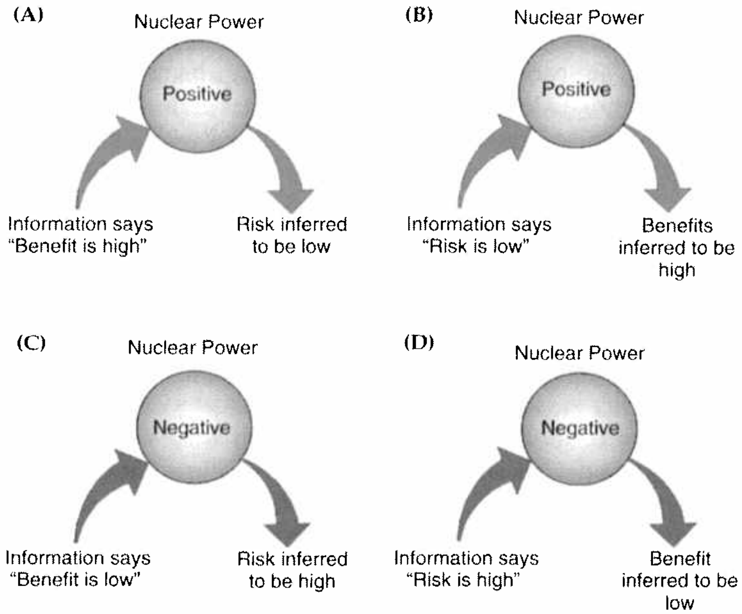

increase perceived benefit. Indeed, Finucane, Alhakami, Slovic, and Johnson (2000) conducted this experiment, providing four different kinds of information designed to manipulate affect by increasing or decreasing perceived risk and increasing or decreasing perceived benefit. In each case there was no apparent logical relation between the information provided (e.g., information about risks) and the nonmanipulated variable (e.g., benefits). The predictions were confirmed. When the information that was provided changed either the perceived risk or the perceived benefit, an affectively congruent but inverse effect was observed on the nonmanipulated attribute as depicted in Fig. 23.3. These data support the theory that risk and benefit judgments are causally determined, at least in part, by the overall affective evaluation.

The affect heuristic also predicts that using time pressure to reduce the opportunity for analytic deliberation (and thereby allowing affective considerations freer rein) should enhance the inverse relationship between perceived benefits and risks. In a second study, Finucane et al. (2000) showed that the inverse relationship between perceived risks and benefits increased greatly under time pressure, as predicted. These two experiments with judgments of benefits and risks are important because they support the contention by Zajonc (1980) that affect influences judgment directly and is not simply a response to a prior analytic evaluation.

Further support for the model in Fig. 23.2 has come from two very different domains – toxicology and finance. Slovic, MacGregor, Malmfors, and Purchase (1999) surveyed members of the British Toxicological Society and found that these experts, too, produced the same inverse relation between their risk and benefit judgments. As expected, the strength of the inverse relation was found to be mediated by these experts' affective reactions toward the hazard items being judged. In a second study, these same toxicologists were asked to make a "quick intuitive rating" for each of 30 chemical items (e.g., benzene, aspirin, second-hand cigarette smoke, dioxin in food) on an affect scale (bad–good). Next, they were asked to judge the degree of risk associated with a *very small exposure to the chemical*, defined as an exposure that is less than 1/100th the exposure level that would begin to cause concern for a regulatory agency. Rationally, because exposure was so low, one might expect these risk judgments to be uniformly low and unvarying, resulting in little or no correlation with the ratings of affect. Instead, there was a strong correlation across chemicals between affect and judged risk of a very small exposure. When the affect rating was strongly negative, judged risk of a very small exposure was high; when affect was positive, judged risk was small. Almost every respondent (95 out of 97) showed this negative correlation (the median correlation was −.50). Importantly, those toxicologists who produced strong inverse relations between risk and benefit judgments in the first study also were more likely to exhibit a high correspondence between their judgments of affect and risk in the second study. In other words, across two different tasks, reliable individual differences emerged in toxicologists' reliance on affective processes in judgments of chemical risks.

In the realm of finance, Ganzach (2000) found support for a model in which analysts base their judgments of risk and return for unfamiliar stocks upon a global attitude. If stocks were perceived as good, they were judged to have high return and low risk; if they were perceived as bad, they were judged to be low in return and high in risk. However, for familiar stocks, perceived risk and return were positively correlated, rather than being driven by a global attitude.

JUDGMENTS OF PROBABILITY, RELATIVE FREQUENCY, AND RISK

The affect heuristic has much in common with the model of "risk as feelings" proposed by Loewenstein et al. (2001) and with dual process theories put forth by Epstein (1994), Sloman (1996), and others. Recall that Epstein argues that individuals apprehend reality by two interactive, parallel processing systems. The *rational* system is a deliberative, analytical system that functions by way of established rules of logic and evidence (e.g., probability theory). The *experiential* system encodes reality in images, metaphors, and narratives to which affective feelings have become attached.

To demonstrate the influence of the experiential system, Denes-Raj and Epstein (1994) showed that, when offered a chance to win a prize by drawing a red jelly bean from an urn, subjects often elected to draw from a bowl containing a greater absolute number, but a smaller proportion, of red beans (e.g., 7 in 100) than from a bowl with fewer red beans but a better probability of winning (e.g., 1 in 10). For these individuals, images of 7 winning beans in the large bowl appeared to dominate the image of 1 winning bean in the small bowl.

We can characterize Epstein's subjects as following a mental strategy of "imaging the numerator" (i.e., the number of red beans) and neglecting the denominator (the number of beans in the bowl). Consistent with the affect heuristic, images of winning beans convey positive affect that motivates choice.

Although the jelly bean experiment may seem frivolous, imaging the numerator brings affect to bear on judgments in ways that can be both nonintuitive and consequential. Slovic, Monahan, and MacGregor (2000) demonstrated this in a series of studies in which experienced forensic psychologists and psychiatrists were asked to judge the likelihood that a mental patient would commit an act of violence within 6 months after being discharged from the hospital. An important finding was that clinicians who were given another expert's assessment of a patient's risk of violence framed in terms of relative frequency (e.g., "Of every 100 patients similar to Mr. Jones, 10 are estimated to commit an act of violence to others") subsequently labeled Mr. Jones as more dangerous than did clinicians who were shown a statistically "equivalent" risk expressed as a probability (e.g., "Patients similar to Mr. Jones are estimated to have a 10% chance of committing an act of violence to others").

Not surprisingly, when clinicians were told that "20 out of every 100 patients similar to Mr. Jones are estimated to commit an act of violence," 41% would refuse to discharge the patient. But when another group of clinicians was given the risk as "patients similar to Mr. Jones are estimated to have a 20% chance of committing an act of violence," only 21% would refuse to discharge the patient. Similar results have been found by Yamagishi (1997), whose judges rated a disease that kills 1,286 people out of every 10,000 as more dangerous than one that kills 24.14% of the population.

Unpublished follow-up studies showed that representations of risk in the form of individual probabilities of 10% or 20% led to relatively benign images of one person, unlikely to harm anyone, whereas the "equivalent" frequentistic representations created frightening images of violent patients (example: "Some guy going crazy and killing someone"). These affect-laden images likely induced greater perceptions of risk in response to the relative-frequency frames.

Although frequency formats produce affect-laden imagery, story and narrative formats appear to do even better in that regard. Hendrickx, Vlek, and Oppewal (1989) found that warnings were more effective when, rather than being presented in terms of relative frequencies of harm, they were presented in the form of vivid, affect-laden scenarios and anecdotes. Sanfey and Hastie (1998) found that compared with respondents given information in bar graphs or data tables, respondents given narrative information more accurately estimated the performance of a set of marathon runners. Furthermore, Pennington and Hastie (1993) found that jurors construct narrative-like summations of trial evidence to help them process their judgments of guilt or innocence.

Perhaps the biases in probability and frequency judgment that have been attributed to the availability heuristic may be due, at least in part, to affect. Availability may work not only through *ease* of recall or imaginability, but because remembered and imagined images come tagged with affect. For example, Lichtenstein, Slovic, Fischhoff, Layman, and Combs (1978) invoked availability to explain why judged frequencies of highly publicized causes of death (e.g., accidents, homicides, fires, tornadoes, cancer) were relatively overestimated and underpublicized causes (e.g., diabetes, stroke, asthma, tuberculosis) were underestimated. The highly publicized causes appear to be more affectively charged, that is, more sensational, and this may account both for their prominence in the media and their relatively overestimated frequencies.

FURTHER EVIDENCE

The studies described above represent only a small fraction of the evidence that can be marshaled in support of the affect heuristic. Although we have developed the affect heuristic to explain findings from studies of judgment and decision making (e.g., the inverse relationship between perceived risks and benefits), one can find related proposals in the literature of marketing and social cognition. For example, Wright (1975) proposed the "affect-referral heuristic" as a mechanism by which the remembered affect associated with a product influences subsequent choice of that product (see also Pham, 1998).

Attitudes have long been recognized as having a strong evaluative component (see, e.g., Thurstone, 1928; Edwards, 1957). Pratkanis (1989) defined *attitude* as "a person's evaluation of an object of thought" (p. 72). He went on to propose that attitudes serve as heuristics, with positive attitudes invoking a favoring strategy toward an object and negative attitudes creating disfavoring response. More specifically, he defined the *attitude heuristic* as the use of the evaluative relationship as a cue for assigning objects to a favorable class or an unfavorable class, thus leading to approach or avoidance strategies appropriate to the class. Pratkanis described numerous phenomena that could be explained by the attitude heuristic, including halo effects not unlike the consistency described earlier between risk and benefit judgments (Finucane et al., 2000).

Other important work within the field of social cognition includes studies by Fazio (1995) on the accessibility of affect associated with attitudes, and by Schwarz and Clore (1988) on the role of affect as information.

Returning to the recent literature on judgment and decision making, Kahneman and Ritov (1994) and Kahneman, Schkade, and Sunstein (1998) demonstrated that responses as diverse as willingness to pay for the provision of a public good (e.g., protection of an endangered species) or a punitive damage award in a personal injury lawsuit seem to be derived from attitudes based on emotion rather than on indicators of economic value.

Hsee and Kunreuther (2000) demonstrated that affect influences decisions about whether to purchase insurance. In one study, they found that people were willing to pay twice as much to insure a beloved antique clock (that no longer works and cannot be repaired) against loss in shipment to a new city than to insure a similar clock for which "one does not have any special feeling." In the event of loss, the insurance paid $100 in both cases. Similarly, Hsee and Menon (1999) found that students were more willing to buy a warranty on a newly purchased used car if it was a beautiful convertible than if it was an ordinary looking station wagon, even if the expected repair expenses and cost of the warranty were held constant.

Loewenstein et al. (2001) provide a particularly thorough review and analysis of research that supports their *risk-as-feelings hypothesis*, a concept that has much in common with the affect heuristic. They present evidence showing that emotional responses to risky situations, including feelings such as worry, fear, dread, or anxiety, often diverge from cognitive evaluations and have a different and sometimes greater impact on risk-taking behavior than do cognitive evaluations. Among the factors that appear to influence risk behaviors by acting on feelings rather than cognitions are background mood (e.g., Johnson & Tversky, 1983; Isen, 1993), the time interval between decisions and their outcomes (Loewenstein, 1987), vividness (Hendrickx et al., 1989), and evolutionary preparedness. Loewenstein et al. (2001) invoke the evolutionary perspective to explain why people tend to react with little fear to certain types of objectively dangerous stimuli that evolution has not prepared them for, such as guns, hamburgers, automobiles, smoking, and unsafe sex, even when they recognize the threat at a cognitive level. Other types of stimuli, such as caged spiders, snakes, or heights, which evolution may have prepared us to fear, evoke strong visceral responses even when we recognize them, cognitively, to be harmless.

Individual differences in affective reactivity also are informative. Damasio relied upon brain-damaged individuals, apparently lacking in the ability to associate emotion with anticipated outcomes, to test his somatic-marker hypothesis. Similar insensitivity to the emotional meaning of future outcomes has been attributed to psychopathic individuals and used to explain their aberrant behaviors (Hare, 1965; Patrick, 1994). Using the Damasio card-selection task, Peters and Slovic (2000) found that normal subjects who reported themselves to be highly reactive to negative events made fewer selections from decks with large losing payoffs. Conversely, greater self-reported reactivity to positive events was associated with a greater number of selections from high-gain decks. Thus individual differences in affective reactivity appear to play a role in the learning and expression of risk-taking preferences.

THE DOWNSIDE OF AFFECT

Throughout this chapter, we make many claims for the affect heuristic, portraying it as the centerpiece of the experiential mode of thinking, the dominant mode of survival during the evolution of the human species. However, like other heuristics that provide efficient and generally adaptive responses but occasionally lead us astray, reliance on affect can also deceive us. Indeed, if it was always optimal to follow our affective and experiential instincts, there would have been no need for the rational/analytic system of thinking to have evolved and become so prominent in human affairs.

There are two important ways that experiential thinking misguides us. One results from the deliberate manipulation of our affective reactions by those who wish to control our behaviors. The other results from the natural limitations of the experiential system and the existence of stimuli in our environment that are simply not amenable to valid affective representation. Both types of problems are discussed below.

MANIPULATION OF AFFECT IN OUR DAILY LIVES

Given the importance of experiential thinking, it is not surprising to see many forms of deliberate efforts being made to manipulate affect in order to influence our judgments and decisions. Consider, for example, some everyday questions about the world of entertainment and the world of consumer marketing:

1. Question: Why do entertainers often change their names?
   Answer: To make them affectively more pleasing. One wonders whether the careers of John Denver, Sandra Dee, and Judy Garland would have been as successful had they performed under their real names – Henry Deutschendorf, Alexandra Zuck, and Frances Gumm. Students of *onomastics*, the science of names, have found that the intellectual products of persons with less attractive names are judged to be of lower quality (Harari & McDavid, 1973; Erwin & Calev, 1984) and some have even asserted that the affective quality of a presidential candidate's name influences the candidate's chances of being elected (Smith, 1997).
2. Question: Why do movies have background music? After all, can't we understand the events we are watching and the dialog we are hearing without music?
   Answer: Music conveys affect and thus enhances meaning even for common human interactions and events.
3. Question: Why are all the models in the mail-order catalog smiling?
   Answer: To link positive affect to the clothing they are selling.
4. Question: Why do packages of food products carry all those little blurbs such as "New," "Natural," "Improved," or "98% fat-free"?
   Answer: These are "affective tags" that enhance the attractiveness of the product and increase the likelihood it will be purchased, much as adding "Save 98%" increased the attractiveness of saving 150 lives.

Clearly, entertainers and marketers of consumer products have long been aware of the powerful influence of affect. Perhaps no corporate entities have more zealously exploited consumers' affective sensitivities than the tobacco companies. An ad for Kool Natural Lights, for example, repeats the word *natural* 13 times in a single half-page advertisement (Brown & Williamson Tobacco Company, 1999). The attractive images of rugged cowboys and lush waterfalls associated with cigarette ads are known to all of us. Indeed, affective associations between cigarettes and positive images may begin forming in children as young as 3 years old (Fischer, 1991). As Epstein (1994) observes, "Cigarette advertising agencies and their clients are willing to bet millions of dollars in advertising costs that the ... appeal of their messages to the experiential system will prevail over the verbal message of the Surgeon General that smoking can endanger one's life, an appeal directed at the rational system" (p. 712). Through the workings of the affect heuristic, as explicated by Finucane et al. (2000), we now have evidence suggesting that cigarette advertising designed to increase the positive affect associated with smoking will quite likely depress perceptions of risk. The factual (impassionate) appeal by the surgeon general will likely have little effect.

Attempts at affective manipulation often work directly on language. Communicators desiring to change attitudes toward stigmatized technologies, for example, created "nukespeak" to extol the virtues of *clean bombs* and *peacekeeper missiles*, whereas promoters of nuclear power coined a new term for reactor accidents: *excursions*. Genetically modified food has been promoted as "enhanced" by proponents and "Frankenfood" by opponents.

Manipulation of attitudes and behavior by persuasive argumentation is often quite effective, but at least it tends to be recognized as an attempt to persuade. Manipulation of affect is no less powerful but is made more insidious by often taking place without our awareness. It is unlikely that Hsee's subjects recognized that what they were willing to pay for the used music dictionary was determined far more by the torn cover than by the more important dimension, number of entries.

Legal scholars such as Hanson and Kysar (1999a, 1999b), paying close attention to research on affect and other judgment heuristics, have begun to speak out on the massive manipulation of consumers by the packaging, marketing, and public relations practices of manufacturers. Such manipulation, they argue, renders ineffective three primary forms of legal control over dangerous products – warning requirements, product liability suits, and regulation of advertising. Hanson and Kysar (2001) point to the need for new regulatory strategies that would take into account the full liability of manufacturers who manipulate consumers into purchasing and using hazardous products.

FAILURE OF THE EXPERIENTIAL SYSTEM: THE CASE OF SMOKING

Judgments and decisions can be faulty not only because their affective components are manipulable, but also because they are subject to inherent biases of the experiential system. For example, the affective system seems designed to sensitize us to small changes in our environment (e.g., the difference between 0 and 1 deaths) at the cost of making us less able to appreciate and respond appropriately to larger changes (e.g., the difference between 570 deaths and 670 deaths). Fetherstonhaugh et al. (1997) referred to this insensitivity as *psychophysical numbing*.

Similar problems arise when the outcomes that we must evaluate change very slowly over time, are remote in time, or are visceral in nature. The irrationality of decisions to smoke cigarettes provides dramatic examples of these types of failure (Slovic, 2000, 2001). Despite the portrayal of beginning smokers as "young economists" rationally weighing the risks of smoking against the benefits when deciding whether to initiate that activity (e.g., Viscusi, 1992), research paints a different picture. This account (Slovic, 2001) shows young smokers acting experientially in the sense of giving little or no thought to risks or to the amount of smoking they will be doing. Instead, they go with the affective impulses of the moment, enjoying smoking as something new and exciting, a way to have fun with their friends. Even after becoming "regulars," the great majority of smokers expect to stop soon, regardless of how long they have been smoking, how many cigarettes they currently smoke per day, or how many previous unsuccessful attempts they have experienced. Only a fraction actually quit, despite many attempts. The problem is nicotine addiction, a condition that young smokers recognize by name as a consequence of smoking but do not understand experientially until they are caught up in it.

The process of becoming addicted appears to begin surprisingly soon after one begins to smoke. Research indicates that adolescents begin to show signs of nicotine dependence within days to weeks of the onset of occasional tobacco use (DiFranza et al., 2000). Loewenstein (1999) explains the process of addiction as being governed by immensely powerful visceral factors or cravings that, from an experiential perspective, are very hard to anticipate and appreciate:

> Unlike currently experienced visceral factors, which have a disproportionate impact on behavior, delayed visceral factors tend to be ignored or severely underweighted in decision making. Today's pain, hunger, anger, etc. are palpable, but the same sensations anticipated in the future receive little weight. (p. 240)

The failure of the experiential system to protect many young people from the lure of smoking is nowhere more evident than in the responses to a survey question that asks smokers: "If you had it to do all over again, would you start smoking?" More than 85% of adult smokers and about 80% of young smokers (ages 14 to 22) answer "No" (Slovic, 2001). Moreover, the more individuals perceive themselves to be addicted, the more often they have tried to quit, the longer they have been smoking, and the more cigarettes they are smoking per day, the more likely they are to answer "No."

We can now address a central question posed by Viscusi (1992): "at the time when individuals initiate their smoking activity, do they understand the consequences of their actions and make rational decisions?" Visusi went on to define the appropriate test of rationality in terms of "whether individuals are incorporating the available information about smoking risks and are making sound decisions, given their own preferences" (p. 11).

The data indicate that the answer to Viscusi's question is "No." Most beginning smokers lack the experience to appreciate how their future selves will perceive the risks from smoking or how they will value the tradeoff between health and the need to smoke. This is a strong repudiation of the model of informed rational choice. It fits well with the findings indicating that smokers give little conscious thought to risk when they begin to smoke. They appear to be lured into the behavior by the prospects of fun and excitement. Most begin to think of risk only after starting to smoke and gaining what to them is new information about health risks.

These disturbing findings underscore the distinction that behavioral decision theorists now make between decision utility and experience utility (Kahneman, 1997; Kahneman & Snell, 1992; Loewenstein & Schkade, 1999). Utility predicted or expected at the time of decision often differs greatly from the quality and intensity of the hedonic experience that actually occurs.

CONCLUSION

We hope that this rather selective and idiosyncratic tour through a mélange of experiments and conjectures has conveyed the sense of excitement we feel toward the affect heuristic. This heuristic appears at once both wondrous and frightening: wondrous in its speed, and subtlety, and sophistication, and its ability to "lubricate reason"; frightening in its dependency upon context and experience, allowing us to be led astray or manipulated – inadvertently or intentionally – silently and invisibly.

It is sobering to contemplate how elusive meaning is, due to its dependence upon affect. Thus the forms of meaning that we take for granted and use to justify immense effort and expense toward gathering and disseminating "meaningful" information, may be illusory. We cannot assume that an intelligent person can understand the meaning of and properly act upon even the simplest of numbers such as amounts of money, not to mention more esoteric measures or statistics, unless these numbers are infused with affect.

Contemplating the workings of the affect heuristic helps us appreciate Damasio's (1994) contention that rationality is not only a product of the analytical mind, but of the experiential mind as well:

> The strategies of human reason probably did not develop, in either evolution or any single individual, without the guiding force of the mechanisms of biological regulation, of which emotion and feeling are notable expressions. Moreover, even after reasoning strategies become established . . . their effective deployment probably depends, to a considerable extent, on a continued ability to experience feelings. (p. xii)

Ironically, the perception and integration of affective feelings, within the experiential system, appears to be the kind of high-level maximization process postulated by economic theories since the days of Jeremy Bentham. These feelings form the neural and psychological substrate of utility. In this sense, the affect heuristic enables us to be rational actors in many important situations. But not in all situations. It works beautifully when our experience enables us to anticipate accurately how we will like the consequences of our decisions. It fails miserably when the consequences turn out to be much different in character than we anticipated.

The scientific study of affective rationality is in its infancy. It is exciting to contemplate what might be accomplished by future research designed to help humans understand the affect heuristic and use it beneficially.

24. Individual Differences in Reasoning: Implications for the Rationality Debate?

Keith E. Stanovich and Richard F. West

The interpretation of the gap between descriptive and normative models in the human reasoning and decision making literature has been the subject of contentious debate since the early 1980s (Baron, 1994; Cohen, 1981, 1983; Evans & Over, 1996; Gigerenzer, 1996; Kahneman, 1981; Kahneman, Slovic, & Tversky, 1982; Kahneman & Tversky, 1983, 1996; Koehler, 1996; Nisbett & Ross, 1980; Stein, 1996), a debate that has arisen because some investigators wish to interpret the gap between the descriptive and the normative as indicating that human cognition is characterized by systematic irrationalities. Due to the emphasis that these theorists placed on reforming human cognition, they have been labelled the *Meliorists* by Stanovich (1999). Disputing this contention are numerous investigators (termed the *Panglossians*; see Stanovich, 1999) who argue that there are other reasons why reasoning might not accord with normative theory – reasons that prevent the ascription of irrationality to subjects (Cohen, 1981; Stein, 1996). First, instances of reasoning might depart from normative standards due to performance errors – temporary lapses of attention, memory deactivation, and other sporadic information processing mishaps. Second, there may be stable and inherent computational limitations that prevent the normative response (Cherniak, 1986; Goldman, 1978; Harman, 1995; Oaksford & Chater, 1993, 1995, 1998; Stich, 1990). Third, in interpreting performance, we might be applying the wrong normative model to the task (Koehler, 1996). Alternatively, we may be applying the correct normative model to the problem as set, but the subject might have construed the problem differently and be providing the normatively appropriate answer to a different problem (Adler, 1984, 1991; Berkeley & Humphreys, 1982; Broome, 1990; Hilton, 1995; Schwarz, 1996).

However, in referring to the various alternative explanations (other than systematic irrationality) for the normative/descriptive gap, Rips (1994) warns that "a determined skeptic can usually explain away any instance of what seems at first to be a logical mistake" (p. 393). The most humorous version of this argument was made by Kahneman (1981) in his dig at the Panglossians, who seem to have only two categories of errors, "pardonable errors by subjects and unpardonable ones by psychologists" (p. 340).

These comments by Rips (1994) and Kahneman (1981) highlight the need for principled constraints on the alternative explanations of normative/descriptive discrepancies. In this chapter, we describe a research logic aimed at inferring