Weber, E. U. & Johnson, E. J. (in press). Constructing preferences from memory. In: Slovic, P., & Lichtenstein, S. (Eds.), The Construction of Preference. New York NY: Cambridge University Press.

Chapter 21

CONSTRUCTING PREFERENCES FROM MEMORY

ELKE U. WEBER AND ERIC J. JOHNSON

Our memories define who we are and what we do. Aside from a few preferences hardwired by evolution, they also define what we like and how we choose. The increasingly more hierarchical coding of experience in memory makes a gourmet out of a gourmand. Grieving as a process of letting-go of departed loved ones involves the extinction of associations between them and a host of daily encountered stimuli that serve as constant reminders of their absence. Memory processes in the construction of value and preference even have a utility of their own. We savor the memory and anticipation of pleasant choice options and may delay choice to lengthen the pleasurable experience (Loewenstein, 1988). Yet despite this wealth of evidence of the involvement of memory processes in preference and choice, their role in preference construction has largely been ignored in existing models of judgment and decision making (Johnson & Weber, 2000; Weber, Goldstein, & Barlas, 1995).

In this chapter, we argue that our view of preference changes if conceptualized explicitly as the product of memory representations and memory processes. We draw on insights about the functions and operations of memory provided by cognitive psychology and social cognition to show that memory plays a crucial role in preference and choice. Economics and behavioral decision research have traditionally been silent on the nature and role of memory processes in preference formation, preferring to treat preference and choice in an "as-if" fashion as mathematical transformations and integrations of the features of choice objects, taking inspiration from psychophysics rather than cognitive psychology. Recently however, a number

of researchers have started to highlight the influence of implicit memory processes (Arkes, 2001) and (multiple) memory representations (see Arkes, 2001; Reyna, Lloyd, & Brainerd, 2003 for a similar discussion). Much of this work has concentrated on inferential processes, such as the notable computational memory process model MINERVA-DM designed to explain probabilistic inference and judgment (Dougherty, Gettys, & Ogden, 1999; Dougherty, Gronlund, & Gettys, 2003), and work on false memories (Reyna & Lloyd, 1997). Within research on preference and choice, image theory (Beach & Mitchell, 1987) has been a pioneer in assigning an important role to memory matching processes, as has naturalistic decision making (Klein, Orasanu, Calderwood, & Zsambok, 1993), which emphasizes recognition priming and the Brunswikian notions of functionalism and representative design (Hammond, 1990). Recent work on decision modes (Ames, Flynn, & Weber, 2004; Weber, Ames, & Blais, 2005; Weber & Hsee, 2000) has documented that choice processes based on recognition of a class of decision situations for which a prescribed best action or production rule exists are just as prevalent as computational choice processes.

In this chapter, we examine memory processes in preference construction and choice at a more "micro" and process-oriented level than previous investigations into the role of memory processes, but at a level that is cognitive and functional, rather than computational. We suggest that a consideration of properties of memory representation and retrieval can provide a unifying explanatory framework for some seemingly disparate preference phenomena.

Preferences as Functional Relationships

Current formal representations of preferences within behavioral decision research map external stimuli to their internal experience, such as prospect theory's value function (Kahneman & Tversky, 1979) and von Neumann-Morgenstern (1947) utilities. Similarly, indifference functions describe equivalent levels of hedonic value or utility generated by combinations of different levels of two or more external attributes, such as reference dependent prospect theory (Tversky & Kahneman, 1991) or standard models of consumer choice (Varian, 1999). This theoretical framework has had a long and important history within economics and psychology. Functional mapping representations of preference have been obviously useful and are particularly amenable to mathematical formalization. They make the following explicit assumptions about the nature of preference in formal axioms and implicit assumptions about the psychology of preferences:

Stability. The mappings between choice attributes and utility are usually assumed to remain constant over time and across measurement procedures, a property often referred to as procedure invariance (Tversky, Sattath, & Slovic, 1988).

Continuity. Almost all representations of preferences use continuous functions to portray the mapping between all attribute levels and utility (in a utility function) or between combinations of attributes (in an indifference curve). Discontinuities do not exist within the defined range of experience.

Precision. Because value, utility, and tradeoff functions are represented by lines, the mappings are infinitely precise. A given level of income, for example, generates an infinitely resolvable and precise amount of utility. For two attributes, an exact equivalence between differing amounts of each attribute can be calculated.

Decision research over the last 20 to 30 years has questioned these assumptions. The first property to be challenged, based upon changes in preference revealed by different response modes, was stability. Preferences were characterized as 'labile,' since they often seemed to vary across response modes (Lichtenstein & Slovic, 1971), leading to the indexing of utility functions.

By the early 1990s, researchers were starting to examine situations where preferences might not exist at all, but had to be constructed in response to preference tasks (Fischhoff, 1991; Payne, Bettman, & Johnson, 1992; Slovic, 1995). As an as-if model of preference and choice, the functional relationship framework is naturally moot on the issue of preference construction, that is, it cannot make predictions about systematic differences in the construction process as a function of variables not included in the model that, nevertheless, have been shown to influence the choice process.

It may be best to appreciate the functional relationships framework as a metaphor that has yielded important insights into decision and choice but that, like any metaphor, is incomplete, highlighting some aspects of the phenomena under study, but letting others reside in the shadows. We suggest that it may be useful to conceptualize preferences under a different metaphor, in particular as the output of the human memory and inference system, since there is little reason that knowledge related to preferences should not possess the properties and characteristics of other types of knowledge. This focus might illuminate characteristics of preferences (as memories) that might be hidden from other vantage points. Our purpose is to lend some unifying theoretical structure to the now widely held view that preferences are often labile and constructed and to suggest possible components of this constructive process. To do so, we adopt some well-documented, high-level properties of human memory and apply them to preference formation. The preferences-as-memory (PAM) approach suggests that preferences are neither constructed from first principles anew on each occasion nor completely stable and immutable.

Conceptualizing preference as the product of memory processes and memory representations predicts systematic deviations from the functional account. Where functional approaches assume stability, PAM suggests that preferences may differ as a function of differences in short term accessibility as the result of priming. Where functional accounts assume continuity and precision, the discrete nature of memory episodes and the presence of excitatory and inhibitory processes in memory (i.e., recalling one aspect of a choice situation may enhance or inhibit recollection of other aspects) suggests that preference can be discontinuous and lumpy, and can take on different values in a path dependent fashion. In contrast to the functional account's idea of precision, PAM suggests that memory and thus preference is reactive, that is, the very act of assessing a preference can change it.

Preferences as Memory (PAM)

The preferences-as-memory (PAM) framework assumes that decisions (or valuation judgments) are made by retrieving relevant knowledge (attitudes, attributes, previous preferences, episodes, or events) from memory in order to determine the best (or a good) action. This process is functionally equivalent to constructing predictions about one's experience of the consequences of different actions, that is, what Kahneman and collaborators (Kahneman & Snell, 1990; Schkade & Kahneman, 1998) call predicted utility. Rather than accessing the equivalent of stable, continuous, and infinitely resolvable utility and indifference curves, we suggest that people make an attempt to retrieve past reactions and associations to similar situations. It is in the nature of human memory that such retrieval attempts do not always generate a single and precise answer. Because retrieval depends upon prior encoding, memory representation, memory query, situational context, and prior attempts at retrieval, the results may vary.

A small but important number of robust empirical properties of the human memory system may play a role in the process of constructing predicted utility. We focus on three aspects of the process: (a) memory interrogation, that is, the queries posed to memory, (b) the

accessibility of information in memory that is relevant to the queries, and (c) memory representation, that is, the structure of memory which reflects the structure of the world and our task environments in a Brunswikian fashion.

Interrogating Memory

A major assumption of PAM's account of how utility predictions (and thus preferences) are generated is that people consult their memory (or the external environment) with a series of component queries about the attributes of choice alternatives, in particular their merits or liabilities. When asked to pick their preferred CD from a choice set of three, for example, people consult their memory about previous experiences with the same or similar CDs. Because it helps with evidence generation and integration, these queries are typically grouped by valence; for example, memory is first queried about what I like about a given CD, and only after no additional positive attributes are generated may a query about negative attributes ensue. We also assume that most tasks suggest a natural way to the order in which queries are posed. Being asked to pick one CD out of three triggers queries about positive attributes first, whereas being asked to reject one of those CDs naturally triggers queries about each CD's negative attributes first (Shafir, Simonson, & Tversky, 1993). A home-owner asked to provide a selling price for her house will first consult her memory about positive features of the house before considering downsides, whereas a potential buyer may pose these queries in the opposite order (Birnbaum & Stegner, 1979). Because of interference processes, described below, the order on which queries are posed is important; that is, it affects the answers provided by memory and thus preference.

We do not assume that such queries are explicit or conscious, although they can be. Instead, we assume that they occur without conscious effort and as a natural part of automatic preference-construction processes. In this sense all decisions are reason based, to borrow Shafir, Simonson, et al.'s (1993) term. The major difference between that work and ours is that we imbed our order-of-queries-based explanation into a broader theory about memory structure and memory processes. We emphasize the measurement and manipulation of reasons as a diagnostic tool and assert that the type, valence, and number of reasons obtained depend systematically upon the valuation or choice task (e.g., pick vs. reject, buying vs. selling) and the accessibility of relevant information in memory. Recent work by McKenzie and colleagues (e.g., McKenzie & Nelson, 2003) suggests that different semantic frames that might be seen as logically equivalent (e.g., a glass being half-full or half-empty) linguistically transmit different information. A PAM interpretation of this view is that different frames lead the decision maker to generate different queries (e.g., 'where did the contents of the glass come from?' vs. 'what happened to the other half of the glass's contents?'). Also relevant is work by Fischer and colleagues (Fischer, Carmon, Ariely, & Zauberman, 1999) suggesting that different response modes have different goals. A PAM interpretation hypothesizes that different goals naturally evoke different queries or different orders in which queries are posed.

Memory Accessibility

Priming

Social cognition and memory research have demonstrated that the presentation of a stimulus can produce a short-term, transient increase in the accessibility of the same stimulus and related concepts (see Higgins and Kruglanski, 1996, for a review). In this phenomenon, called priming, previous activation of a memory node affects later accessibility, resulting in both shorter reaction times and greater likelihood of retrieval. A recent field experiment (North, Hargreaves, & McKendrick, 1999) linked priming and revealed preference. A supermarket where wine was sold played equally pleasant music with either strongly stereotypical French or

German cultural associations on alternative days and measured the amount of wine that was sold from each country of origin. French wines sold more strongly on the days French music was played, and German wines sold more briskly when German music was played. Type of music accounted for almost a quarter of the variance in wine sales, even though respondents were unaware that the music affected their choice, with a majority denying such an effect. Priming also produced differences in behavior in a study by Bargh, Chen, and Burrows (1996), who primed (in an "unrelated experiments" paradigm) constructs such as "politeness" or "being elderly." When primed with the construct elderly (by being given an adjective rating task that involved adjectives related to the construct), for example, respondents walked away from the experiment more slowly than those who were not so primed. Walking speed was not mediated by other possible factors such as a change in mood, and the priming effect again appeared to occur without awareness. In an important application of priming, Siegel, Krank, and Hinson's (1987) work on classical conditioning mechanisms in physical addictions suggests that the best predictor of continued abstinence after a rehab program is physical relocation (even to non-drugfree environments), since relocation removes the environmental cues associated with previous drug use that elicit withdrawal symptoms in recovering addicts.

Mandel and Johnson (2002) examined the effects of priming on consumer choice by selectively priming particular product attributes by using the background ("wallpaper") of the initial page of an online-shop website. Preliminary studies had shown that a background of a blue sky with clouds primed the attribute "comfort" while leaving the accessibility of other attributes unchanged. When deciding between different couches, respondents who had seen this background on the initial store webpage were more likely than control subjects to choose a more comfortable but more expensive couch over a less comfortable but cheaper couch. Similar results were shown for other primes and products.

Priming also has the potential to explain semantic framing effects reported in the behavioral decision literature, such as Levin and Gaeth's (1988) studies of differential preference for fried ground beef described as either "being 90% lean" or "containing 10% fat"; Bell and Loftus' (1989) study of differential speed estimates provided for a car observed in a videotaped car accident depending on whether respondents were asked to estimate the speed of the car when it "hit" versus "smashed into" the parked truck; and McKenzie's work described above. In these studies, the details of the way in which fat content, traveling speed, or glass-contents are described prime different components of the objects' representation in memory. These differences in short-term activation result in different answers to explicit evaluative questions, despite identical sensory information.

Memory is Reactive

When you listen to a CD on your stereo, the fact that you have accessed this string of 1's and 0's does not change the music recorded on the disk. This is not true of human memory. "Tests of memory are not neutral events that merely assess the state of a person's knowledge; tests also change or modify memories.... Changes can be positive, aiding later retentions, or negative, causing forgetting, interference, and even false recollection" (Roedinger & Guynn, 1996, p. 225). Negative effects fall into the category of interference effects, which are mostly limited to short time spans. Here we concentrate on positive effects, namely, how access can both increase short-term accessibility and change the long-term content of memory. The picture that emerges is that memory, and therefore preferences, are reactive to measurement. By analogy

to Heisenberg's uncertainty principle, the measurement can influence the behavior of the object under observation.

Short-term effects. Studies of anchoring suggest that memory accessibility and preference can be changed by asking a prior question, even if the answer to this question should be irrelevant to subsequent tasks. Chapman and Johnson (1999) asked a group of students if their selling price for a gamble was greater or less than the last four digits of their Social Security number, converted into a price. While normatively irrelevant, this comparative judgment influenced their selling prices for the gamble, despite respondents' denial of any such effect. Contemporary accounts of anchoring attribute this effect to a short-term increase in the accessibility of information in memory. The Selective Accessibility Model (Strack & Mussweiler, 1997) and the Anchoring as Activation (Chapman & Johnson, 1999) perspective suggest that, despite its substantive irrelevance, an anchor makes anchor-consistent information more accessible and therefore more likely to be included in a subsequent judgment. Chapman and Johnson (2002) reviewed a large body of empirical evidence supporting this interpretation (Chapman & Johnson, 1994, 1999; Mussweiler & Strack, 2001b; Mussweiler, Strack, & Pfeiffer, 2000; Strack & Mussweiler, 1997) and found information consistent with the anchor to be more accessible than inconsistent information, as revealed by reaction time differences. For anchoring to occur, the anchor must be used in a preliminary judgment (and not just be present). Anchoring can be enhanced by increasing the knowledge base relevant to the judgment. Finally, asking people to consider other judgments can debias anchoring. While accessibility may not be sufficient to explain all anchoring effects (Epley & Gilovich, 2001), accessibility-mediated anchoring effects are strong and robust, and persist in the presence of significant accuracy incentives, experience, and market feedback (Ariely, Loewenstein, & Prelec, 2003).

Long-term effects. Queries about possible choice options seem to not only generate short-term changes in the accessibility of related information, but also actually change memory representation in a more permanent fashion. Students asked prior to election day whether they intended to vote predicted that they would do so in numbers greater than the actual voting percentage later observed for comparable control-group subjects who were not asked that question (Greenwald, Carnot, Beach, & Young, 1987). More interestingly, answering the question affirmatively increased students' subsequent actual voting behavior, months after the initial question had been posed and without any conscious memory of ever having answered it. Hirt and Sherman (1985) term this widely replicated phenomenon the self-correcting nature of errors in prediction. In the context of consumer choice, Morwitz and collaborators (Fitzsimons & Morwitz, 1996; see also Morwitz, 1997; Morwitz, Johnson, & Schmittlein, 1993; Morwitz & Pluzinski, 1996) have demonstrated that measuring purchase intentions can change purchases. Actual purchases of a personal computer increased by about one third as the result of answering a single purchase intent question several months previously. Just as we can induce "false memories" in people (e.g., of being lost as a child in a shopping mall; Loftus & Pickrell, 1995), we can induce "false preferences" in people by asking them for behavioral intentions in ways that will result in a biased answer, capitalizing on such things as social desirability or wishful thinking. Such effects cannot be explained by conscious attempts at consistency over time, as explicit memories of prior measurement episodes typically do not exist.

Interference and Inhibition

If priming increases accessibility, other tasks or events can affect memory accessibility negatively. The classic memory phenomenon of inhibition or interference is illustrated by the following example. Imagine you have just moved and have not yet completely committed your

new telephone number to memory. You try to have the newspaper delivered to your new address and are asked for your new phone number. You almost had it recalled when the customer service representatives reads out your old number, asking you if that is it. You now find it impossible to recall the new number.

While 80 years of research on interference effects have produced a plethora of explanations and theoretical mechanisms, they all share a common idea: When one component of a memory structure is recalled, the recall of others components that could have been response competitors is temporarily suppressed. Inhibition as the result of prior recall of related and competing material is one of the oldest and most developed memory phenomena (see M. C. Anderson & Neely, 1996, for a review of theories and results). Interference effects have been shown in many different experimental paradigms, including the effect of new material on the recall of old (retroactive interference) and the effect of old material on new (proactive interference). Interference has been shown to occur for both semantic and episodic memory and for verbal as well as non-verbal materials, such as visual stimuli and motor skills. Recent studies have demonstrated conditions under which implicit memory also shows interference effects (Lustig & Hasher, 2001).

Are there functional advantages to memory processes such as priming or inhibition? Like other species, humans did not evolve to ponder and analyze, but to act in ways that maximize survival and inclusive fitness. The effective use of heuristics documented by Payne, Bettman and Johnson (1990) and Gigerenzer, Todd, and The ABC Group (1999) suggest that humans are probably wired to pick a good action/option quickly. Excitatory and inhibitory memory activation processes facilitate the fast emergence of a response/action/decision that, in the past, has been associated with a successful outcome. Task goals and choice context determine the focus of attention, which translates into a series of implicit queries of external and internal knowledge bases, often executed in a task-specific order. Queries, in turn, result in increased activation (priming) of response-consistent information and decreased activation (inhibition) of response-inconsistent information. Russo and colleagues' (Carlson & Russo, 2001; Russo, Meloy, & Medvec, 1998; Russo, Meloy, & Wilks, 2000) evidence of very early pre-decision polarization and bias in the exploration of choice alternatives in favor of an initially selected alternative is also consistent with a strategic semantic priming and inhibition explanation. An early focus on one choice alternative (perhaps because it scored highly on an attribute that was randomly examined first) will result in greater accessibility of features consistent with it and reduced accessibility of features inconsistent with it.

Memory interference may help explain established phenomena in the behavioral decision literature and provide predictions for other effects. The basic idea is that the natural and sequential order of different queries used in the process of constructing preference or other judgments produces interference, in the sense that the responses to earlier queries inhibit possible responses to later queries. Reyna et al. (2003) suggested, for example, that interference can occur between verbatim, gist, and procedure representations. Hoch (1984, 1985) examined interference in the prediction of preferences and purchase intentions. Respondents were asked to provide reasons why they would or would not buy a consumer product in the future. Hoch counterbalanced the order of the two tasks in a between-subject design, arguing that the first task caused interference with the second. Consistent with this hypothesis and the PAM framework, he found that the first task generated more reasons than the second. Respondents were more likely to predict that they would purchase the item when they generated reasons for buying it first, even though everyone answered both types of questions. To demonstrate that this effect was due to

memory interference, Hoch separated the reasons-generation task in time from the judgment of purchase intention. Consistent with the fact that interference is a transient phenomenon, he no longer found an effect of reasons-generation on purchase intentions.

While not motivated by a PAM (and, in particular, memory interference) perspective, several studies have manipulated the natural order of queries, often in an attempt to debias judgments. Koriat, Lichtenstein and Fischhoff (1980) argued that, when asked for a confidence judgment in a general knowledge forced-choice quiz, people naturally ask first why their choice is correct and only then examine reasons why their choice might be wrong. A PAM interpretation suggests that this order will inhibit generation of con-reasons, leading to a biased ratio of pro/con evidence and thus overconfidence. The interference account suggests that overconfidence occurs because people naturally generate pro-reasons first, not necessarily because they have a motivational stake in generating more pro- than con-reasons. Consistent with this explanation, Koriat et al. showed that asking people to generate reasons why an answer might be wrong before generating reasons why the answer might be right diminished overconfidence. Similarly, asking respondents in anchoring studies to first generate reasons why an anchor might be irrelevant minimized anchoring bias (Chapman & Johnson, 1999; Mussweiler et al., 2000). Current research being conducted in our lab suggests that a similar decomposition of judgment and choice tasks into a series of component queries, combined with interference processes (by which the answers to earlier queries inhibit the answers to later queries) may be responsible for such phenomena as loss aversion and time discounting.

The Structure of Representations

To better predict how memory processes affect accessibility and preferences, we need a fuller understanding of the structure of memory representations that are relevant to preference construction. Not all knowledge is interconnected. For reasons of cognitive economy, concepts vary in their degree of interconnectedness, which reflects, at least in part, the natural co-occurrence of concepts in our daily experience in natural environments (Sherman, McMullen, & Gavanski, 1992). Most commodities (e.g., a house you plan to buy) have rich, highly structured, strongly hierarchical representations. Currencies (e.g., money, time), on the other hand, another important category in trades and other transactions, probably have a very different (more impoverished, flatter, and less connected) representation. Memory theorists argue that the number of connections between a concept node and associated subordinate nodes determines the likelihood of retrieval of any given subordinate node, a phenomena sometimes called the fan effect (J. R. Anderson & Reder, 1999; Reder & Anderson, 1980) or cue overload (M. C. Anderson & Spellman, 1995). The basic intuition is that for a given memory cue, the probability that associated information is retrieved is a diminishing function of the number of connected nodes. Thus, ceteris paribus, memory cues with many associates will result in less recall of any one associate than cues with fewer associates. The hierarchical structure of human memory representation is probably an adaptive response to this constraint. In order to recall rich sets of facts, people organize this information into hierarchies of information and sub-information, a hallmark of expertise. A hierarchical category structure ensures that only a manageable number of relevant items are interconnected at any given level. Consider the category "mug." It may consist of subcategories such as coffee mug and beer mug, each subcategory containing fewer links than a representation without subcategories. Properties of the category as a whole (mugs have handles and are heavier and larger than cups) are associated with the higher level category, producing a more efficient representation, less prone to cue overload. We hypothesize that representations of currencies, such as money or time, differ from the hierarchical representations

of commodities such as mugs. The non-hierarchical representation and non-systematic patterns of connections for currencies, and in particular for functional properties of currencies (e.g., "things to do with $100 or 5 hours of time") lead to impaired retrieval. Imagine you were asked to name an approximate price for the following objects (possible responses are in parentheses): A newsstand copy of Time Magazine ($2.95), a vending machine can of Diet Coke ($1.00), a round trip airline ticket from Los Angeles to Chicago ($400). Introspection and pilot testing in our lab indicates that this is a fairly easy task. Respondents name prices that they feel confident about within 2 to 3 seconds. Now consider the opposite task, naming something that costs $2.95, $1.00, or $400, respectively. Introspection and pilot data indicate this is a much harder task, requiring much more time, presumably because of the hypothesized asymmetry in memory connectivity as the result of the structure of memory. For Western respondents, chances are that any commodity that has a price-tag associated with it has a hierarchical structure. On the other hand, we don't have natural (or even easily generated ad-hoc) categories of the kind "things to buy with $400."

Other evidence of this commodity/currency retrieval asymmetry comes from statements made while generating buying versus selling prices in endowment effect experiments conducted in our lab. While people in both the buying and selling condition easily generate many statements (both good and bad) about the mug, statements about good or bad attributes of the money they could receive or keep are much less frequent, and mentions of alterative uses of the money are almost non-existent.

Differences in the structure of memory for different classes of trading objects may lie at the root of other phenomena in decision research, such as people's failure to attend sufficiently to the opportunity costs of their choices or actions (Thaler, 1999). While it may be easy to generate a large number of reasons for the purchase of a particular vacation home, the structure of memory for currencies and their uses makes it very hard to generate alternative uses for the required purchase price. A deeper understanding of the way in which people organize and cluster their knowledge about the world may be helpful in modeling behavior in many preference and choice tasks. The current inability of behavioral decision researchers to provide ex-ante predictions about such things as the reference point(s) used in a given task (Fischhoff, 1983) or about the number and type of mental accounts used seriously curtails our ability to apply our most popular behavioral models. We suggest that experimental paradigms and procedures designed to reveal the organization and structure of memory in a particular content domain may help with such predictions.

Many behavioral decision models use similarity (e.g., between choice outcomes, between presented outcomes and some memory referent) as a variable in their predictions, typically without specifying how such similarity could or should be assessed. Judgment heuristics such as representativeness (Tversky & Kahneman, 1974) make implicit use of the category structure of human knowledge, including people's ability to generate and use ad-hoc categories ("feminist bank tellers"). A better understanding of the organization and structure of preference-related knowledge and memory will lend greater predictive power to such models.

Caveats

Alternative Theories

We do not (necessarily) conceive of the PAM interpretations of the behavioral decision research phenomena discussed in this chapter as alternative explanations to those offered previously by others, but instead as a cognitive process-level account of the essence of other explanations or descriptions (e.g., myopia or hyperbolic discounting to describe time discounting). What does a cognitive process-level account add to high-level descriptive labels such as myopia or loss aversion? First, we hope that PAM explications of such constructs will

provide unifying theory across disparate phenomena. For example, similarities in effects between temporal discounting and loss aversion have been noted as a curious phenomenon (Prelec & Loewenstein, 1991). A PAM account of those two classes of effects may show that this similarity in effects is the result of similarity in fundamental memory processes in preference construction. Second, process-level explanations are of crucial importance when addressing the normative status of violations of traditional preference models and in designing effective interventions. Understanding the causes of loss aversion would help us characterize it as either a decision error in generating predicted utility or as a valid factor in predicting utility. Further, if we choose to change elicitation procedures, a process-level explanation will be helpful in engineering better alternatives.

In other cases, PAM accounts of phenomena *are* offered as alternative explanations that differ from others accounts (e.g., a desire for self-consistency to explain priming effects), by postulating that the phenomena can occur without the awareness of the decision maker. An implicit null hypothesis of the PAM theoretical structure is that many phenomena that may seem motivational in nature are actually sufficiently explained by basic cognitive processes. The PAM framework does not exclude the possibility that motivational variables may play important roles. It simply tries to see how far it can get without them.

Level of Analysis

The PAM framework outlined in this proposal is intended as much more than simply a metaphor; rather, it is a process-level account, albeit one at a relatively high level of abstraction. It is probably premature to worry about the precise nature of memory micro-processes involved in preference formation. Thus, we do not address relevant theoretical controversies within each of the areas of memory research on which we drew in this chapter to explain preference phenomena. In the case of priming, theoretical accounts range from spreading activation (J. R. Anderson, 1983; Collins & Loftus, 1975) to connectionist models (Ratcliff & McKoon, 1997). Interference effects (Postman & Underwood, 1973) have been explained in two ways, the first of which is known as the cue-overload principle (M. J. Watkins & Tulving, 1978; O. C. Watkins & Watkins, 1975), which is related to research on fan effects (J. R. Anderson & Reder, 1999; Reder & Anderson, 1980). This approach emphasizes that new information competes with old for associations in memory. More recent accounts describe interference as occurring from a process of selective retrieval with the explicit inhibition of competing material, sometimes called negative priming (M. C. Anderson & Neely, 1996; Milliken, Joordens, Merikle, & Seiffert, 1998). Later stages of a PAM research program may usefully revisit these distinctions, in light of data that may speak to one or more of them. Much of the memory literature that we reviewed and utilized falls into the growing category of implicit memory. While space constraints prevent us from reviewing the explicit/implicit memory distinction, it is the unconscious and automatic nature of the PAM processes that provides their appeal. Finally, we think that developments in cognitive neuroscience over the next couple of years will inform the PAM perspective and further theoretical developments and research in future years.

Affect and Memory

The PAM framework might, at first sight, appear to be overly cognitive, a "cold" model of decision making. However, this claim would be based on an artificial dichotomy between memory and affect, whereas in reality the two concepts are closely intertwined. First, affect determines what is recalled. Information consistent with basic or core emotions (Russell, 2002) is more available in memory (Forgas, 1995; Johnson & Tversky, 1983). Second, affect is often recalled or generated from memory. In addition to retrieving the factual event information,

cognition often delivers feelings, either through explicit or implicit retrieval or mental simulation, in what Russell (Russell, 2002) terms the virtual reality hypothesis. People have been shown to use both task-related and task-misattributed affect as information in when making judgments and choices (Loewenstein, Weber, Hsee, & Welch, 2001; Schwarz & Clore, 1983). Current work in our lab explores the relationship between affect induction and PAM framework memory processes in preference construction, to see whether recently reported affect-related moderation of endowment effects (Lerner, Small, & Loewenstein, 2004) can be explained as priming effects and affect-consistent recall, or whether non-memory processes need to be invoked. Complex preference and choice phenomena like loss aversion and time discounting may well have multiple, independent determinants. However, parsimony dictates that we should search for the smallest number of necessary explanatory processes.

Significance, Implications, and Applications

A better understanding of the cognitive process underlying preference construction serves several functions. Process explanations help to understand when predictions of utility are accurate predictors of experienced utility. Knowing why and how different preference elicitation methods lead to different predictions will help with the design of effective preference elicitation procedures, that is, interventions that produce "better" outcomes from either a societal or individual perspective. In addition, a memory-process orientation provides insights into how and why effects such as loss aversion, time discounting, and tradeoff difficulty may vary across people, domains, and decision context. Work in our lab currently examines developmental changes in preference (e.g., in the degree of loss aversion) as the result of known changes in susceptibility to interference with age (N. D. Anderson & Craik, 2000). A better understanding of the relationship between memory processes and preference may lead to the design of interventions that minimize socially harmful consequences of changes in memory performance on preference and choice (e.g., increased susceptibility to persuasion and advertising) in older populations.

Acknowledgments

This paper reports work supported by the National Science Foundation under Grant No. SES-0352062. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation. We thank NSF for their support.