

5

Knowledge Activation: Accessibility, Applicability, and Salience

E. TORY HIGGINS

Reprinted from *Social Psychology: Handbook of Basic Principles*, edited by E. Tory Higgins and Arie W. Kruglanski. Copyright 1996 by The Guilford Press, 72 Spring Street, New York, NY 10012.

When people identify, interpret, or, more generally, respond to a stimulus, what knowledge will be activated and used in their response? For half a century, social psychologists have addressed this fundamental issue in terms of two concepts—*accessibility* and *salience*. Krech and Crutchfield (1948), for example, described how a person's beliefs and attitudes vary in "salience," where a salient belief or attitude stands out as prominent in the cognitive field and enters thought more readily. They suggest that one might regard saliency as involving different thresholds. A salient belief or attitude would have conditions of activation that are easily satisfied—it would be "on the top of his mind" (p. 164). Thus, salient beliefs or attitudes would occupy the person's attention and be more frequently verbalized. As another example, Bruner (1957a, 1957b) proposed that the greater the "accessibility" of stored categorical knowledge, the more likely that it would be used to categorize stimulus information even when the stimulus information was impoverished, distantly related, or fit an alternative category better. As accessibility of the stored category increased, stimulus information would be more easily and swiftly recognized as a member of the category.

The concepts of salience and accessibility had a resurgence within social psychology in the 1970s. The proposed determinants of these psychological variables expanded during this resurgence. In Taylor and Fiske's (1978) review of salience, for example, salient information is that information which is most easily brought to mind and produces "top of the head phenomena." Determinants of salience include properties of stimuli, such as brightness and contrast, properties of situations, including instructional sets, and properties of perceivers, including both temporary need states and enduring cognitive and motivational differences. In another well-known review of salience, McArthur (1981) also includes both properties of stimuli and properties of perceivers, including their

expectations and arousal level, as determinants of salience. In Higgins and King's (1981) review of accessibility, accessible constructs are stored constructs that are readily used in information processing. Determinants of construct accessibility include momentary and chronic expectations, goals, and needs, recent and frequent prior activation of the construct, and the construct's interconnectedness with other stored constructs.

It is clear from even this brief review of the concepts of salience and accessibility that there has been considerable overlap in how social psychologists have used these concepts. There is nothing inherently wrong with using the terms salience and accessibility as synonyms for the same basic concept. But something would be lost by doing so. Specifically, treating these terms as synonyms precludes asking questions such as "What is the relation between accessibility and salience?" or "What are the independent contributions of accessibility and salience to knowledge activation?" Addressing such questions might be fruitful (see also Sherman, 1987). Moreover, given the increasingly frequent use of accessibility and salience as explanatory concepts in psychology, it might be helpful to clarify their unique contributions to knowledge activation.

The first section of this chapter proposes a set of distinctions for sharpening and defining key concepts and elucidating the role of different variables in knowledge activation. These distinctions will then be used in the subsequent sections to consider the nature and consequences of knowledge activation and use. The second section reviews when and how the accessibility of stored knowledge influences responses to stimuli. It also considers different models and proposals to account for accessibility effects. Knowledge activation depends not only on the accessibility of stored knowledge, but also on stimulus information and its relation to stored knowledge. The third section reviews the literature on applicability and salience, which concerns the contribution of stimulus information to

knowledge activation and use. The final section reviews the consequences of knowledge activation variables and considers the impact of their interrelations.

A SET OF DISTINCTIONS FOR KNOWLEDGE ACTIVATION AND USE

The purpose of this section is to propose a conceptual framework for considering how stored knowledge is activated and used when responding to stimulus information. The distinctions proposed here will be used in later sections to consider the nature and consequences of knowledge activation variables.

Availability and Accessibility

Knowledge cannot be activated or brought to mind unless it is present in memory. *Availability* refers to whether or not some particular knowledge is actually stored in memory. Tulving and Pearlstone (1966) distinguished between accessibility and availability when considering the difference between free recall and cued recall. Information that is not retrieved in free recall may nevertheless be retrieved with cued recall. Thus, information can be stored in memory, that is be "available," but not be easily retrieved or "accessible." (Tversky and Kahneman's, 1973, use of the term availability in their "availability heuristic" is different and will be discussed later.) Availability is a necessary condition for accessibility. If knowledge is not available, it has zero accessibility. If knowledge is available, it has at least some accessibility (Higgins & King, 1981). Distinguishing between availability and accessibility is especially useful when considering chronic differences among people, such as personality, developmental, or cultural differences, where both availability and accessibility differences are often involved (e.g., Higgins & Wells, 1986).

Accessibility can be defined as *the activation potential of available knowledge*. The term *potential* nicely captures several characteristics of accessibility (see *Merriam-Webster's Ninth New Collegiate Dictionary*, 1989, pp. 920-921). First, potential means being "capable of development into actuality" and "existing in possibility." This captures the characteristics of accessible knowledge that it is capable of being activated (and then used) but exists in a latent rather than in an active state. Second, the "potent" root of potential refers to "having influence," "bringing about a particular result." This captures the characteristic of accessibility that it, along with stimulus properties, is effective in increasing the likelihood that some stored knowledge will be activated. Third, potential includes notions of intensity, chemical effectiveness, electrical energy, and energy from the position of a piece of matter in an arrangement. As will be seen, these concepts cover the range of metaphors underlying current models of accessibility.

Knowledge Activation and Knowledge Use

As mentioned earlier, Bruner (1957b) defined the accessibility of stored categorical knowledge in terms of the likelihood that the knowledge would be used to categorize

stimulus information, and Higgins and King (1981) proposed that accessible constructs are stored constructs that are readily used in information processing. These definitions accurately reflect the general empirical relation between the accessibility of knowledge and the likelihood that the knowledge will be used in some way. The drawback is that these distinctions fail to distinguish conceptually between the activation and the use of knowledge. It is useful to distinguish between these concepts because there are variables that influence knowledge use beyond those involved in knowledge activation. It might be considered inappropriate to use activated knowledge, for example, when it was activated as part of an earlier task that is supposed to be kept separate from the current task (see Martin, 1986) or the activated knowledge is judged to be inappropriately stereotypic (see Gilbert & Hixon, 1991). It is preferable, therefore, to define accessibility in terms of potential for knowledge *activation* rather than potential for knowledge *use*. Moreover, although there are alternative models of accessibility (described later), conceptualizing accessibility in terms of activation potential is compatible with all current models.

Increased Accessibility from Priming and Priming Effects

There has been a tendency in the literature to confuse priming effects with the effects of increased accessibility from priming. Cramer (1968) defines *priming* as "a change in antecedent conditions which is specifically designed to increase the probability of a particular response being given to a particular stimulus" (p. 82). Tulving (1983) defines priming more broadly as "the facilitative effect of performing one task on the subsequent performance of the same or similar task" (p. 100). Tulving (1983) notes that priming has also been used to describe the underlying processes responsible for priming effects. In addition, priming has been used to refer to a procedure involving priming events that stimulate or activate some stored knowledge. It has been found that such events (under conditions to be discussed later) can *decrease*, as well as increase, the probability of a particular response being given to a particular stimulus (e.g., Lombardi, Higgins, & Bargh, 1987; Martin, 1986). In this chapter, the term priming refers to procedures that stimulate or activate some stored knowledge.

Priming can be accomplished in a variety of ways with different effects. One of its effects is to increase accessibility. The priming event activates a stored knowledge unit, which temporarily increases the accessibility of the unit. The increased accessibility, in turn, increases the likelihood that this knowledge unit will be activated by subsequent stimulus information. Some consequences of priming procedures can be attributed to increased accessibility. But not all, because priming procedures can do more than just increase accessibility. The influence of priming on subsequent responses, therefore, might not be due simply to increased accessibility. A priming procedure, for example, could induce an active response set that influences subsequent judgments independent of any accessibility effect per se (e.g., Ferguson & Wells, 1980). Other priming effects as well need not be attributed to accessibility per

se, especially if one accepts Tulving's (1983) broad definition of priming mentioned earlier. By this definition, priming would be involved when a subject produces a response in order to be consistent with an earlier response, which could involve a motivational consistency effect rather than an accessibility effect.

There has been confusion about priming effects especially for evaluative impressions of target persons. When priming a trait construct increases the accessibility of that construct, it increases the likelihood that the construct will be used later to identify a target person's behavior to which it is applicable. The *connotative implications of this identification* might then be used to form an evaluative impression of the target (see Sinclair, Mark, & Shotland, 1987). Although it would be appropriate, given the preceding definitions, to refer to the influence of priming on the evaluative impression as a "priming effect," it is misleading to claim that this priming effect was an accessibility effect per se. The priming effect on initial identification would be an accessibility effect, but the priming effect on the subsequent evaluative impression would be an identification-based inference effect.

Accessibility and Applicability

Two basic variables influence the likelihood that some stored knowledge will be activated—the accessibility of the stored knowledge prior to stimulus presentation and the *fit* between the stored knowledge and the presented stimulus (Bruner, 1957b). Like *fit*, *applicability* refers to the relation between the features of some stored knowledge and the attended features of a stimulus (where the features are typically categorical in nature; see Higgins & Chaires, 1980). The greater is the overlap between the features of some stored knowledge and the attended features of a stimulus, the greater is the applicability of the knowledge to the stimulus and the greater is the likelihood that the knowledge will be activated in the presence of the stimulus (see Higgins, 1989a).

It should be emphasized that applicability concerns the "attended features" of a stimulus. As Medin (1989) has noted, the concept of "similarity" alone is not sufficient because not all features of a stimulus are equally relevant or important. And different features can be relevant or important at different times. Various factors can influence which features are attended to at any point in time (see also Smith & Zarate, 1992, for a similar point regarding social perception). Some factors have to do with the perceiver's characteristics, such as his or her needs or goals (see Martin & Achee, 1992). Other factors have to do with situational characteristics, such as the alternative stimuli present. As discussed later, the current function of the stimulus in relation to the knowledge is also important because it determines whether only matches between the stimulus and the knowledge are relevant or both matches and mismatches are relevant.

Applicability and Salience

The concept of salience could be used to capture the notion that not all features of a stimulus receive equal attention at any point in time. By relating salience to applicability in

this way, the independent contributions of accessibility and salience to knowledge activation could be examined. Moreover, conceptualizing salience this way would not be a radical departure from the literature. Although the concepts of salience and accessibility have been confused in the literature, as noted earlier, salience has more often been thought of as having something to do with stimulus information than as having something to do with stored knowledge.

The salient features of a stimulus event are those features that receive selective attention, those features that draw, grab, or hold attention relative to alternative features (see McArthur, 1981; Taylor & Fiske, 1978). This property of salience concerns focal attention—selecting one object or set of features in a stimulus event for examination, rather than alternatives, prior to knowing what that object or set of features might represent (see Neisser, 1967). Again, the features are typically categorical. This aspect of salience reflects the everyday meaning of salient as "standing out conspicuously; prominent; striking" (Merriam-Webster, 1989, p. 1038).

There is an implicit aspect of this conceptualization of salience that needs to be made explicit: Salience refers to something about the stimulus event that does not occur until exposure to the stimulus, and that occurs without a prior set for a particular kind of stimulus, such as a belief about or search for membership in a particular category. Thus, salience is conceptualized here as *something about a stimulus event that occurs on exposure, without a prior set for a particular kind of stimulus, that draws attention selectively to a specific aspect of the event*. As mentioned earlier, previous perspectives on salience have included any variable that can influence selective attention as part of the phenomenon of "salience." Again, there is nothing inherently wrong with such an approach. The drawback is that the distinct contribution of salience to knowledge activation can be lost. For example, the distinction between salience and accessibility could be lost since previous reviews have included accessibility as one determinant of selective attention (see McArthur, 1981; Taylor & Fiske, 1978). Equally important, perceivers' needs, goals, and expectancies have also been included as part of the phenomenon of salience, as noted earlier. Because needs, goals, and expectancies can determine knowledge activation before the stimulus is even presented, the distinct contribution of salience to knowledge activation becomes unclear.

To clarify the distinct contribution of salience to knowledge activation, it is necessary to restrict the concept of salience to properties of the stimulus event only, independent of any prior set for a particular kind of stimulus. Properties of the perceiver, for example, whether goals, expectancies or accessibility, must be excluded. Properties of the stimulus event include properties of particular objects in the stimulus event, properties of the immediate context in which the objects appear, and the relations among these properties. Restricting the concept of salience to properties of the stimulus event means that other previously included variables would be excluded from the concept. For example, both statistical novelty as a property of stimuli and instructional set as a property of situations have been included as conditions of salience (e.g., Taylor & Fiske, 1978). But both of these

properties can have a direct influence on knowledge activation prior to and independent of stimulus exposure. Thus, such properties need to be excluded from the concept of salience.

In contrast, the property of “comparative distinctiveness” would be included in the concept of salience. In a typical salience experiment, for example, subjects are presented with an array of objects in which one object has a characteristic that differs from most or all of the other objects. This “comparative distinctiveness” draws attention selectively to the object, thus making it salient in the immediate context.

According to this conceptualization, salience is one distinct determinant of selective attention. Instructional sets, momentary goals, active expectancies, and so on, are additional distinct determinants of selective attention. Thus, salience involves selective attention but is not the same as selective attention. Selective attention can occur for reasons other than salience. Salience is a source of selective attention that arises from properties of the stimulus event. Because of its impact on how features of the stimulus event are attended, salience influences applicability and plays an important role in knowledge activation (to be described later).

Applicability and Judged Usability

It is also important to distinguish between applicability, or the “goodness of fit” between some stored knowledge and the attended features of a stimulus (Bruner, 1957b), and *judged usability*, which is the judged appropriateness or relevance of applying stored knowledge to a stimulus (see Higgins, 1989a). (In this chapter, I have replaced the term “judged usability” for the term “perceived applicability” that I, and others, have used previously because the latter term is too easily confused with “applicability” per se and misleadingly suggests a process of perception rather than the actual judgmental process involved.) There is substantial evidence that perceivers’ judgment of the relevance or appropriateness of particular information for their response can determine their use of the information (see, e.g., Bar-Hillel, 1980; Higgins & Bargh, 1987; Kruglanski, Friedland, & Farkash, 1984; Sherman & Corty, 1984). Variables that influence judged usability include apparent causal significance (e.g., Ajzen, 1977), framing of a task or problem (e.g., Trope & Ginossar, 1988; Zukier, 1986), conversational norms (e.g., Strack, Martin, & Schwarz, 1988), and perceived reliability and representativeness (e.g., Kassin, 1979; Zukier & Pepitone, 1984).

Judged usability concerns both the judged relevance of some knowledge for a current task, such as using an activated exemplar as a standard of comparison, and the judged appropriateness of some knowledge for directly assigning meaning to a stimulus, such as using an activated concept to categorize a stimulus. As mentioned earlier, judged usability is important because it influences whether activated knowledge will actually be used (see Martin & Achee, 1992). Even when stored information is activated because of its accessibility and applicability to a stimulus, it might not be consciously used if it is perceived as irrelevant or inappropriate. During a criminal trial, for

example, jurists will attempt not to use evidence in their deliberations that has been ruled inadmissible by the trial judge. Nonprejudiced persons will attempt not to use stereotypic information in judging others when the momentary situation facilitates usability judgments (e.g., Devine, 1989). It has been suggested that information perceived as inappropriate is likely to be suppressed and that this can produce contrast effects of priming (see Martin, 1986). Rather than being an accessibility or applicability effect per se, such contrast would be an effect of judged usability.

Accessibility Experiences

People also have subjective experiences of accessibility and knowledge activation, such as whether it was easy or difficult. In addition, people have theories about these subjective experiences and draw inferences from them; that is, they have metacognitions about *accessibility experiences* (see Schwarz, this volume; Strack, 1992). The now classic “availability heuristic” of Tversky and Kahneman (1973; see also Kahneman & Tversky, 1973) is one such metacognition. Jacoby’s “becoming famous overnight” phenomenon provides another example (see Jacoby & Kelley, 1990). The more easily instances of some knowledge class can be activated or brought to mind (i.e., the higher the availability of class instances), the higher the judged frequency or likelihood of that class. This increase in presumed base rate of the class might, in turn, increase the likelihood that the stimulus will be judged as an instance of the class. Thus, as with judged usability, metacognitions about accessibility experiences influence judgments of stimuli by influencing how activated knowledge is used rather than by influencing the likelihood of activating that knowledge in the first place.

Stimuli as Stimulants and Stimuli as Targets

The literature on accessibility and salience has described various determinants of the relation between stored knowledge and stimulus information. In considering these determinants, it would be useful to distinguish two different directions of this relation—the direction *from stored knowledge to stimulus information* versus the direction *from stimulus information to stored knowledge*. Stimulus information serves a different function depending on which direction is involved.

According to Bruner (1957b), accessibility involves “perceptual readiness” that reflects both goal pursuit and the learned probabilities of the occurrence of events in everyday life. The expectancies associated with the learned probabilities of occurrence involve hypotheses about what to look for, and stimulus information is checked to test the hypotheses (see Bruner, 1951). The goals are predispositions that produce a readiness to respond to certain types of goal-relevant stimuli. The direction in these cases is from stored knowledge to stimulus information. The stimulus information serves *two different target functions*: (1) stimulus information as *evidence* in the case of expectancies or hypothesis testing; and (2) stimulus information as *need fulfillment* in the case of goal pursuit.

Bruner's (1957b) perceptual readiness perspective on how accessible knowledge relates to stimulus information can be contrasted with the perspective on accessibility offered by the *synapse* model of Higgins and his colleagues (see Higgins, 1989a; Higgins & Brendl, 1995). According to the *synapse* model (to be described more fully later), the process begins with the stimulus, which increases the excitation level of stored knowledge as a function of the match between the attended features of the stimulus and the features of the stored knowledge. If the excitation level of the stored knowledge before stimulus exposure is sufficiently high (i.e., high accessibility), then the additional excitation provided by the attended features of the stimulus will activate the knowledge. The direction in this case is from stimulus information to stored knowledge, and the stimulus information serves a *stimulant function*.

Thus, depending on the direction of the relation between stored knowledge and stimulus information that is emphasized, stimulus information can function as a *target* or as a *stimulant*. The relevance of different features of the stimulus information varies depending on whether it functions as a target or as a stimulant. When stimulus information functions as a stimulant, only those features that match stored knowledge are relevant. In addition, the stimulus information will increase the excitation level of all stored knowledge units whose features are similar to the features of the stimulus information. In contrast, when stimulus information functions as a target, both features that match and features that mismatch stored knowledge are relevant. Moreover, only one knowledge unit might be considered at a time. For example, perceivers expecting a particular event might assess the extent to which a stimulus confirms or disconfirms that expectancy.

When stimulus information functions as a target, the stored knowledge might already be activated prior to stimulus exposure. A person might expect or seek out a particular object or event, for example. Because stored knowledge is *already activated* in such cases, any effect of the knowledge on responding to a stimulus would *not* be an accessibility effect given that accessibility refers to activation *potential*. As Bruner (1957b) suggested, however, such effects do involve perceptual readiness. Thus, although accessibility and perceptual readiness have been treated almost as synonyms in the literature, they are not the same variable and it would be useful to distinguish between them. For example, the effect on subsequent responding of an active set created by experimental instructions, such as the instruction to expect a "warm" or a "cold" instructor in Kelley's (1950) classic study of labeling effects, is more appropriately described in terms of perceptual readiness than accessibility. In contrast, the effect of subliminal priming on subsequent stimulus categorization is clearly an accessibility effect (e.g., Bargh & Pietromonaco, 1982).

Salience was conceptualized as something about the stimulus event that occurs on exposure without a prior set for a particular kind of stimulus. When stimulus information functions as a stimulant, perceivers are not yet set for a particular kind of stimulus, and thus salience is a likely source of selective attention. When stimulus information functions as a target, it is less clear whether selective

attention is due to salience or to expectations, for example, because perceivers might already be set for a particular kind of stimulus even before it occurs. Thus salience is a more likely source of selective attention and knowledge activation when stimulus information functions as a stimulant than when it functions as a target.

Although a stimulus can function as a target or as a stimulant at any one time, the same stimulus can function as both a target and as a stimulant over time. In a controlled process of categorization, for example, a stimulus might first function as a stimulant and activate a construct that could be used to categorize the stimulus. Then, prior to selecting this construct, the stimulus might function as a target of the hypothesis that the stimulus is a category member. Depending on the extent of hypothesis-confirming and hypothesis-disconfirming features, the stimulus will or will not be categorized in terms of the construct. It should also be noted that some features of a stimulus event might function as a stimulant while other features are functioning as a target, perhaps because a stimulus has features that are not represented in the target-related goal or expectancy.

The proposed distinction between "stimuli as targets" and "stimuli as stimulants" differs from the classic distinction between "top-down" and "bottom-up" processing (see Bobrow & Norman, 1975). Top-down processing is driven by motives and goals and fits input to expectations. Bottom-up processing finds knowledge units in which to embed the input. In this sense, processing stimuli that function as targets could be relatively top-down or theory-driven, or it could be relatively bottom-up or data driven. A perceiver, for example, could either be motivated to confirm a prior belief by fitting input to the belief (more theory-driven) or be motivated to use input as evidence in testing alternative hypotheses about the input's category membership (more data-driven). Likewise, processing stimuli that function as stimulants could be more top-down or bottom-up. For example, only minimal stimulation might be needed from a stimulus to produce construct activation because some construct has such high prior accessibility (more theory-driven) or a highly salient and unambiguous stimulus might itself produce such strong stimulation that a construct need only be available to be activated (more data-driven).

This first section has proposed a set of distinctions with the aim of clarifying the unique contribution of different variables to knowledge activation and use. These distinctions will provide the basis for the subsequent review of accessibility, applicability, and other variables influencing knowledge activation and use. Social psychological research and ideas will be reviewed as they relate to the variables *as defined in this section*, regardless of how they might have been labeled or characterized in the literature. This is necessary if processes and principles are to take precedence over historical labels.

The next two sections will review the conditions of knowledge activation and use (i.e., the "when" and the "how"), and will discuss different conceptualizations of knowledge activation variables and their effects. These sections will be followed by a section on the consequences of these variables and their interrelations. Consequences

for attention, memory, feelings, judgment, and action will be considered that have significant implications for a wide variety of social phenomena.

ACCESSIBILITY

This review of the principle of accessibility begins with a survey of the basic findings in the social-personality literature on when and how accessibility influences responses to stimuli; that is, the conditions for accessibility effects. Different models and proposals to account for these accessibility effects are then considered.

When does knowledge accessibility influence people's responses to stimuli and how does it influence the responses? Although research relevant to these "second generation" questions (Zanna & Fazio, 1982) only began relatively recently, some basic conditions for accessibility effects have already been identified.

Accessibility Effects from Priming

Before considering the conditions of accessibility effects from priming, the investigator must establish that there *are* accessibility effects. Thus, the initial question was, "Does prior activation of stored knowledge increase the probability of using that knowledge when responding to a subsequent stimulus?" An initial study by Higgins, Rholes, and Jones (1977) addressed this question for the case of people categorizing stimulus behaviors.

Effects from Recent Priming

Subjects in the Higgins et al. (1977) study were initially exposed to one or another set of trait-related constructs as an incidental aspect of a study on perception. The subjects later participated in a supposedly unrelated study on "reading comprehension," where they were asked to characterize the ambiguous behaviors of a target person. For example, one ambiguous behavior that they read about was as follows: "Once Donald made up his mind to do something it was as good as done no matter how long it might take or how difficult the going might be. Only rarely did he change his mind, even when it might well have been better if he had." Pilot testing had indicated that this behavior was equally likely to be characterized as "stubborn" or as "persistent." In the earlier "perception" task, half of the subjects had been exposed to the word "stubborn" as an incidental part of a Stroop task (i.e., a supposed memory "load" to make the task more difficult) and the other half of the subjects had been exposed to the word "persistent." The study found that the subjects were significantly more likely to use the trait-related constructs primed in the initial perception task to categorize the target person's behaviors in the second "reading comprehension" task than to use the equally applicable alternative constructs.

This basic phenomenon of activating stored knowledge and thereby increasing the likelihood that this knowledge will influence subsequent responses to stimuli has been replicated many times using a wide variety of priming methods, types of responses, and types of stimuli (e.g.,

Bargh, Bond, Lombardi, & Tota, 1986; Carver, Gannellen, Froming, & Chambers, 1983; Erdley & D'Agostino, 1988; Fazio, Powell, & Herr, 1983; Herr, Sherman, & Fazio, 1983; Higgins, Bargh, & Lombardi, 1985; Higgins & Chaires, 1980; James, 1986; Martin, 1986; Rholes & Pryor, 1982; Sedikides, 1990; Sherman, Mackie, & Driscoll, 1990; Sinclair et al., 1987; Skowronski, Carlston, & Isham, 1993; Smith & Branscombe, 1987; Srull & Wyer, 1979; Trzebinski & Richards, 1986; Williams, 1993; Wyer & Hartwick, 1980). These studies have provided a clear affirmative answer to the first generation question, "Is there a phenomenon?" They have also suggested an initial "When?" condition for accessibility effects.

In the Higgins et al. (1977) study, as in many subsequent studies, the activated constructs were common trait-related concepts. It is likely that all subjects had each of the alternative constructs available to them. Given that these constructs were available, they had at least some activation potential. In addition, each behavior could be categorized by the alternative constructs. To reveal an accessibility effect, therefore, it is necessary that one construct be *more* accessible than an alternative construct that is equally applicable to the stimulus. These studies demonstrate that a construct can be made more accessible than alternatives by priming that construct prior to stimulus exposure—*recent priming* is one condition for accessibility effects.

Delay from Priming to Stimulus Exposure

Recent priming increases the likelihood that the primed construct will be used to judge a subsequent stimulus. But how recent must the "recent" priming be? More generally, what is the effect of the delay from priming to stimulus exposure on the strength of accessibility effects? (The effect of delay from the initial encoding of a stimulus to subsequent stimulus-related responses is an encoding effect rather than an accessibility effect of priming, and thus will not be reviewed.)

Srull and Wyer (1979) examined this issue in a couple of studies that used a "scrambled sentences" technique to prime trait-related constructs. Subjects were given a scrambled sentence and were asked to underline three of the words that would make a complete sentence. Each possible sentence was related to the target trait, such as the scrambled sentence, "leg break arm his," for the trait "hostility." Following this priming task, the subjects participated in another supposedly unrelated "impression formation" task in which they read a set of moderately hostile behaviors. The subjects read the essay either about 5 minutes after the priming task (i.e., the time taken by the new experimenter to introduce the impression formation task), one hour after, or 24 hours after.

The effect of priming on trait-related judgments of the behaviors significantly decreased over time. Indeed, when the priming was the weakest (i.e., 6 hostility-related primes among 30 scrambled sentences), there was little or no effect of priming on behavioral judgments once the delay was an hour or more. This basic effect of priming-to-stimulus delay was replicated by Srull and Wyer (1980), and the results of other studies support the conclusion

that the effect of recent priming tends to decrease as the delay from priming to stimulus exposure increases (e.g., Higgins et al., 1985; Higgins & Brendl, 1995). These more recent studies have also found that when there is a competing accessible construct, the effect of recent priming on judgment can disappear in a couple of minutes or less.

There are also cases, however, when priming effects can remain strong for a very long time. When a knowledge unit, such as an ACT^{*}-type production, is frequently activated and used in categorizing a specific behavior, then the effect of priming on later judgments of that same behavior can last several days. Smith, Stewart, and Buttram (1992), for example, found that judgments of specific behaviors in relation to the same trait were quicker when the behaviors had been judged before than when they were new even when the repetition was separated by a 7-day delay. Moreover, prior experience categorizing an "ambiguous" behavior in relation to a particular trait made it more likely that the same behavior would be evaluated 7 days later according to the desirability of the previously primed trait. It is also possible that a single prime would cause people to think about something for a long time afterward, especially if a personal concern was primed (see Schwarz & Strack, 1981). Thus, a single prime could produce frequent activation of the prime-related construct, which would yield a priming effect on construct-related responses even after a long delay.

Effects from Frequent Priming

A stored knowledge unit can be primed only once or many times prior to stimulus presentation. It has been recognized for a long time that the frequency of prior activation is another important variable of knowledge use. In the early nineteenth century, for example, Thomas Brown proposed both recency and frequency as basic conditions under which the general principle of association operates (see Heidbreder, 1933). And in the late nineteenth century, Edward Lee Thorndike (1898) proposed in his law of exercise that the more recently and the more frequently a stimulus-response bond is exercised, the more effectively it is stamped in.

Srull and Wyer (1979) examined the relation between frequency of priming and the strength of accessibility effects on behavioral judgment. Using their "scrambled sentences" technique, they primed the target construct (e.g., "hostile") either 48 times, 24 times, 12 times, or 6 times. The priming effect on behavioral judgments increased as the frequency of priming increased, especially when the stimulus information was ambiguous and there was little priming-to-stimulus delay. This basic effect of increasing the frequency of priming was also replicated by Srull and Wyer (1980). The Srull and Wyer (1979, 1980) studies also showed that with highly frequent priming (i.e., 35 or 48 primes), there could be an effect of priming on judging "ambiguous" behaviors even when there was a priming-to-stimulus delay of 24 hours.

A conceptual replication of Srull and Wyer's (1979, 1980) studies was conducted by Bargh and Pietromonaco (1982). In their study, however, the priming events were presented subliminally, outside the subjects' awareness.

The subjects, therefore, were not only unaware of the relation between the priming events and the subsequent judgment task as in previous studies but were unaware of even the presence of the priming events. In their study as well, the strength of subjects' prime-related judgments increased as the frequency of priming increased. Using subliminal priming like Bargh and Pietromonaco (1982), Devine (1989) also found that when constructs stereotypically related to the social category of African Americans were primed more frequently, subjects' subsequently judged a target person's vaguely hostile behaviors as more "hostile."

Effects from Recent Priming versus Frequent Priming

Most priming studies have primed only one construct per subject that could be used to judge the stimulus information. Higgins et al. (1985), however, primed two alternative constructs that were equally applicable to an ambiguous behavior. In this way, the effects of recent and frequent priming could be pitted against each other. Using a modified "scrambled sentences" paradigm in which the sentences included trait terms rather than behaviors, one of the alternative constructs was primed more frequently (four times), and the other alternative construct was primed only once but most recently at the very end of the priming task. Following the final prime, there was either a 15-second or 120-second delay (filled by a "counting backward" interference task) before subjects read the ambiguous target behaviors.

After the shorter delay, the subjects tended to judge each behavior using the more recently primed construct. But after the longer delay, the subjects tended to judge each behavior using the more frequently primed construct. This basic pattern of results was replicated by Lombardi et al. (1987).

Chronic Accessibility Effects

Many psychologists have proposed that there are individual differences, including cross-cultural differences, in the "theories" or viewpoints people possess about human nature and personality (e.g., Kelly, 1955; Sarbin, Taft, & Bailey, 1960; Tagiuri, 1969; Tajfel, 1969). Some of these differences concern the availability of particular constructs. But chronic differences in the accessibility of constructs have interested personality and social psychologists as well.

Kelly (1955), for example, seemed to have differences in chronic accessibility in mind when he wrote, "Construct systems can be considered as a kind of scanning pattern which a person continually projects upon his world. As he sweeps back and forth across his perceptual field he picks up blips of meaning" (p. 145). The basic notion behind the Thematic Apperception Test, developed by Murray (1938) and used extensively to measure need achievement (e.g., McClelland & Atkinson, 1948), was that thoughts related to a strong motive should be chronically accessible (see Sorrentino & Higgins, 1986). Research inspired by the "New Look" in perception examined how chronically

accessible values influenced responses to stimulus information, such as subjects with strong religious versus economic values perceiving tachistoscopically presented pictures differently (Bruner, 1951; cf. Cacioppo, Petty, & Sidera, 1982). The defense mechanism of "projection" described in the psychodynamic literature includes the notion of heightened responding to others in terms of chronically accessible thoughts or feelings (see Cameron, 1963).

Chronic versus Nonchronic

How individual differences in chronic accessibility influence the processing of behavioral information was directly examined by Higgins, King, and Mavin (1982). In one study, subjects' chronically accessible constructs were measured by asking them to list the traits of a type of person that they liked, disliked, sought out, avoided, and frequently encountered. Chronic accessibility was defined in terms of output primacy. For a given trait-related construct, individuals would be "chronic" on a construct if they listed that construct first in response to one or more questions, and they would be "nonchronic" on a construct if they did not list that construct in response to any question.

About one week later, subjects participated in a supposedly unrelated study on "psycholinguistics" conducted by a different experimenter. Each subject read an individually tailored essay containing trait-related descriptions of a target person. Each subject was chronic for half of the traits and nonchronic for the other half of the traits. Using a "quasi-yoking" design to control for the content of the trait-related descriptions, some subjects were chronics and some subjects were nonchronics for each of the traits described in the essays.

On both a measure of subjects' spontaneous impressions of the target person and a measure of their recall of the behavioral descriptions, the study found that subjects were significantly more likely to include information related to traits on which they were chronic than information related to traits on which they were nonchronic. These basic effects were also found in another study by Higgins et al. (1982) in which chronic accessibility was defined in terms of the frequency of output to different questions rather than the primacy of output. This study also found that the effects of chronic accessibility on impressions and memory can last for at least a couple of weeks. Moreover, using an alternative operationalization of chronic accessibility, Lau (1989) has found that chronically accessible constructs are relatively stable even over years and that they guide the processing of information about a wide variety of political objects.

Bargh and Pratto (1986) used the Stroop color-naming paradigm to test directly whether the activation potential of constructs on which individuals are chronic is higher than constructs on which individuals are nonchronics. Subjects were classified as chronic or nonchronic on different constructs according to the Higgins et al. (1982) criteria. Bargh and Pratto (1986) found that subjects' chronically accessible constructs were indeed at a higher level of activation readiness than their inaccessible constructs, as revealed by greater interference on the color-

naming task from the chronically accessible constructs. (For additional evidence of individual differences in the chronic accessibility of constructs, in this case interpersonal constructs, see Baldwin, Fehr, Keedian, Seidel, & Thomson, 1993.)

Bargh and Thein (1985) tested whether chronic accessibility is associated with automatic processing. Again, subjects were classified as chronic or nonchronic on different constructs according to the Higgins et al. (1982) criteria. For constructs on which the subjects were chronic, Bargh and Thein (1985) found evidence of the processing efficiency expected with automatic processing for both subjects' impressions and recall of a target person's behaviors. No such evidence of processing efficiency was found for constructs on which the subjects were nonchronic.

Although not originally conceptualized in terms of individual differences in chronic accessibility, research comparing individuals for whom a construct is central (*schematics*) or is not central (*aschematics*) to their self-description can be interpreted as providing additional evidence for the role of chronic accessibility in social information processing. Markus (1977) classified subjects as being independent schematics, for example, if they responded at the high end of both a self-descriptiveness and an importance scale for at least two of three independence-related traits, and classified subjects as being independent aschematics if they responded in the middle of these scales. According to such criteria, schematics probably differ from aschematics in a number of ways (see Higgins & Bargh, 1987). Still, it is likely that independent schematics, for example, are chronic on the construct "independent" whereas aschematics are nonchronic.

Markus (1977) found that independent schematics (i.e., independent chronics) processed independent-related stimulus information faster and more consistently than did independent aschematics (i.e., independent nonchronics). Bargh (1982) also compared the processing of independent schematics and aschematics, classified according to Markus's (1977) criteria. Subjects were given a dichotic listening task in which independent-related words either appeared in the to-be-attended or the to-be-ignored channel. A probe reaction time was used to measure the amount of resources allocated to the listening task. Bargh (1982) found that compared with independent aschematics (i.e., independent nonchronics), independent schematics (i.e., independent chronics) used less resources when independent-related words appeared in the to-be-attended channel but used more resources when these words appeared in the to-be-ignored channel. These results support the notion that chronic accessibility is associated with automatic processing.

Chronicity as a Continuous Variable

The initial studies on the effects of chronic accessibility on judging and remembering social behavior compared the effects of individuals who were chronic versus nonchronic on the target constructs. In a couple of recent studies, Higgins and Brendl (1995) have investigated the effects of chronic accessibility as a continuous variable with multiple

levels of chronicity. Using the Higgins et al. (1982) measure of chronic accessibility, individuals were again classified as nonchronic on a construct if they did not list that construct in response to any question. An operationalization of continuous levels of chronicity was introduced that involved a frequency of output score. This score was obtained by counting the total number of construct-related responses (in this case, "conceited" and its synonyms) to all five questions and dividing this number by the total number of all responses to these questions. Using the full range of subjects' chronicity scores for "conceited," obtained weeks before the experiment, Higgins and Brendl (1995) found a strong positive relation between higher chronicity scores and stronger impressions that the vague or ambiguous behaviors of a target person reflected conceitedness.

Lau (1989) used a couple of similar measures of levels of chronicity to examine the relation between chronic accessibility for political constructs and evaluative responses to political candidates. One study investigated subjects' chronically accessible constructs for political candidates by measuring how frequently each subject mentioned person-related traits when answering open-ended questions about the good and bad points of the two major parties and their candidates. Lau (1989) found that subjects' preselection candidate opinions were more strongly related to their choice of candidate when the chronic accessibility of their candidate constructs was higher. Lau (1989) also provided evidence that the chronic accessibility of political constructs remains quite stable even over a period of four years.

Priming and Chronicity

It is evident from the preceding review that greater temporary accessibility as a function of priming and greater chronic accessibility as an individual difference both predict stronger responses to stimulus information in terms of the accessible construct. What happens when both of these sources of accessibility are present? Bargh et al. (1986) addressed this question by selecting subjects who were chronic or nonchronic on the construct "kind" weeks before the experiment, and then priming or not priming (subliminally) that construct a few minutes before the subjects' judged "ambiguously kind" behavioral descriptions. Bargh et al. (1986) found that subjects' construct-related impressions of the target person were stronger both when the subjects were chronic on the construct and when the construct was primed (see also Graham & Hudley, 1994). Most important, there was a reliable effect of chronic accessibility even within the priming condition. In one of their studies, Higgins and Brendl (1995) also found that subjects with higher levels of chronic accessibility for "conceited" made stronger judgments that ambiguous behavior was conceited. Moreover, as in the Bargh et al. (1986) study, this effect did not interact with either priming (or delay) and it was reliable even within the priming condition.

The chronicity effect was (nonsignificantly) stronger within the priming condition than within the no-priming condition in both the Bargh et al. (1986) and Higgins and

Brendl (1995) studies. In addition, Higgins and Brendl (1995) also found that subjects judged extremely vague behavior as conceited only when their chronic accessibility for conceited was high, conceited had been primed, there was a short priming-to-stimulus delay, and they did not remember the priming events (a variable to be discussed later). Thus, chronicity and priming can have independent and interactive effects on judgments.

Recent Priming versus Chronicity

Just as recent priming can be pitted against frequent priming, recent priming can be pitted against chronic accessibility. Using the Higgins et al. (1982) criteria for classifying chronicity, Bargh, Lombardi, and Higgins (1988) selected subjects who were either chronic on "inconsiderate," chronic on "outgoing," or nonchronic on both inconsiderate and outgoing. Using the Higgins et al. (1985) scrambled sentences task, the subjects had either the construct inconsiderate primed or the construct outgoing primed in an experimental session held weeks after chronicity was measured. For the chronic subjects, the construct that was primed was the construct on which they were nonchronic (e.g., "inconsiderate" chronics were primed with "outgoing"). After either a 15-second, a 120-second, or a 180-second delay following the last priming trial, subjects judged what type of person was exemplified in a behavioral description ambiguous for inconsiderate and outgoing.

Bargh et al. (1988) found that nonchronics' judgments did not change as a function of the priming-to-stimulus delay. In contrast, chronics' judgments reflected the primed construct slightly more than the chronic construct after the 15-second delay, but the reverse was true after the 120-second delay. Thus, the pattern of judgments over time for recent priming versus chronicity closely resembled the pattern of judgments over time for recent priming versus frequent priming found by Higgins et al. (1985).

Accessibility Experiences

As defined earlier, accessibility is the activation potential of available knowledge, the likelihood that stored knowledge will be activated. Such likelihood of activation is an objective variable. But accessibility and activation also have a subjective component. People can experience the retrieval and activation of stored knowledge as being easy or difficult, for example. These accessibility experiences also influence people's responses.

Availability Heuristic

As noted earlier, knowledge cannot be accessible if it is not even available in memory. Such events are not just difficult to activate, they are impossible to activate. Thus, the availability of episodes, in the sense of being stored in memory or not, should influence frequency estimates. The news media, for example, are more likely to report instances where the cause of death was an accident, which is newsworthy, than instances where the cause of death was a stroke. Although there are many more deaths by stroke than by accident, subjects in a study by Lichtenstein,

Slovic, Fischhoff, Layman, and Combs (1978) judged death by accident as more likely than death by stroke.

Accessibility can also influence likelihood estimates. Tversky and Kahneman (1973; see also Kahneman & Tversky, 1973) used the term "availability" in a different way than used here to refer to accessibility experiences. They proposed that one way in which people infer the frequency or likelihood of a type of event is by their experience of how quickly or easily instances of that type of event can be activated or brought to mind—the "availability heuristic." It has been suggested, for example, that words are organized in long-term memory by initial consonant. Given this, it would be easier to retrieve instances of words that begin with the letter "K" than words that have "K" as the third letter, even though the former words are actually less frequent than the latter. Consistent with the prediction of the availability heuristic, Tversky and Kahneman (1973) found that subjects judged the former words to be more frequent than the latter.

Gabrielcik and Fazio (1984) used a subliminal priming paradigm to test more directly whether such word frequency estimates are really due to the subjective experience of increased accessibility. A set of target words were tachistoscopically flashed at a rate too fast to be accurately recognized. In the priming condition all of the words contained the letter "T." In the no-priming, control condition the words were replaced with strings of asterisks. When later asked whether more words in the English language contain the letter "T" compared with another letter (e.g., "D" or "S"), the primed subjects judged "T" to occur more frequently than did the control subjects.

According to the availability heuristic, it is the subjective experience of easily activating instances of a category that produces judgments of category instances being frequent or a category being common. In several studies of the availability heuristic, however, the judgments could have been based on the actual information that was activated rather than being based on accessibility experiences. Even in the Gabrielcik and Fazio (1984) study, for example, the priming manipulation could have resulted in more instances being activated of one word type than the alternative word type at the time of the frequency estimation (independent of any episodic memory for the priming events themselves). Thus, the subjects might have based their judgments on the relative frequency of each word type that was available in working memory. The biased judgments, then, could be due to priming producing actual differences in knowledge accessibility that yielded different sample sizes, which would be a direct accessibility effect, rather than priming producing different accessibility experiences underlying the use of the availability heuristic.

Schwarz, Bless et al. (1991) introduced a methodology that more directly tested for judgmental effects of accessibility experiences independent of direct accessibility effects. In a study supposedly on "autobiographical memory," the subjects were asked either to describe situations in which they had behaved very assertively and felt at ease or to describe situations in which they had behaved unassertively and felt insecure. In each of these conditions, half of the subjects were told to describe 6 examples

and the other half were told to describe 12 examples. Pretesting indicated that generating 6 examples of a type of behavior was experienced as an easier task than generating 12 examples, and this was confirmed on an ease of retrieval manipulation check taken at the end of the experiment. The subjects later judged their own assertiveness. Consistent with the availability heuristic, Schwarz, Bless et al. (1991) found that the subjects who recalled 6 examples of assertive behavior judged themselves to be more assertive than the subjects who recalled 12 instances of assertive behavior (and vice versa for recalling instances of unassertiveness), despite the latter subjects actually recalling *more* examples of assertive behavior. In this study, then, the judgment effect is better explained as an accessibility experience effect than as a direct accessibility effect (see also Stepper & Strack, 1993).

Schwarz, Bless et al. (1991) provided additional evidence that accessibility experiences produced the judgment effect in another study that manipulated the opportunity for subjects to misattribute their experiences. The subjects listened to "meditation music" over headphones as they worked on the recall task. They were told that the music was known to facilitate recalling a particular kind of autobiographical memory—either experiences characterized by assertiveness or by insecurity. Schwarz, Bless et al. (1991) found that the results of the previous study were replicated when the music did not provide an explanation for the relative ease of recalling a particular kind of autobiographical memory (i.e., did not provide an explanation for their accessibility experience). But when the music did provide an explanation, the earlier results disappeared. Indeed, when the accessibility experience could be attributed to the music, the subjects based their judgments on the number of examples of each type of behavior that they actually recalled.

Jacoby, Kelley, Brown, and Jasechko (1989) describe familiarity effects on judgments of a person's fame that are analogous to Tversky and Kahneman's (1973) account of likelihood estimations based on the availability heuristic. They argue that familiarity is the attribution of variation in processing to a particular source. The name of a person, for example, will seem familiar if it can be easily brought to mind. If this feeling of familiarity is attributed to the name having been frequently encountered outside the context of the experiment, then the name will be judged as famous (see also Jacoby & Kelley, 1990).

One way to increase the ease with which a name can be brought to mind (or can be read fluently) is to have prior exposure to it. In a study supposedly concerned with people's speed and accuracy in pronouncing words, Jacoby et al. (1989) had subjects read aloud a list of nonfamous names taken from a telephone book. Half of the nonfamous names were presented once and the other half were presented four times. The subjects made fame judgments either immediately or after a 24-hour delay for a set of names that included famous names, old nonfamous names, and new nonfamous names. Before making the fame judgment, all subjects were reminded that all of the names that they had read before were nonfamous names.

Jacoby et al. (1989) found that on the immediate test subjects were more likely to judge new nonfamous names as

being famous than old nonfamous names, but after the 24-hour delay they were more likely to judge old nonfamous names as famous than new nonfamous names. Interestingly, the latter delayed effect was stronger in the single-prime condition than in the four-prime condition. This suggests that subjects were more aware at the time of the fame judgment test that the old nonfamous names had appeared in the earlier experimental list when they had been repeated in the list four times. And the subjects knew that all the names from the experimental list were nonfamous. Indeed, the subjects were quicker to reject old nonfamous names as being famous when they had been repeated four times versus only once in the experimental list. Thus, when the nonfamous names were repeated four times, subjects could have attributed their experience of familiarity to having read the names in the experimental list rather than to the general fame of the name.

The importance of awareness of source was demonstrated in another study by Jacoby et al. (1989) involving the same basic paradigm. In this study, half of the subjects were asked only to make the fame judgment of the test names. The other half were asked to make both the fame judgment *and* a judgment of whether the name had appeared on the experimental list or not. In the fame judgment test, old nonfamous names were more likely to be judged as famous than new nonfamous names when the fame judgment was the only judgment, but this effect disappeared (and even slightly reversed) when subjects also made a judgment of the appearance of each name on the earlier experimental list.

The Jacoby et al. (1989) studies suggest that when a stimulus is easily brought to mind (or processed fluently), it is judged to be familiar. And experiencing it as familiar can produce the judgment that it is famous, as long as the familiarity is not attributed to the momentary experimental situation. In these studies, the nonfamous names were probably not stored prior to the experiment. Thus, exposure to the experimental list probably did not activate the names. This does not mean, however, that exposure to the experimental list did not increase the accessibility of the names. If the names were not stored prior to exposure, then exposure made the names available to subjects. And, as discussed earlier, a change from nonavailability to availability necessarily increases accessibility. In the Jacoby et al. (1989) studies, then, previous exposure to nonfamous names allows them to be more easily brought to mind than new nonfamous names given that the latter cannot be brought to mind at all. The experience of some accessibility makes the old nonfamous names seem familiar, which increases the likelihood of judging these names to be famous. As in the Schwarz, Bless et al. (1991) misattribution study, the disappearance of this effect when subjects can attribute the familiarity to exposure to the experimental list indicates that it is the experience of accessibility and not accessibility per se that drives the "becoming famous" effect.

The accessibility experience effect associated with the availability heuristic involves people inferring that a category (including its name) is easy to retrieve, or that a category's instances are easy to retrieve, because that category has been instantiated many times in the past. Underlying this inference is the reasonable assumption that frequently

occurring categories are likely to become highly accessible. It would also be reasonable to assume that important or relevant types of knowledge are likely to become highly accessible. Thus, people might infer that highly accessible knowledge must be motivationally significant even though the accessibility was actually due to recent priming (see Higgins, 1989a; Sherman et al., 1990). Consistent with this possibility, Sherman et al. (1990) found that target person information related to dimensions that were recently primed had more influence on subsequent evaluations and preferences, perhaps because the experienced high accessibility of these primed dimensions led subjects to infer that they must be especially relevant or important.

Awareness of the Priming Event: Contrast versus Assimilation

As described earlier, priming increases the accessibility of available knowledge such that the stored knowledge is more likely to be activated and used in subsequent processing of stimulus information. In the early category priming studies, increased accessibility from priming increased the likelihood that the subsequent stimulus behavior would be judged as belonging to the primed category—an "assimilation" effect. The early priming studies were explicitly designed to minimize the possibility that subjects would have the priming events in mind at the time that they judged the behavioral stimulus. The studies on accessibility experiences make clear that subjects' awareness of potential sources of accessibility can influence their likelihood estimates or their judgments of fame. How does awareness of priming events influence categorizations of target person behaviors? By now, several studies have addressed this question.

Martin (1986) specifically examined the issue of what happens when stored constructs are primed blatantly rather than subtly. To increase the likelihood that the subjects would associate construct activation with the priming events, the priming method blatantly activated the constructs. The subjects were given a series of phrases and were asked to decide which of two personality traits they exemplified. They also indicated their choice by placing the number of the phrase under a column headed by the trait they selected. In the positive priming condition, "boldness" and "self-assurance" were each primed four times. In the negative priming condition, "foolhardiness" and "egotism" were each primed four times. Following two distractor tasks, the subjects read the ambiguous target person behaviors used in Higgins, Rholes, and Jones (1977) and then gave their impressions of the target person by rating him on scales bounded by the alternative traits (e.g., reckless/adventurous).

Martin (1986) found with this blatant priming task that subjects rated the target person more on the positive end of the scales in the negative priming condition than in the positive priming condition—a "contrast" effect. Higgins et al. (1985) had found that the assimilation effect of recent priming on categorization was stronger among subjects who did *not* remember the recent prime. To examine this finding more systematically, Lombardi, Higgins, and Bargh (1987) replicated the basic design of Higgins et al.

(1985) and divided subjects into two groups—those who did and those who did not remember the priming events (assessed by a free recall measure taken after the categorization task). A strong assimilation effect of priming was found among those subjects who did not remember any of the priming events, replicating the basic results of Higgins et al. (1985). But an equally strong *contrast* effect was found among those subjects who *did* remember one or more priming events.

The results of the Martin (1986) and Lombardi et al. (1987) studies suggest that a contrast effect of priming rather than an assimilation effect can occur when it is likely that subjects are aware of or remember the priming events at the time that they make their stimulus judgment (see also Philippot, Schwarz, Carrera, De Vries, & Van Yperen, 1991). Subsequent studies have provided more direct evidence to support this conclusion.

Newman and Uleman (1990) used a modified version of the Higgins et al. (1977) paradigm to prime two evaluatively positive trait constructs and two evaluatively negative trait constructs. After priming, the subjects read 12 sentences that included a set of four ambiguous behaviors that were related to the primes. After reading all the sentences, the subjects were asked to recall them using trait-related recall cues. When the subjects did not recall the primes, they remembered the sentences better when the recall cue matched the prime-related interpretation. The subjects in this case apparently used the primed construct to interpret the ambiguous behavior, and thus a recall cue related to this interpretation facilitated remembering the sentence—an assimilation effect. When the subjects did recall the primes, however, they remembered the sentences better when the recall cue matched the prime-unrelated alternative interpretation. The subjects in this case apparently did not use the primed construct to interpret the ambiguous behavior but instead used the alternative construct, and thus a recall cue related to this alternative interpretation facilitated remembering the sentence—a contrast effect. In another study, Newman and Uleman (1990) found a contrast effect when the primes were blatantly given to subjects just prior to their being explicitly asked to interpret the ambiguous behaviors.

Moskowitz and Roman (1992) examined the judgmental effects of prior construct activation from spontaneous versus conscious inferences. All subjects were given different sentences that implied different trait-related constructs. Half of the subjects were explicitly asked to “form an impression” of the actor in order that their trait-related inferences would be conscious. The other half of the subjects were asked simply to “memorize the sentences” in order that their trait-related inferences would be spontaneous. In both conditions, the subjects made trait-related inferences. For all subjects, therefore, the trait constructs were activated and increased in accessibility during the first phase of the experiment. However, the subjects were either aware or unaware, respectively, of having made the trait-related inferences. Following this first phase of the experiment, all subjects were given the ambiguous behavioral descriptions used in Higgins et al. (1977) and later rated the target person’s behaviors using scales bounded by the applicable trait dimensions. The other experimental

manipulation was the type of sentences that subjects received in the first phase. One-third of the subjects received sentences implying the negative end of the trait dimensions (i.e., “stubborn,” “conceited”), one-third received sentences implying the positive end (i.e., “persistent,” “confident”), and one-third received sentences irrelevant to these trait dimensions.

Moskowitz and Roman (1992) found that the construct activation or priming that occurred in the first phase of the experiment produced an assimilation effect on subjects’ subsequent judgments of the ambiguous behavioral descriptions when the priming involved spontaneous inferences of which they were unaware. But when the priming in the first phase involved conscious impression formation, it produced a contrast effect on subjects’ subsequent judgments. To control for any possibility that the priming produced a “halo” effect or evaluative orientation, Moskowitz and Roman (1992) replicated their first study but changed the priming conditions so that all subjects had both one positive trait and one negative trait primed. The same pattern of results were obtained as in the first study.

In the most direct test of the effect of memory for the priming events on subsequent judgments, Strack, Schwarz, Bless, Kubler, and Wanke (1993) primed either “helpful” or “dishonest” with four synonyms using a modified version of the Higgins et al. (1977) priming procedure. Following a distractor task, the subjects were or were not assigned to a reminding condition in which they were asked to remember the words that had functioned as primes in the first task. All subjects then read a story about a target person who behaved in a manner that could be interpreted as either “helpful” or “dishonest” and then rated him on a scale bounded by the alternative traits. Strack et al. (1993) found an assimilation effect of priming for the subjects who were not reminded of the priming events prior to their behavioral judgment. But a contrast effect of priming was found for the subjects who were reminded of the priming events.

Together, this set of studies indicates that either contrast or assimilation effects of priming can be obtained depending on people’s awareness of the priming events when they subsequently make judgments of stimulus information. As varied as these different studies have been, their findings, as well as those from other studies (e.g., Erdley & D’Agostino, 1988; Smith & Branscombe, 1987; Smith, Stewart, & Buttram, 1992), suggest that assimilation effects of priming are more likely when people are *not* aware of the priming events at the time they make their judgment of the stimulus.

It is when people *are* aware of the priming events at the time that they make their stimulus judgment that the story becomes complicated. From the results of these studies, one might be tempted to conclude that when people are aware or remember the priming events at the time of stimulus judgment, increased accessibility from recent priming produces a contrast effect. Other findings, however, indicate that this is not the case. In one of the conditions of Martin’s (1986) study, for example, the subjects were prepared to respond to 12 phrases but they were interrupted and told to proceed to the next task after completing

only 8. The priming was just as blatant in this condition, but it produced an assimilation effect on subjects' subsequent judgments. Lombardi et al. (1987) also found that when the priming task was interrupted, an assimilation effect of increased accessibility from priming was obtained even for subjects who later remembered the priming events. Thus, awareness or memory of the priming events does not inevitably produce contrast effects.

Martin, Seta, and Crelia (1990) again used a blatant priming procedure in a series of studies that manipulated people's motivation and capacity for expending effort in forming an impression of a target person who was described in evaluatively ambiguous terms. One study reduced capacity by asking half of the subjects to perform a distracting task while trying to form their impressions (i.e., increasing cognitive load). Another study reduced motivation by telling half of the subjects that their impressions would be averaged with other members of their group (i.e., inducing "social loafing"). A third study selected subjects who were low in need for cognition and thus were chronically nondisposed to expend cognitive effort. In the alternative conditions of high capacity and high motivation, Martin (1986) found the usual contrast effects of blatant priming. But in the conditions where capacity and motivation were reduced, Martin et al. (1990) found assimilation effects of increased accessibility from blatant recent priming!

What is clear from these and other studies (e.g., Herr et al., 1983; Philippot et al., 1991) is that several factors beyond knowledge accessibility determine how increased accessibility from recent priming will influence subsequent stimulus judgments. One such general factor is people's experiences of the accessibility and activation, including their awareness of the priming events themselves. Different proposals for how such experiences might influence the use of activated knowledge will be considered later (see also Schwarz & Clore, Chapter 15, this volume).

Accessibility Specificity

From the studies reviewed thus far, investigators might assume that accessibility effects would be the same for people who had the same construct primed, or had the same construct chronically accessible, and then responded to the same stimulus information. This would be a reasonable assumption if knowledge accessibility and its effects only concerned the activation potential of separate, relatively general constructs, such as individual lexical items in semantic or declarative memory. As described next, however, there is considerable evidence indicating that this is not the case.

Accessibility of Specific "Trait-Behavior" Linkages

Smith and Branscombe (1987) tested whether the relation between the specific content of the prime and the content of the stimulus information influences the strength of the priming effect on behavioral judgments. Trait-related constructs, such as "hostility" for example, were frequently

primed, either by using Srull and Wyer's (1979) procedure where subjects form sentences describing hostility-related behaviors or by having subjects circle pairs of trait words that had the same meaning and were mostly hostility-related. After a delay of approximately 15 minutes, the subjects were given Srull and Wyer's (1979) paragraph describing a target person's "ambiguously" hostile behaviors and rated the target person's hostility.

Smith and Branscombe (1987) found that priming by trait words did not influence behavioral judgments after the 15-minute delay, but priming by behavioral descriptions produced an assimilation effect. In a second study, they found that priming by trait words produced an assimilation effect after a 15-second delay but had no effect after 3 minutes, whereas priming by behavioral descriptions produced an assimilation effect after both delays. Smith and his colleagues (e.g., Smith, 1989; Smith, Branscombe, & Bormann, 1988; Smith, Stewart, & Buttram, 1992) have now conducted several studies that compellingly demonstrate just how specific priming effects can be as a function of prime-stimulus similarity.

Subjects in a study by Smith (1989), for example, were given behavior verbs, such as "help" or "advise," and were told to consider the word as a behavior that one person might do to another. They then judged whether or not the behavior could be characterized by a target trait, such as "friendly" or "intelligent." Some of the specific behaviors appeared twice. The dependent measure was reaction time to make each judgment. Even when there were as many as 16 words between repetitions, judgments of repeated behaviors were significantly faster. Moreover, judgments of repeated behaviors were significantly faster than judgments of new behaviors even when the repetition occurred 24 hours apart. This facilitation occurred only when the repeated behavior was judged in relation to the same target trait as for the initial judgment.

Smith et al. (1992) also examined specific linkages between a particular trait construct and a particular stimulus behavior made accessible from past experiences of judging the behavior in terms of the trait. They found again that judgments of particular behaviors in relation to the same trait were quicker when the behaviors had been judged before than when they were new, even when the repetition was separated by a 7-day delay. Smith et al. (1992) also had subjects judge a "mixed" behavior that could be characterized either negatively (e.g., "unfriendly") or positively (e.g., "intelligent") or both. Even after a 7-day delay, subjects tended to evaluate a mixed behavior in the direction of whichever trait had been previously used to make judgments. Most important, there were two effects of prior trait judgments. There was a general practice effect of prior experience making judgments in relation to a particular trait, which was evident in subjects' evaluations of new mixed behaviors (see also Smith, 1989). And there was also a stronger, specific effect on evaluations of making prior judgments of the same mixed behavior in relation to the same trait. The results of these studies indicate that increased accessibility from priming concerns not just the accessibility of separate, general constructs, such as constructs in semantic or declarative memory. It also concerns the accessibility of procedures for relating traits to

different behaviors *and* the accessibility of highly specific linkages between particular traits and particular behaviors (see Smith, 1990).

Accessibility of Specific "Attitude Object-Evaluation" Linkages

Fazio (1986) has also proposed that object-evaluation linkages can vary in accessibility. He views an attitude as a specific association between a particular attitude object and its evaluation that can vary in strength. In one experiment by Fazio, Chen, McDonel, and Sherman (1982), half of the subjects completed a single attitude measure concerning how interesting they found each puzzle to be for a set of puzzles, and the remaining subjects expressed their interest in each type of puzzle an additional two times. Following this manipulation, all of the subjects performed a response-time task in which each of the puzzles was randomly paired with one positive and one negative evaluative adjective (e.g., "Letter Series: Interesting?"). The subjects had to respond as quickly as possible, "Yes" or "No," by pressing the appropriate button to indicate whether or not the adjective expressed their attitude toward that puzzle.

Fazio et al. (1982) found that the subjects were significantly quicker to respond when they had previously expressed their interest in the puzzles repeatedly than when they had previously expressed their interest only once. Thus, frequent activation of the specific association between a particular attitude object and evaluation increased the accessibility of that association (see also Powell & Fazio, 1984). In another study, Fazio et al. (1982) examined whether increased accessibility of a specific "attitude object-evaluation" linkage from prior frequent activation influences subsequent responding to the attitude object. As in the earlier study, the subjects expressed their interest in each puzzle either a single time or three times. Following this manipulation, the subjects had a "free play" period for 15 minutes when they could play with any of the puzzles as much as they wished. Both the order on which the puzzles were worked and how much of each puzzle's problems was attempted were measured and a combined score computed. For the sample as a whole, approximately one-fourth of the subjects had a correlation greater than .65 between their prior expressed attitudes toward the puzzles and their responses to the puzzles during the free play period. Over twice as many "repeated expression" subjects as "single expression" subjects were among these high correlation subjects.

The results of this study suggest that frequent activation of a specific "attitude object-evaluation" linkage can increase the consistency between people's attitude toward an object and their behavioral response to that object. Further support for this conclusion is provided in a study by Fazio and Zanna (1978). The subjects were in an introductory social psychology course and knew what volunteering as an experimental subject entailed. Most of the subjects had been given the opportunity the previous year to be in psychology experiments, but they varied in the number of experiments in which they had previously

participated. Several weeks after the start of the course, the subjects' attitudes toward participating in psychology experiments were measured. The measure of subjects' behavioral intentions consisted of their indicating the number of experiments in which they were willing to participate during the current year. The correlation between attitudes and behavioral intentions was found to be much higher for those subjects who were relatively high in prior experience with the attitude object (i.e., high past participation in experiments) than for those subjects who were relatively low in prior experience (see also Houston & Fazio, 1989). If greater direct experience means greater activation of the "attitude object-evaluation" linkage (see Fazio et al., 1982), then the results of this study support the conclusion that frequent activation of a specific "attitude object-evaluation linkage" can influence subsequent responses to the attitude object.

Accessibility of Specific "Actual Self-Desired Self" Linkages

Discrepancies between people's self-concept and their personal standards or desired selves are motivating and emotionally significant (e.g., Duval & Wicklund, 1972; James, 1890/1948; Rogers, 1961). These "actual self-desired self" discrepancies can be conceptualized as chronically accessible cognitive structures (see Higgins, 1989b). For each individual, there are specific linkages between particular actual-self attributes and particular desired-self attributes that involve matches or mismatches. A desired attribute, such as "kind" or "successful," might be represented as a hope or a wish for one individual (i.e., as an ideal-self attribute) but as a duty or responsibility for another individual (i.e., as an ought-self attribute). Higgins and his colleagues have shown (e.g., Strauman & Higgins, 1987) that priming desired-self attributes automatically produces distinct emotional syndromes depending on whether the subjects' actual-self attributes are discrepant from primed ideal or ought desired-self attributes. Even though the priming events in such studies involve positive traits, the emotions produced are negative when the subjects possess discrepancies. Moreover, the same priming event can make people feel sad (an ideal-self discrepancy) or feel worried (an ought-self discrepancy) or feel little at all (no discrepancy). These findings suggest that the priming activates specific "actual self-desired self" linkages and not just positive trait constructs.

Studies by Higgins and his colleagues have demonstrated that the accessibility of these "actual self-desired self" linkages can influence how people respond to the life events of another person (Higgins, Roney, Crowe, & Hymes, 1994; Higgins & Tykocinski, 1992). In these studies, the subjects read a story about a few days in the life of a target person and, after a filler task, were asked to recall all of the story events. Higgins and Tykocinski (1992) found that individual differences in the chronic accessibility of an "actual self-ideal self discrepancy" linkage versus an "actual self-ought self discrepancy" linkage influenced which events in the story were remembered. Higgins et al. (1994) found that priming actual self-ideal

self discrepancy linkages versus actual self-ought self discrepancy linkages also caused subjects to recall different story events. These accessibility effects on memory were not due to the content of the story events but were produced by the regulatory significance of the specific actual self-desired self linkages.

Some studies by Higgins, Van Hook, and Dorfman (1988) provide more direct evidence that specific actual self-desired self linkages can influence responding. The subjects performed a modified Stroop task. Each target word was preceded by a "memory" word that subjects had to repeat after naming the target word's color. This memory word functioned as the prime. The words used in the Stroop task were selected idiographically for each subject so that some prime-target word pairs were part of that subject's self system while other prime-target word pairs were not. Among the self-related word pairs, some of the prime-target word pairs involved discrepancies for the subject whereas other word pairs were nonproblematic. Evidence for linkage (i.e., slower latencies from Stroop interference) was found for a control condition of semantically related object-category pairs, which replicated previous cognitive studies. As expected, there was no evidence of linkage for the self-unrelated pairs. For the self-related pairs, however, the subjects' response latencies were slower for the self-related problematic pairs than for the self-related nonproblematic pairs.

Research on accessibility specificity is important because it highlights that accessibility effects are not restricted to the effects on judgment of increasing the accessibility of general semantic or declarative knowledge. Specific episodes and specific procedural linkages can increase in accessibility as well, and when they do, this can influence memory, feelings, and behavior.

Models and Proposals for Accessibility Effects

A variety of different accessibility effects were reviewed in the previous section. Not surprisingly, no one model has been developed to account for all of these effects. Moreover, there are no clearly distinct alternative models for some of these effects. For example, all current accounts of accessibility specificity are in terms of "associations," "linkages," or "connections" that permit the activation of one knowledge unit to influence the activation potential of another (cf. Smith, 1993). However, alternative proposals have been offered to account for certain accessibility effects, and these are discussed in the following sections.

Recent Priming

As reviewed earlier, a stored knowledge unit can be made temporarily more accessible by priming it, and such recent priming increases the likelihood that this knowledge unit, rather than alternatives, will be used in processing a subsequent stimulus. It should be noted explicitly that "knowledge unit" refers to *any kind of knowledge unit*, including general constructs (such as those associated with "semantic" or "declarative" memory), specific episodes and exemplars (such as those associated with "episodic memory"), and productions (such as those associated with

"procedural" memory). To account for recent priming effects on behavioral judgments, two basic types of models (or, perhaps, metaphors) have been proposed—*mechanistic models* and *excitation transmission models* (see Higgins et al., 1985). Mechanistic models account for recent priming effects in terms of the arrangement and the working of component parts. In contrast, excitation transmission models account for recent priming effects in terms of the heightening and the dissipation of excitation (or energy levels) from stimulation and decay.

Perhaps the most explicit mechanistic model was originally proposed by Wyer and Srull (1980, 1981). According to their "storage bin" model, long-term memory consists of a set of content-addressable storage bins. To explain past priming effects, the most relevant bins are the semantic bins that function as "mental dictionaries" of constructs that can be used to interpret social objects and events. One semantic bin includes constructs that can be used in judging a person's traits, states, or physical attributes. Another semantic bin includes constructs that can be used in judging a person's behaviors.

The model proposes that knowledge units, such as constructs, are stored in these bins, with more recently deposited units being stored on top. When a construct is drawn from the bin (i.e., is activated), it is then returned to the top of the bin regardless of its previous position in the pile. The critical implication of this proposal is that a construct in a given bin is more likely to be retrieved (i.e., activated) and used if it has been retrieved and used in the recent past. More generally, the constructs in each bin are stored in layers in the order in which they were previously activated. When stimulus information appears for interpretation, the relevant bin is searched from the top down so that constructs at the top are more likely to be retrieved and used. Thus, when several constructs are potentially applicable to judge a stimulus, the most recently used construct is most likely to be retrieved first from the top of the bin and applied to the stimulus. A conceptualization very similar to Wyer and Srull's storage bin model was also proposed by Forbach, Stanners, and Hochhaus (1974).

Accessibility effects of recent priming have also been interpreted in terms of various forms of excitation transmission (e.g., Collins & Loftus, 1975; Higgins & King, 1981; Marcel & Forrin, 1974; Reder, 1983; Warren, 1972; Wyer & Carlston, 1979). These models generally share the postulates that recent stimulation of a knowledge unit, such as a stored construct, increases that unit's excitation (or energy level), and that a unit's excitation level must reach a certain, minimal threshold for that knowledge to be activated and used to judge a subsequent stimulus. The critical implication in these models is that when several stored constructs are potentially applicable to judge a stimulus, the most recently activated construct is likely to have the highest excitation level and to reach activation threshold most quickly or easily in the presence of the stimulus.

This excitation transmission proposal for recent priming effects involves basically the same assumptions as Anderson's ACT* procedural strengthening model (Anderson, 1983, 1987; see also Smith, 1993). Specifically, when an ACT* production—as another type of knowledge unit—

has been recently used, it accrues "strength"; and increased strength makes such knowledge units faster to execute in the presence of subsequent stimuli to which they are applicable. Thus, when several productions are potentially applicable to judge a stimulus, the most recently used production is likely to have the greatest strength and be executed faster in the presence of the stimulus.

Both the excitation transmission models and the ACT* procedural strengthening model propose that recent activation of a knowledge unit *energizes* that unit so that a subsequent stimulus will activate it *more quickly and easily*. The storage bin model, in contrast, proposes that recent retrieval of a knowledge unit *improves its hierarchical position* in the bin such that it will be *tried first* in responding to a subsequent stimulus. These distinct proposals have different implications for the relation between stored knowledge and stimulus information that will be considered not only in this section of the chapter but in later sections as well.

A Historical Note. Newcomb, Turner, and Converse (1965, p. 38) used a "thinking machine" or computer metaphor to account for accessibility effects that combined some features of each of the preceding models. Their proposal resembled a storage bin model of productions that varied in speed of execution as follows:

Thus the machine storage of information is organized into "levels" of more or less rapid accessibility. Quite naturally, an efficient programming of the machine stores information that is going to be operated on most frequently in the most rapidly accessible storage level. The immediate situation or recent experience may serve to render certain types of stored information or cognitive contexts more readily accessible or salient for the individual.

Delay from Priming to Stimulus Exposure

As described earlier, the effect of recent priming of trait-related constructs on judgments of subsequent stimuli generally decreases as the priming-to-stimulus delay increases. This decrease can be relatively fast or slow depending on the frequency of the recent priming and whether or not alternative applicable constructs have also been primed. The preceding models have also attempted to account for this accessibility finding. According to Wyer and Srull's (1980, 1981) storage bin model, a construct will remain at the top of a bin for a substantial period as long as other constructs in the bin are not activated during the priming-to-stimulus interval. As the priming-to-stimulus delay increases, however, other constructs in the bin are more likely to be retrieved and then placed back on top of the bin. Therefore, the primed construct is less likely to remain on top of the bin as the priming-to-stimulus delay increases, and thus is less likely to be tried first in judging the stimulus.

Excitation transmission models generally assume that the excitation level of a construct (or knowledge unit) that has been increased by priming begins to decay after the priming. As the priming-to-stimulus delay increases, the excitation level of the construct decreases. Thus a subsequent

stimulus will activate the construct less quickly and easily. The ACT* procedural strengthening model also assumes that the strength of a production will decrease with increasing priming-to-stimulus delay, but the decrease can be quite small even for large priming-to-stimulus delays, especially when the production has been frequently primed in the past (see Smith et al., 1992).

Frequent Priming

Past studies have found that the more frequently a construct is primed, the more likely it is that the construct will be used in judging a subsequent stimulus. According to Wyer and Srull's (1980, 1981) storage bin model, if a construct is primed only once, then it can lose its position on top of the bin with increasing priming-to-stimulus delay. But if a construct is primed frequently over time, then at any particular point it is more likely to have been the most recent construct used and thus be back on top of the bin. According to this model, then, the effect of frequent priming is understood in terms of the likelihood of there having been a recent priming effect.

Excitation transmission models account for the effect of frequent priming quite differently. These models propose that the more frequently a construct (or knowledge unit) is stimulated, the more likely that an excitation level near the threshold for activation will be maintained. Higgins et al. (1985) point out, however, that for excitation transmission models to be as specific as Wyer and Srull's (1980, 1981) storage bin model about the underpinnings of priming effects, additional assumptions must be made regarding two parameters: *level of activation* and *decay function*. For level of activation, one could assume either that activation level is binary (i.e., all or none) or that activation level is continuous, such that constructs can build up different levels of activation above the minimal threshold of activation (e.g., Wyer & Carlston, 1979).

For decay function, two variables need to be considered. First, one could assume either that the decay function is the same for all activated constructs (i.e., a *uniform* decay function) or that the decay functions vary depending on their prior history of activation (i.e., *nonuniform* decay functions). Second, one could assume either that the decay occurs at a *constant* rate (or amount) or that there are varying, *nonconstant* rates (or amounts) of decay depending on a construct's prior history of activation. A uniform and constant decay function for primed constructs was proposed by Marcel and Forrin (1974). A nonuniform and nonconstant decay function that varies with different levels of activation is implicit in Wyer and Carlston's (1979) proposal that residual excitation decreases more rapidly initially than later on, with higher decay rates for higher levels of activation regardless of the source of the level of activation.

In their synapse model, Higgins et al. (1985) propose another alternative—a binary level of activation (i.e., either activated or not, although preactivation excitation levels vary up to some maximum level needed for activation), and a decreasing rate of decay with increases in the frequency or intensity of past priming. The difference between this model and alternative models becomes clear

when one considers what each model would predict for the Higgins et al. (1985) experimental design discussed earlier in which one applicable construct was frequently primed whereas an alternative applicable construct was primed only once but most recently. According to Wyer and Carlston's (1979) "battery"-like model, whichever construct has the higher level of activation immediately after the final priming will continue to predominate regardless of the priming-to-stimulus delay. According to Wyer and Srull's (1980, 1981) storage bin model, the recently primed construct would predominate over the frequently primed construct regardless of the priming-to-stimulus delay as long as no applicable construct was retrieved during the delay period. The Higgins et al. (1985) study had subjects perform a difficult interpolated task (i.e., "counting backward") during the delay period to prevent the retrieval of an applicable construct during that period.

In contrast to these alternative models, the "synapse" model predicts a reversal of construct predominance as a function of priming-to-stimulus delay. According to this model, both the frequently primed construct and the recently primed construct have the same maximal level of excitation when the construct is primed for the final time. But the frequently primed construct has its final priming before the recently primed construct. Immediately after the final priming of the recently primed construct, therefore, the excitation level of the recently primed construct will have decayed less than the excitation level of the frequently primed construct. The recently primed construct will thus predominate at this point. But the frequently primed construct has a slower rate of decay than the recently primed construct. Given sufficient priming-to-stimulus delay, therefore, the frequently primed construct will predominate at a later point. As described earlier, Higgins et al. (1985) obtained this predicted crossover interaction. And Bargh et al. (1988) found the same basic pattern in their conceptual replication in which chronic accessibility substituted for frequent priming.

After the Higgins et al. (1985) study, Wyer and Srull modified their "storage bin" model in a manner that permits the model to account for this pattern of results (see Wyer & Srull, 1986, 1989). First, they propose that the probability of identifying any given construct during the top-down search of a semantic bin might be rather low, such as $p = .3$. Second, they propose that *each time* a construct is used a *copy* of the construct is placed back on top of the bin. Therefore, a construct that is frequently primed will have several copies in the bin. By combining these two proposals, it is possible to account for why the frequently primed construct predominates. As Wyer and Srull (1989) point out, however, it is now more difficult to account for why the recently primed construct would ever predominate. They propose that the interpolated task clears working memory completely after a long delay but only partially after a very short delay (e.g., 15 seconds). In addition, during the short delay, the copies of the frequent primes are cleared before the copy of the most recent prime. After the short delay, therefore, a copy of the recently primed construct is more likely to be retained in

working memory than copies of the frequently primed construct, and an applicable construct in working memory is likely to be used to encode the behavioral stimulus.

Another factor needs to be considered, however. If a primed construct remains in working memory before and during encoding, subjects are more likely to remember the priming event and to be aware that the construct was already activated *before* stimulus exposure. As reviewed earlier, such accessibility and activation experiences could make a contrast effect of priming more likely than the predicted assimilation effect.

Priming and Chronicity

In an earlier version of the synapse model, Higgins and his colleagues (Higgins & King, 1981; Higgins et al., 1982) proposed that chronic individual differences in construct accessibility function like temporary accessibility differences produced by priming. In this way, the insights of Bruner (1957a, 1957b) and Kelly (1955) could be integrated for the first time. Chronic accessibility was considered to be like construct accessibility from highly frequent priming. It was proposed that to the extent that people's social experiences vary in the frequency with which particular constructs are activated, then individual differences in chronic accessibility should develop. And the more frequently a construct has been activated in the past, the longer its relatively high accessibility should persist.

As discussed earlier, the decay functions of accessible constructs (or knowledge units) vary depending on their past history of activation. Thus, the persistence of a construct's accessibility over time depends on its source of activation (e.g., just recent priming or frequent priming). According to the synapse model, however, the effect of an accessible construct on stimulus judgment, *at the time of judgment*, depends solely on its level of accessibility. The source of the accessibility is irrelevant (see Bargh et al., 1986; Higgins, 1989a). A construct at a particular level of accessibility will have the same impact whether the construct just reached that level from recent priming or had maintained that level for a long time because of past frequent priming. By the same token, different sources of accessibility can combine, and what matters is the final level of accessibility obtained by the time of stimulus judgment and not how much each source contributed to that level. In this sense, different sources of a construct's accessibility combine additively (see Bargh et al., 1986; Higgins, 1989a). At the time of judgment, different combinations of sources of accessibility that yield the same final level are not distinguishable. For the reasons discussed earlier, however, different combinations do matter for how long a particular level of accessibility will persist.

As Bargh et al. (1986) noted, Wyer and Srull's (1980, 1981) storage bin model posited no separate and independent effects of accessibility from recent priming versus from chronic accessibility or frequent priming. Chronic accessibility or frequent priming is postulated simply to increase the likelihood of obtaining an effect of recent priming. If recent priming already has an effect, therefore, no additional effect from frequent priming would be expected. As reviewed earlier, however, there is evidence

that recent priming and chronic accessibility can combine to increase accessibility effects (Bargh et al., 1986; Higgins & Brendl, 1995).

The original storage bin model cannot account for this effect. As discussed earlier, the revised storage bin model (Wyer & Srull, 1986, 1989) predicts that accessibility effects will result from the frequent activation of a construct unless there is a short priming-to-stimulus delay following recent priming, such as the 15-second delay used in Higgins et al. (1985). Wyer and Srull (1989) also suggest that recent priming can continue to predominate beyond a short delay if the recently primed construct is primed twice rather than just once. Subjects in the short delay condition of Higgins and Brendl's (1995) study had the most recently primed construct primed three times and were exposed to the stimulus description 15 seconds after the final priming. Under such circumstances, the revised storage bin model should predict that recent priming would be the source of any accessibility effect. The addition of chronic accessibility for the same construct should contribute little, if anything. Instead, Higgins and Brendl (1995) found a highly significant effect of levels of chronic accessibility in the short delay, priming condition (with no recall). It is difficult for the storage bin model to account for this effect of levels of chronic accessibility beyond both recent and frequent priming.

Contrast versus Assimilation Effects of Priming

Herr et al. (1983) found that when subjects were primed with moderate exemplars of animal size or ferocity (using a modified version of the Higgins et al., 1977, priming paradigm), and then judged the size or ferocity of a fictitious animal (referred to as "ambiguous" stimuli), an assimilation effect of priming on judgment was found. But when subjects were primed with extreme exemplars of animal size or ferocity, the priming produced contrast effects both for fictitious animals and for real animals (referred to as "unambiguous" stimuli).

Herr et al.'s (1983) explanation for these effects was in terms of principles of social judgment theory (Sherif & Hovland, 1961). They suggest that the primed animals become salient category exemplars, such as salient exemplars of "animal size," and together function as an anchor for the subsequent judgment. When the anchor is moderate or relatively neutral, a fictitious target animal, which by definition has highly vague features, can fall within the latitude of acceptance and be assimilated to the anchor. But when the anchor is extreme, any target animal, fictitious or not, is likely to fall within the latitude of rejection and be contrasted from the anchor. This "social judgment and anchoring" explanation can also account for the comparable findings in Herr's (1986) subsequent study of the effects of exemplars varying in hostility on judgments of the "ambiguous" target person's behaviors used by Srull and Wyer (1979, for an alternative perspective, see Schwarz & Bless, 1992).

Higgins (1989a) proposed an explanation of contrast effects more generally that was similar to Herr et al.'s (1983) explanation, although episodic memory for the

priming events themselves received more emphasis. When priming events are remembered, implicitly or explicitly, the events themselves can function like standards or reference points relative to which stimulus behaviors are judged (see also Wyer & Srull, 1989). When priming events describe a person, as in the Srull and Wyer (1979) and Higgins et al. (1985) paradigms used in many studies, subjects are provided with statements or labels describing people who possess or instantiate a particular trait, such as hostility. According to the rules of language comprehension (see, e.g., Huttenlocher & Higgins, 1971), such statements are likely to be represented as a "greater than average" amount of the trait. Moreover, the phrases used in many studies have involved relatively extreme behaviors (e.g., "leg break arm his"; "his favorite topic of conversation is himself"). The subsequent use of these *priming exemplars* as standards when judging less extreme or less clear targets is likely to produce contrast effects. Philippon et al. (1991) provide a discussion of *when* priming exemplars simply increases the accessibility of the specific exemplars themselves or increases the accessibility of the exemplar-related category, and *how* this mediates contrast versus assimilation effects of priming.

As a general explanation, then, the "use of priming exemplars as standards" can account for various studies that have found contrast effects of priming. This explanation, however, is not sufficient alone to account for when priming produces assimilation versus contrast. According to this explanation, it is the relation between the extremity of the standard and the extremity of the target that determines when contrast versus assimilation occurs. But several studies have found both contrast and assimilation for the *same* priming events and targets. Indeed, this is what Martin (1986) found under conditions where *all* the subjects were likely to be fully aware of the priming events. When the subjects completed the priming task, a contrast effect was found. But when the subjects believed that the priming task had been interrupted, an assimilation effect was found.

Martin's (1986) explanation for contrast effects when people are aware of the priming events at the time of judgment is quite different from the explanation in terms of the "use of priming exemplars as standards." Martin (1986) suggests that, when priming is very subtle, people may not even be aware that the primed construct has been recently activated in them. Thus, when this construct comes to mind in the presence of the subsequent stimulus, they are likely to infer that it is just their own spontaneous reaction to the stimulus. In other words, the applicability of the construct to the stimulus, rather than construct accessibility from contextual priming, is treated as the source of construct activation in the presence of the stimulus (see Higgins & Stangor, 1988).

In contrast, when a construct is blatantly primed and individuals remain aware of this priming at the time of judgment, they might infer that the construct is activated in the presence of the stimulus because of their prior exposure to the priming events rather than because of the close relation between the construct and the stimulus. They might perceive the priming events of the prior task to be biasing their

response to the stimulus in the new task, and thus actively avoid using the primed construct to interpret the stimulus. Such suppression or active inhibition of the primed construct would make alternative applicable constructs more likely to be used to judge the stimulus. This would yield a contrast effect of priming for stimulus behaviors to which bipolar constructs are applicable.

Martin (1986) proposes another variable to account for the assimilation effect of priming in the "interrupted" condition. He suggests that when a task is interrupted, thoughts related to the task are more likely to perseverate. When a priming task is interrupted, therefore, the constructs activated by the priming persevere. The subjects remain in a state of readiness to use the constructs in responding to subsequent stimuli for which the constructs are applicable. The key variable here is response perseveration (cf. Ford & Kruglanski's, in press, "need for closure" interpretation of assimilation effects of priming).

For the "use of priming exemplars as standards," as well as the "suppression of priming exemplars," individuals can be aware of the priming events at the time of judgment. But these two explanations make opposite assumptions about the implications of such awareness. According to the "use of priming exemplars as standards" explanation, individuals consider the priming events to be contextual exemplars that are *relevant* for judging the target stimulus. But according to the "suppression of priming episodes" explanation, the priming events are potential sources of encoding bias and thus are *not appropriate* for judging the stimulus (see Martin, 1986; Martin & Achee, 1992). These different assumptions have not yet been directly tested, such as by eliciting subjects' judgments of relevance and appropriateness. These judgments, however, are *not* the same. On the one hand, subjects might believe that the conspicuousness of the primed category could bias their judgment of the relation between the stimulus and that category, and thus believe it would be inappropriate to identify the target in terms of the primed category. On the other hand, they might believe that the priming episodes are representative exemplars and thus are relevant as standards of comparison. The suppression of priming episodes concerns the former belief, whereas the use of priming exemplars as standards concerns the latter belief.

Each of these explanations provides a reasonable account of much of the contrast versus assimilation effects of priming. As mentioned earlier, however, the use of priming exemplars as standards explanation does not account for why assimilation occurs when the blatant priming task is interrupted given that the priming events and target stimulus remain the same. By adding a "disinhibition" variable, and considering multiple effects of priming events, Higgins (1989a) suggests how this might occur. Priming might both provide specific exemplars when the episodes are remembered *and* increase the accessibility of the stored constructs. When subjects actively use the priming exemplars as standards, the direct assimilation effect of accessibility on judgment is inhibited because an active set inhibits a passive process (see Higgins & King, 1981; Neely, 1977; Posner & Warren, 1972). The "interruption" condition used by Martin (1986) and by Lombardi et al. (1987; Study 2) is

unpleasant and disruptive. In this condition, therefore, active use of the primes as standards might be inhibited or interfered with in some way because of associations with negative prior experiences. By inhibiting the active use of the exemplars as standards, the passive accessibility effect on identification would be disinhibited and produce the assimilation effects found by Martin (1986) and others (e.g., Lombardi et al., 1987).

Priming exemplars themselves can be not only blatant but so unpleasant that the primes are inhibited because of *negative responses to them*, such as priming subjects with words like "idiot," "moron," and "dimwit" prior to their judging a target person who, in addition to displaying "mixed" behaviors, is explicitly labeled as mentally retarded (see Skowronski, Carlston, & Isham, 1993). Ironically, this kind of extreme, blatant priming might produce an assimilation effect because inhibiting the use of the exemplars as standards could disinhibit the passive assimilation effect of increased accessibility. In the labeling condition, Skowronski et al. (1993) found on a general evaluative judgment of the target person that priming with moderate exemplars produced a slight contrast effect, but priming with the extreme exemplars produced, if anything, an assimilation effect (relative to a no priming condition). In this regard, the "suppression of priming episodes" explanation described earlier could account for a weaker contrast effect for such unpleasant, blatant primes (see Martin & Achee, 1992, for how a perceived priming bias is partialled out), but it could not account for an assimilation effect.

The preceding extension of the use of priming exemplars as standards explanation implicitly raises an important point that Strack (1992) has explicitly addressed (cf. Martin & Achee, 1992). What is being suggested is that simple awareness of the priming episodes does *not* necessarily lead to their suppression. Their use might be suppressed or inhibited if they are associated with negative experiences, such as prior task interruption or distasteful content. They might also be suppressed if they are perceived to be unrepresentative. But if they are perceived to be representative, they could be used as standards regardless of subjects' awareness of their appearance in the priming events (see Strack, 1992). Thus, extreme exemplars in priming events might still be considered relevant as standards of comparison because they could be perceived as representative. Individuals could be aware that these exemplars appeared in the previous task but *not* be aware later that their relatively high accessibility was due to this factor. Instead, the experience of their higher accessibility could produce a *feeling of representativeness* or of familiarity. Such feelings could be sufficient for individuals to consider the exemplars to be relevant to their judgment.

Judged Usability and Inhibitory Processes

The literature on assimilation and contrast effects of priming highlights the possibility that there are inhibitory (and disinhibitory) processes involved in knowledge accessibility and use. These processes could be controlled or

automatic. The discussion thus far has focused on controlled processes.

Judged usability, such as the perceptions of relevance and appropriateness reviewed earlier, involves a controlled process in which people judge whether activated knowledge is relevant or appropriate to use in responding to a stimulus. Judged usability occurs *after* knowledge activation but *before* knowledge use. To take the strongest case, when activated knowledge is perceived to be both inappropriate and irrelevant, its use in responding to a stimulus is likely to be actively inhibited or suppressed. The examples of judged usability previously discussed concern the use of activated knowledge to categorize a target person's behavior. There are other cases as well (see also Higgins & Bargh, 1987). A study by Devine (1989), for example, found that both high-prejudice and low-prejudice White subjects had equal knowledge of African American pejorative stereotypes, but low-prejudice White subjects were less likely to express these stereotypic beliefs when the social context of the study made them conscious of race as a person characteristic. Apparently, low-prejudice subjects suppressed expressing stereotypic beliefs when they were aware of the inappropriateness of doing so (see also Biernat & Vescio, 1993). Among subjects who were low-prejudice toward the opposite sex, Higgins and King (1981; Study 1) also found that when the social context made gender salient, subjects were less likely to respond in a stereotypic manner.

Controlled inhibitory processes could also occur because individuals are motivated not to respond with the first applicable knowledge that is activated. Kruglanski (1989) has proposed that individuals, temporarily or chronically, can be motivated not to use the first thought that comes to mind because of a strong need to avoid closure, such as a strong motivation for accuracy or a fear of invalidity. When so motivated, these individuals will actively inhibit the immediate use of any one possibility and seek alternatives. Thus, even when priming increases the accessibility of some stored knowledge and the likelihood that it will be subsequently activated first, this priming effect might not be evident in the individuals' responses because the activated knowledge is inhibited to avoid (premature) closure. Indeed, some recent studies have found that priming effects on subjects' subsequent responses are weaker when the subjects are motivated to avoid closure (see Ford & Kruglanski, in press; Schuette & Fazio, in press; Thompson, Roman, Moskowitz, Chaiken, & Bargh, 1994).

There is also evidence suggesting that conversational norms can influence priming effects on knowledge use because of judged usability. Strack et al. (1988) found that a question about "happiness with marriage" primed specific episodic knowledge that later influenced responses to a second question about "happiness with life-as-a-whole." When the conversational norm of nonredundancy was highlighted by making the two questions distinct, then the priming effect was eliminated (see also Schwarz, Strack, & Mai, 1991). These results suggest that when the conversational norm was highlighted, the subjects perceived using the primed knowledge in responding to a different question to be inappropriate and thus inhibited its use.

Wegner and his associates (e.g., Wegner, 1989; Wegner & Erber, 1992; Wegner & Schneider, 1989) have examined another kind of controlled inhibitory process—active suppression—that has intriguing implications for both knowledge activation and accessibility. The results of their studies indicate that subjects can use active strategies such as distraction to inhibit activation of a particular construct, at least temporarily. Their results also suggest that this process of inhibiting the activation of a construct actually increases the accessibility of the construct such that the construct will be *more* likely to be activated and used subsequently, especially under conditions that make any further suppression difficult—the "rebound" effect (see Wegner & Erber, 1992). These findings suggest that controlled inhibitory processes can increase the activation potential of a stored construct while successfully decreasing current activation, with the increased activation potential being revealed once disinhibition occurs.

Macrae, Bodenhausen, Milne, and Jetten (1994) use the synapse model to account for such rebound effects. They suggest that the ironic monitoring process described by Wegner and Erber (1992) produces frequent priming of the unwanted construct, which results in the activation potential of the unwanted construct decaying more slowly over time for suppressors than nonsuppressors. Suppressors, then, should be more likely to use the unwanted construct in their subsequent responses. The results of their studies both support this interpretation and extend the rebound phenomenon to the area of stereotypic beliefs. They found that suppression of stereotypic beliefs at one point in time increased the likelihood that these beliefs would influence subsequent evaluations of and behavior toward a stereotyped target. They also found that a prior period of suppressing a stereotype increased the subsequent accessibility of the stereotype.

Thus far, only controlled inhibitory processes have been considered. But automatic inhibitory processes can also occur. In the early nineteenth century, for example, Johann Friedrich Herbart suggested that every idea has the tendency to maintain itself and to drive out ideas with which it is incompatible. He proposed that ideas already in possession of the field regularly repel uncongenial ideas, thrusting them below the threshold of consciousness (see Heidbreder, 1933). Contemporary theories also state that already activated ideas can automatically inhibit competing ideas (e.g., Logan, 1980; Milner, 1957; Posner, 1978; Shallice, 1972). In a now classic demonstration of such inhibition, Neely (1977) had subjects perform a word-nonword classification task. Each visually presented target letter string was preceded by a signal word. In one condition, whenever the signal word "body" appeared, it was usually followed by the name of a building part as target (e.g., "door"). The signal word body, then, both primed the body category and induced an active set for the building category. Once the active set for instances of building is established as an idea, it should inhibit competing ideas, such as body category instances. Consistent with this prediction, the signal word body had the usual facilitative priming effect on responding to body category instances when the postsignal delay was too short to establish the active set (250 msec), but it had an inhibitory effect when

the postsignal delay was long enough to establish the active set (2,000 msec).

A study by Sedikides (1990) combined a priming manipulation with a manipulation of communication goals used in previous studies of the "communication game" (see Higgins, 1981). When communication goals were absent, the usual effects of priming on judgments were found. But these effects disappeared when communication goals were active (i.e., when there was a strong motivation to tailor the message to suit the audience's attitude). Apparently, an active set to attain communication goals (or follow communication rules) automatically inhibited the effects of competing constructs made highly accessible by priming. The subjects in this study, like those in Neely's (1977) study, did not actively intend to inhibit or suppress the constructs made highly accessible by priming.

Studies such as those by Neely (1977) and Sedikides (1990) concern the inhibitory effects of momentary active sets or goals on the activation and use of constructs made temporarily accessible through priming. Although not designed to examine inhibitory processes directly, the results of some personality studies have been interpreted as involving automatic inhibitory effects between chronically accessible constructs and persistent active sets. As discussed earlier, Higgins, King, and Mavin (1982) found in a couple of studies that subjects' memory and impressions of a target person's behaviors were influenced by their chronically accessible constructs. In Study 2, the subjects were divided into three groups according to their level of cognitive differentiation (see Zajonc, 1960)—high, moderate, and low. The low differentiation subjects were found overall to have a positivity bias in their impressions, consistent with the literature's finding that low differentiation individuals are intolerant of judgmental ambiguity (see Higgins et al., 1982). A positivity bias can function like a persistent active set to form clear, positive impressions of other people. As such, it should inhibit the activation and use of competing chronically accessible constructs. Higgins et al. (1982) found a significant effect of chronic accessibility on impressions for high and moderate differentiation subjects but not for low differentiation subjects.

In a study on the effects of self-discrepancies, Alexander and Higgins (1993) measured married subjects' chronically accessible self-discrepancies weeks before their first child was born. Affective measures were taken both weeks before and weeks after the child's birth. It was hypothesized that the new demands and child-care responsibilities associated with the new role of parent would have two distinct effects. First, the new demands would interfere with individuals attaining their own hopes and wishes and disrupt the interdependent goals of the couple. This should increase dejection-related suffering (e.g., disappointment, dissatisfaction) in parents who were high in ideal(own) discrepancy before becoming a parent. Second, because child-care obligations would now take center stage, there would be a shift in attention from the role of spouse and its distinct duties to the role of parent and its distinct duties. By involving the *same* type of regulatory system, the persistent set to fulfill the new parental responsibilities would

directly compete with the old spousal responsibilities and thus inhibit them. Such inhibition of ought(spouse) self-guides should *decrease* agitation-related suffering (e.g., worry, nervousness) in parents who were high in ought(spouse) discrepancy before becoming a parent. Both of these predicted effects were found.

Most recently, Brendl, Higgins, and Lemm, (1995) examined the inhibitory effects between chronically accessible self-discrepancies and persistent event orientations induced by framing. Subjects were asked to imagine either monetary gains or nonlosses (depending on framing condition) or to imagine monetary losses or nongains (again depending on framing condition). A psychophysical technique was used to measure subjects' discrimination among different amounts of gains or losses—a measure of sensitivity. Consistent evidence was found that discrimination was enhanced when the motivational features of the imagined event (i.e., their valence and their outcome focus) were either both the same as or both different from an individual's self-discrepancy. Discrimination was reduced when only one of the motivational features was different. The results were consistent with a model that postulated, in addition to excitatory effects, both inhibitory *and* disinhibitory effects.

There is another possible kind of inhibitory process as well. It could be that on exposure to a stimulus to which alternative constructs are applicable, constructs with higher excitation levels automatically inhibit competing constructs with lower excitation levels (see Higgins, 1989a). As discussed earlier, the results of accessibility studies in the existing literature could be explained by simply assuming that a primed or chronically accessible construct has a higher excitation level than alternative constructs and thus is tried or activated first. It is also possible, however, that when a stimulus increases the excitation levels of alternative constructs, such that these constructs become competing alternatives for activation, whichever construct reaches the highest level (from a combination of prior accessibility and current applicability) predominates at the moment and inhibits the activation of any alternative. If this were the case, then not only would the predominate construct be more likely to be activated, but any competing alternative would be *less* likely to be activated (compared with a baseline "no priming" condition). For example, if the construct "adventurous" were primed, followed by an ambiguous "adventurous/reckless" behavioral description, and then followed immediately by a Stroop task, would subjects be *quicker* to name the ink color of the printed word "reckless" because the construct "reckless" was temporarily inhibited? (And would this *not* occur for a vague "adventurous" behavioral description because, with only the construct "adventurous" being applicable, there would be no response competition?)

No direct test of this possibility has been conducted. Newman and Uleman (1990), however, report some initial findings that are consistent with this possibility. The study used cued recall for behavioral descriptions (i.e., the "spontaneous trait inference" paradigm of Winter & Uleman, 1984). In the normal priming condition, recall of an ambiguous behavioral description was actually worse when the retrieval cue was an applicable construct that

competed with the primed construct than when the same construct was a retrieval cue in a baseline, "no priming" condition.

APPLICABILITY AND SALIENCE

Knowledge activation does not depend solely on the accessibility of stored knowledge units. It also depends on the relation between these units and the stimulus information. Compared with the attention given to the theory-driven aspect of knowledge activation, relatively little attention has been paid to the contribution of the data. This section will review what little research and conceptualization there has been on this issue, and will discuss different relations between accessibility and stimulus information in general and between accessibility and salience in particular.

Applicability

Higgins, Rholes, and Jones (1977) compared the effect on subjects' responses to a behavior of priming stored constructs that could be used to characterize the behavior (i.e., "applicable" constructs) versus priming equally positive or negative constructs that could not be used to characterize the behavior (i.e., "nonapplicable" constructs). They found that priming constructs increased the likelihood that they would be used to characterize subsequent behaviors when they were applicable but not when they were nonapplicable (see also Sedikides, 1990, for a replication of this basic effect). Higgins and Chaires (1980) described a construct as being applicable to a stimulus if there was sufficient match between the features of the construct and the features of the stimulus (where the features are typically categorical). More generally, the greater the overlap between the features of some stored knowledge and the attended features of a stimulus, the greater is the applicability of the knowledge to the stimulus and the greater is the likelihood that the knowledge will be activated in the presence of the stimulus (see Higgins, 1989a).

As mentioned earlier, not all the features of a stimulus have equal weight in determining the applicability of the stored knowledge to the stimulus. In addition, not all the features of the stored knowledge have equal weight in determining applicability. Depending on situational variables and how the knowledge features themselves are represented (e.g., a feature's structural centrality), some knowledge features may be more important than others in determining applicability (see Herr et al., 1983; Higgins & Chaires, 1980). Herr et al. (1983), for example, suggest that the applicability of their primed extreme-animal exemplars was greater for the animals' size features than for their ferocity features. Studies by Banaji, Hardin, and Rothman (1993) have demonstrated the importance of specific applicability in knowledge activation and use. They found that a target person's social category feature had to match the social category feature of primed stereotyped trait constructs in order for the constructs to be used to judge the target's (somewhat vague) behaviors, even though the constructs otherwise were applicable to

the target and had been made highly accessible through recent priming.

Applicability should not be confused with consistency. A construct whose features are denotatively inconsistent with the features of a stimulus would still have greater applicability to that stimulus than a construct whose features were totally unrelated to that stimulus. Inconsistency implies at least some relatedness, some overlap of features. A "dishonest" behavior, for example, is inconsistent with the category "honest" but shares the dimension of honesty. Although a primed "inconsistent" construct would rarely be used directly to characterize a stimulus, it is more likely to be used as a standard and produce contrast effects than a primed "unrelated" construct.

Studies of accessibility effects have involved different types of applicability. These different types will be briefly described in the next section. In addition, this section will review different proposals for how accessibility and applicability interrelate and what factors influence applicability.

Types and Levels of Applicability

A behavioral description is "unambiguous" as stimulus information when only one alternative construct is applicable to it. The unambiguous descriptions used in previous research have been strongly related to one construct, such as the "extremely hostile" description that Srull and Wyer (1979, 1980) also used in their research. Unambiguous stimulus information needs to be distinguished from two basic alternatives—"ambiguous" stimulus information and "vague" stimulus information.

A behavioral description is "ambiguous" as stimulus information when at least two alternative constructs are equally applicable to it. The ambiguous descriptions typically used in past studies have been strongly and equally related to two alternative trait constructs, such as the ambiguous "persistent"/"stubborn" description used by Higgins et al. (1977). A different kind of ambiguity, a mixed description, was used by Smith (1989) that involved behavioral descriptions containing some information strongly implying one construct (e.g., intelligent) and other information strongly implying a different, evaluatively opposite construct (e.g., unfriendly). Another version of a mixed ambiguous description was used by Skowronski, Carlston, and Isham (1993). In their studies, the behavioral descriptions contained some information implying a construct at one pole of a trait-related dimension (e.g., unintelligent) and other information implying a construct at the opposite pole of the dimension (i.e., intelligent).

In contrast, a behavioral description is "vague" as stimulus information when no construct has more than weak applicability to it. The vague descriptions typically used in past studies have been weakly related to one trait construct, such as the moderately vague "hostile" description used by Srull and Wyer (1979, 1980), which they called "ambiguous." The fictional animals used by Herr et al. (1983) exemplify another kind of vague stimulus. Higgins and Brendl (1995) constructed an extremely vague behavioral description for "conceited" that received the same

conceitedness ratings as a baseline description that contained no conceited-related information at all.

There are two additional ways that stimulus information can relate to a particular construct. As discussed briefly earlier, a behavioral description is "inconsistent" as stimulus information (i.e., provides disconfirming or incompatible facts), when the construct is clearly *not* applicable to it. The adventurous/reckless stimulus description used by Higgins et al. (1977) in relation to their nonapplicable prime "listless" exemplifies this type of stimulus. A behavioral description is "contrary" as stimulus information (i.e., strongly implies an opposite alternative), when the applicability of a competing alternative construct is both strong and stronger than the target construct's applicability. Higgins and Brendl (1995), for example, constructed a behavior description that received not only high "self-confident" ratings but significantly higher "self-confident" ratings than "conceited" ratings, and thus was contrary as stimulus information (i.e., an evaluatively opposite alternative) for the target construct conceited. They found that when the stimulus information was contrary, subjects' judgments of conceitedness did not become stronger as the construct conceited increased in accessibility.

An important dimension distinguishing among these different types of applicability, although not the only dimension, is the amount of overlap between the features of a stored construct and the attended features of the stimulus. As the stimuli vary, a construct's level of applicability will vary. As stated earlier, the higher a construct's level of applicability to a stimulus, the greater is the likelihood that the construct will be activated in the presence of that stimulus. The Higgins et al. (1977) findings described earlier are consistent with this "applicability level" hypothesis. Srull and Wyer (1979, 1980) compared "hostile" (or "kind") behavioral descriptions that were unambiguous (called "hostile"), vague (called "ambiguous"), and inconsistent (called "nonhostile"). The applicability levels of these three types of descriptions ranged from high to moderate to very low, respectively. Consistent with the applicability level hypothesis, the subjects' hostility-related judgments increased as applicability level increased. In one study, Higgins and Brendl (1995) compared subjects' "conceited" judgments for "conceited" behavioral descriptions that were unambiguous and vague, which ranged from high to moderate in applicability level. In another study, the behavior descriptions were ambiguous, contrary, and vague, which ranged from high to low to extremely low in applicability level. In both studies, there was a strong applicability effect, with stronger conceited-related judgments for higher levels of applicability.

Proposals for Applicability Effects

In comparison with the proposals for accessibility effects, the proposals for applicability functions have been less developed. Still, there are some clear differences in conceptualizations of applicability. In Wyer and Srull's (1986, 1989) bin model, whichever knowledge unit is on the top of the semantic bin is activated and then tested first for its

applicability to the stimulus information. If the features of a construct resemble those of the stimulus enough to justify the conclusion that the stimulus exemplifies the construct, then the construct is used to encode the stimulus. Otherwise, the executor performs a top-down search of the semantic bin until the first applicable construct is found.

There are four noteworthy aspects of this proposal, two explicit and two implicit. One explicit aspect of the proposal is that knowledge activation and applicability are independent. A stored construct is retrieved from the semantic bin as a function of the hierarchical positions of its copies. After the construct is retrieved, it is tested for its applicability to the stimulus. Applicability determines whether the construct is used to encode the stimulus, but it does not influence whether the construct is activated. A second explicit aspect of the proposal is that applicability is a serial process. The construct on top of the bin is tried first. If it is not applicable, a top-down search retrieves the next candidate, and so on until an applicable construct is found.

An implicit aspect of the search process is that the direction of the applicability process is from activated stored knowledge to stimulus information. The construct on top of the semantic bin is retrieved, or is already in the work space, and then is compared with the stimulus to test whether its features are enough like those of the stimulus to justify the conclusion that the stimulus exemplifies the construct. Thus, the stimulus information in this proposal functions as a target. This leads to another implicit aspect of the proposal. If the stimulus information functions as a target, then one would expect that the proposed test for "sufficient similarity" would be like that described in previous inferential or hypothesis-testing models of assigning objects to categories (e.g., Smith & Medin, 1981; Tversky & Gati, 1978).

Wyer and Srull (1989, pp. 258–260) describe a process of inferring category membership that is like these hypothesis-testing models, but less Bayesian in nature. They suggest that a conditional inference is computed of the likelihood that, assuming something is a category member, it would have the features possessed by the stimulus (e.g., assuming that a person is an engineer, what is the likelihood that he or she would like beer and science fiction movies [or wine and poetry] like the stimulus person does?). They assume that judging a stimulus as being a category member involves identifying the category's prototypic features and comparing the features of the stimulus with these category features, or searching for features of the stimulus that define or characterize members of the category. The critical characteristic of this model and hypothesis-testing models more generally is that features of the stimulus that disconfirm category membership would be taken into account as well as confirming features (although they might be given less weight). That is, both matches *and* mismatches between the features of the construct and those of the stimulus are involved in the applicability process.

The conceptualization of applicability proposed by the synapse model (see Higgins, 1989a) differs from these

models in various respects. The normal direction of the applicability process in the synapse model is from stimulus information to stored knowledge. With the stimulus information functioning as stimulant, it is only the matches between the features of the stimulus and those of the stored constructs that matter. Although not stated explicitly, applicability would be a parallel rather than a serial process. All the alternative constructs that shared features with the stimulus would be stimulated by it. The greater the number of shared features, the more the stimulus would contribute to a construct's excitation level. In the synapse model, then, applicability is *not* independent of knowledge activation. By contributing to excitation level, applicability influences knowledge activation together with accessibility.

The synapse and bin proposals make similar predictions for ambiguous stimulus information where two (or more) constructs have features that are strongly related to the stimulus. From the bin perspective, whichever construct is tested first is likely to be used to categorize the stimulus given that both constructs will meet the rules for category assignment. From the synapse perspective, whichever construct has a higher excitation level is likely to be used to categorize the stimulus given that the stimulus will contribute the additional stimulation needed for activation. The two proposals make different predictions for *vague* stimulus information, however. From the bin perspective, even a highly accessible construct that was activated and tested first in relation to a stimulus would not be used to categorize an extremely vague stimulus that had almost no feature similarity to the construct. From the synapse perspective, however, stimulus information can be extremely vague in relation to a particular construct and the construct will still be used to categorize the stimulus as long as the accessibility of the construct compensates for the low stimulant value of the stimulus. The results of the Higgins and Brendl (1995) study are consistent with the latter proposal.

Saliency

As discussed earlier, not all the features of a stimulus receive equal attention at any point in time. Prior goals or expectancies can create an active set in a perceiver that will direct more attention to some features of a stimulus than to others. In hypothesis testing, more attention might be given to hypothesis-matching features than hypothesis-mismatching features of a target stimulus (for a review, see Higgins & Bargh, 1987). Some features or exemplars of a construct might be more typical than others, and those features of a stimulus that match the more typical properties of a construct are likely to receive more attention and have more influence on construct activation (see Smith & Medin, 1981; Smith & Zarate, 1992). Some features or exemplars of a construct might also be more frequently activated than others creating differences in accessibility that influence selective attention (Smith & Medin, 1981). Thus, various *preexposure* characteristics of stored knowledge and knowledge activation can influence which features of a stimulus will receive more attention. Saliency is another variable that influences which

features of stimulus information will receive more attention, but saliency has to do with something about the stimulus event itself that occurs *on* exposure without a prior set for a particular kind of stimulus. To appreciate the distinct contribution of saliency, consider the case where stimulus information is functioning as a stimulant. Those features of the stimulus that match the features of stored constructs have the potential to increase the excitation levels of those constructs. But it is the salient features that will actually contribute most to increasing the constructs' excitation levels. Thus, different constructs might reach activation threshold sooner and be used to judge the stimulus depending on whose features match the stimulus features that happen to be salient at the moment of exposure.

Saliency, then, is something about a stimulus event that occurs upon exposure to the stimulus, without a prior set for a particular kind of stimulus, that draws attention selectively to a specific object of perception or thought. Newcomb et al. (1965) similarly described salient information as information at the forefront of perceivers' conscious thought that is prompted by the immediate situation in which they find themselves. An object of perception or thought can be salient either because of something about its absolute properties (i.e., its *natural prominence*) or because of something about its properties in comparison with the properties of other objects in the immediate situation (i.e., its *comparative distinctiveness*).

Effects of Natural Prominence

The extent to which an object (or set of features) is conspicuous or stands out is always influenced by its relation to the immediate surroundings. Thus, the saliency of an object is never absolute and prominence is never fully independent of comparative distinctiveness. Still, there are relatively invariant natural properties of objects that increase the likelihood that attention will be drawn selectively to them. Classic early studies revealed that bright, moving, or "complex" objects were more likely to draw attention selectively (e.g., Berlyne, 1958; Koffka, 1935; Titchener, 1908/1966). McArthur and Post (1977) manipulated these properties for social objects. Subjects observed a getting-acquainted conversation between two actors. One study manipulated brightness by having one actor seated in a bright light and the other actor seated in a dim light. A second study manipulated movement by having one actor seated in a rocking chair and the other actor seated in a regular chair. A third study manipulated "complexity" by having one actor wear a shirt with a black and white pattern and the other actor wear a solid grey shirt (with reflectance equated). In these studies, *both* interacting actors were distinctive relative to each other. Thus, the actors differed in natural prominence but not in distinctiveness.

These studies found that the behavior of an actor was attributed more to situational causes when that actor's interaction partner (functioning as the situation) was more salient. Arkin and Duval (1975) also found that an actor's choice of art works was attributed more to the situation (the art works) when the situation was dynamic and changing (a videotaped portrayal of the art works) than

when the situation was stable (photos of the art works). More direct evidence of salience effects on attention *per se* was found by McArthur and Ginsberg (1981), who report that the duration of subjects' selective visual fixation was greater when an actor on a color TV monitor wore a striped shirt versus a simple grey shirt, or had visible red hair versus being shown on a black-and-white TV monitor.

The natural properties in these studies that made some objects salient varied in the extent to which they were permanent properties of the objects. Although the "complexity" of the salient shirt was relatively permanent, the brightness of the salient actor was a momentary effect of being seated in a bright light. Like the bright light manipulation, other studies have manipulated something about the momentary properties of an object in the immediate situation that increased its natural prominence. A strong form of such manipulations is to vary the natural ease or likelihood of attending particular objects (see, e.g., Duval & Wicklund, 1972; Fiske, Kenny, & Taylor, 1982; Taylor, Crocker, Fiske, Sprinzen, & Winkler, 1979; Taylor & Fiske, 1975).

In Taylor and Fiske's (1975) first study, for example, six observers watched a dyadic interaction in different seating positions where two observers could see the behaviors of one dyadic partner more easily, two other observers could see the behaviors of the other dyadic partner more easily, and the remaining two observers could see the behaviors of each dyadic partner equally easily. When one dyadic partner could be seen more easily (i.e., was more salient), that partner was perceived to have been more influential in the interaction. In a study of objects of thought rather than objects of perception, Salancik and Conway (1975) used a situational linguistic manipulation to vary the ease (and likelihood) of retrieving different samples of a subject's past behaviors such that different information became consciously available for attention. This manipulation of selective attention to particular objects of thought had a strong effect on subjects' subsequent evaluative judgments.

Effects of Comparative Distinctiveness

Something about an object's properties in comparison with the properties of other objects in the immediate situation can also draw attention selectively to that object. In the typical experiment, subjects are presented with an array of objects in which one object has a property that differs from most or all of the other objects. It is this comparative distinctiveness that draws attention selectively to the object and thus makes it salient in the immediate context. Taylor, Fiske, Close, Anderson, and Ruderman (1977), for example, varied the race or sex composition of a group of adults having a discussion. In different conditions, there was a solo black male in an otherwise white male group, a solo white female in an otherwise white male group, and a solo white male in an otherwise white female group. Taylor et al. (1977) found that subjects who heard each discussion while viewing slides of the participants judged the solo individual as being more influential than the other individuals in the group. Using a similar group composition paradigm, Biernat and Vescio (1993) describe other effects

of comparative distinctiveness on social judgment and memory of target persons from different social categories (see also McCann, Ostrom, Tyner, & Mitchell, 1985; cf. Oakes & Turner, 1986).

There is also evidence that the distinctiveness of one's own properties in comparison with the properties of others increases the salience of those personal properties. McGuire and Padawer-Singer (1976) found, for example, that an elementary school student was more likely to include a particular personal attribute (e.g., green eyes, foreign birthplace) in his or her spontaneous self-description when that attribute was relatively distinctive in that child's classroom. The distinctiveness of an attribute within a momentary group also makes possession of that attribute salient. And the salient attribute, such as distinctive gender, can influence how an individual processes additional attribute-related information about self or others, such as gender-stereotypic information (see Higgins & King, 1981).

An object can also be distinctive because it is comparatively rare in each of a variety of contexts or settings. Members of minority groups within a community are relatively distinctive for this reason. Likewise, negative behaviors are comparatively distinctive because they tend to be less common in many different contexts or settings. Hamilton and Gifford (1976) examined how this source of salience can produce "illusory correlations." There were two abstract groups of people, Group A and Group B, and two classes of behaviors, positive or negative. The proportion of positive and negative behaviors for the members of each group was the same. In different conditions, however, either Group A or Group B was made distinctive by having less group members, and either the positive or negative behaviors were made distinctive by occurring less frequently. Hamilton and Gifford (1976) found that the distinctive group and the distinctive class of behavior were perceived to have the strongest co-occurrence, consistent with the prediction that this unit or combination would receive the most attention.

Contextualized "novelty" as a condition of salience can involve distinctiveness in a variety of forms. Some property of an object, such as a person's gender, can be novel because the other objects in the immediate situation possess a different property in common. The shared property could be perceptual, such as possessing a particular skin color, or it could be more functional, such as being a living room object. For example, the salience of a cow versus a couch would vary depending on its location in an apartment versus a pasture. Whether perceptual or functional, the object's property would be distinctive in comparison to the common property possessed by the other objects in the immediate situation.

It is also possible for the immediate situation to activate stored objects associated with that situation. A particular object in the immediate situation might then be distinctive compared with these mental objects. Indeed, in developmental studies of dishabituation or "release from proactive inhibition" (see Wickens, 1972), selective attention was found to be greater for a new stimulus that was categorically distinctive compared with the immediately preceding stimuli (i.e., the mental objects in the immediate situation)

than for a new stimulus that shared the categorical property of the preceding stimuli (e.g., Bornstein, 1979; Ross, 1980; cf. Fiske, 1980).

The "How?" of Salience Effects

Two major types of variables have been examined in the literature as potential mediators of salience effects—*amount of recall* and *experienced ease of retrieval* (or "availability"). In addition to these variables, this section will consider two additional possible mediators—*applicability effects* and *salience experiences*.

Amount of Recall

Taylor and Fiske (1975) suggested that because salient stimulus information received relatively more attention, it might be better recalled, and greater overall recall of the salient information could increase its impact on subsequent responding. Although there is some evidence that salient information is better remembered, this is not a strong or reliable effect (see reviews by Fiske & Taylor, 1991; McArthur, 1980, 1981; Taylor & Fiske, 1978). Moreover, salience effects have been found to be independent of the mere quantity of salient information recalled (e.g., Fiske, Kenny, & Taylor, 1982; Taylor & Fiske, 1975).

Fiske, Kenny, and Taylor (1982) proposed that among the salient stimulus information that is better recalled overall, one subset might be particularly relevant for subsequent responses. And it is recall of this subset of relevant information, and not total recall, that functions as the mediator. In making judgments of causal influence, for example, salience effects would be mediated by better recall of just that subset of stimulus information which is causally representative (see Kahneman & Tversky, 1972) rather than the overall quantity of information recalled. Fiske et al. (1982) found some support for this proposal, especially for recall of visual information. They also noted, however, that the salience manipulation effect on attention had an impact on subsequent responses that was independent of this "relevant recall" mediator.

Experienced Ease of Retrieval

Pryor and Kriss (1977) proposed an alternative mediator for salience effects. They suggested that the reason increased attention from salience had an impact on subsequent responses was not because greater attention increased the amount of recall but because greater attention increased the ease of later retrieval. And applying Tversky and Kahneman's (1973) availability heuristic, they proposed that the greater availability of salient stimuli (i.e., the greater ease of retrieving previously salient stimuli) could make them seem more likely as agents of influence (see also McArthur, 1981; Taylor & Fiske, 1978). To test this, they manipulated the salience of either the person or the entity in a person-entity sentiment relation by varying which appeared first as the topic of the sentence (e.g., "Joe likes the film" versus "The film was liked by Joe"). In one study, they found that subjects' recognition responses for the person relative to the entity were quicker when the person versus the entity had been made

linguistically salient, and in another study that subjects were more likely to make entity attributions when the entity versus the person was linguistically salient. These studies provide some support for the notion of availability (i.e., experienced ease of retrieval) as a mediator in salience effects, but because the measures of such availability and of attribution were obtained in separate studies, they do not provide direct evidence for the mediational role of availability.

Applicability Effects

Higgins (1989a) suggested another possibility for how salience influences subsequent responses. It is not so much that salient objects receive much more attention than nonsalient objects but rather that salient objects receive sufficient attention to permit judgments of them, whereas nonsalient objects do not (see also Taylor & Thompson, 1982). As stated earlier, applicability concerns the extent of overlap between the features of some stored knowledge and the *attended* features of a stimulus. And the greater the applicability of a knowledge unit to the stimulus, the greater is the likelihood that the knowledge will be activated in the presence of the stimulus. It is the attended features of a stimulus event, therefore, that determine applicability and influence knowledge activation. Salience, then, can impact subsequent responses by influencing which features of a stimulus event receive attention, and this in turn will influence which stored knowledge units are likely to be activated in the immediate situation.

This "salience-applicability-knowledge activation" chain could impact on subsequent responses for a couple of reasons. One reason is that this chain could influence how a stimulus event is initially judged or identified, and initial judgments have been shown to influence subsequent responses (for a discussion of encoding effects, see Higgins & Stangor, 1988; Wyer & Srull, 1989). Thus, salience would impact on subsequent responses because of differences in the initial judgments of salient versus nonsalient information rather than because of differences in the episodic recall of salient versus nonsalient information. For example, even though perceivers in Taylor and Fiske's (1975) first study might have stored the detailed behaviors of both interactive partners and thus could later recall these episodes, they might have made and stored judgments of mainly the behaviors of the salient partner. And it might be the storage of these initial judgments rather than episodic memory for the specific behaviors that increased the strength of subsequent judgments of causal influence (cf. Taylor & Thompson, 1982).

The "salience-applicability-knowledge activation" chain could also influence subsequent responses because the activation of particular knowledge units could increase the accessibility of those units, and, as reviewed earlier, knowledge accessibility influences subsequent responses. Moreover, this relation between salience and accessibility could have chronic as well as relatively temporary effects if the characteristic of the immediate situation that produces salience is stable over time or even becomes institutionalized. For example, the aforementioned effect on spontaneous self-descriptions of children's distinctiveness in their

classroom (McGuire & Padawer-Singer, 1976) probably involves such a "stable salience-chronic accessibility" relation. Regarding institutional stability, symbolic interaction models describe normative situations in which particular roles or social objects are chronically salient (e.g., McCall & Simmons, 1978; Stryker, 1980; see also Stryker & Statham, 1985). The salience of these social objects increases the likelihood that particular knowledge units, such as specific aspects of self, will be repeatedly activated, which increases their chronic accessibility. This chronic accessibility, in turn, can have a major impact on how people respond in social situations.

As suggested in the earlier discussion of inhibitory processes, increasing a stored construct's excitation to a level higher than competing constructs may result in it inhibiting the activation of the alternative constructs. Salience, then, might not only increase the likelihood that a construct is activated, which would increase that construct's *subsequent* accessibility, but prior to activation it could inhibit the activation of currently accessible constructs that compete with the construct made predominant by salience. Thus salience, as well as other sources of selective attention, could actually inhibit accessibility effects. Once again, this points up the need to distinguish "perceptual readiness" from accessibility.

Salience Experiences

Another possible mediating variable of salience effects is people's experiences of salience and the inferences they draw from these experiences. That is, there might be effects of salience experiences that are analogous to the effects of accessibility experiences described earlier. When some aspects of a stimulus event are more salient than other aspects, people experience their attention being drawn to and held by particular objects. People might spontaneously make inferences about the sources of this experience, and they might infer sources beyond or other than the natural prominence or comparative distinctiveness of particular objects. These inferences might then be used in their subsequent responses. For example, in the studies of solo group members by Taylor et al. (1977), subjects might have inferred that the powerful, dynamic forces expressed in the solo member's appearance and behaviors accounted for his or her attention-grabbing nature. This inference, in turn, would have increased the likelihood that subjects would subsequently evaluate the solo member as being more influential than others in the group. That is, an "attention-grabbing" heuristic might have been used to infer importance.

A finding in the literature is especially interesting in light of this proposed mediator. According to this proposal, it is inferences about potential sources of the salience experience that impact subsequent responding. It would be reasonable for people to infer, as in the preceding example, that an individual's forceful expressions make him or her attention-grabbing, which would imply that the individual was influential in that immediate situation. But it would not be reasonable for people to infer that an individual's underlying dispositions make him or her attention-grabbing. After all, underlying dispositions

per se are rarely directly observable and cannot themselves grab attention (cf. McArthur, 1980). This could account for why the literature has rarely found evidence that salience affects perceivers' dispositional attributions (e.g., McArthur & Post, 1977; Pryor & Kriss, 1977; Taylor & Fiske, 1975; cf. Strack, Erber, & Wicklund, 1982).

CONSEQUENCES OF KNOWLEDGE ACTIVATION VARIABLES AND THEIR INTERRELATIONS

This chapter has emphasized the conditions of knowledge activation and use rather than their consequences. This focus has the benefit of considering accessibility, applicability, and judged usability as basic principles, which is consistent with the theme of this book. It is also important, however, that the reader appreciate the significant contribution of these variables to social psychological phenomena. Such appreciation might best be obtained by illustrating how these variables influence people's responses to life events. Thus, this last section of the chapter will briefly consider how accessibility, applicability, and judged usability influence attention, memory, feelings, judgment, and action. Another purpose of this section is to illustrate how these effects depend on the *interrelations of these different variables*. Other reviews are also available that illustrate the significant impact of knowledge activation variables on information processing and social behaviors, including stereotyping and discrimination, emotional responses, prosocial and antisocial behavior, and social interaction (e.g., Fiske & Taylor, 1991; Sedikides & Skowronski, 1991; Wyer & Srull, 1989).

Attention

The higher the accessibility of a stored knowledge unit, the more likely that stimulus information related to that knowledge will receive attention. Evidence for this relation between accessibility and attention has been found in several studies. Bargh and Pratto (1986) demonstrated that individuals who possessed high chronic accessibility for a particular trait construct were less able to ignore a word related to that construct when it appeared in a Stroop task whose goal was to attend to and name only the color of each word presented.

Roskos-Ewoldsen and Fazio (1992) found accessibility effects on selective attention (and incidental learning) from both individual differences in chronic accessibility and contextual priming. In one study, for example, arrays of pictorially represented objects were briefly presented to subjects. The investigators found that individuals whose chronic attitude accessibilities for particular objects (e.g., squirrels, bicycles) were relatively high were more likely to notice those objects compared with objects for which they had less accessible attitudes. In another study, the attitude accessibility of some objects was temporarily increased by having subjects make repeated like-dislike evaluations of only these randomly selected objects. (To control for simple exposure to the pictures themselves,

subjects made animate-inanimate judgments of the comparison objects.) This study also demonstrated that higher attitude accessibility was associated with more objects being noticed.

Memory

Results such as those previously reviewed indicate that stimulus information related to a more highly accessible construct is likely to receive more attention. Given this, one would also expect that the higher the accessibility of a stored knowledge unit, the more likely that stimulus information related to that knowledge will be remembered. This was the reasoning behind the studies by Higgins et al. (1982) described earlier. As previously noted, Kelly (1955) had suggested that personal constructs, which would be chronically accessible, might function like a kind of projected scanning pattern that would pick up "blips of meaning" (p. 145) as it was swept back and forth across the perceptual field. If so, then stimulus information related to these more chronically accessible constructs should be remembered better than stimulus information related to less accessible constructs. The results of the Higgins et al. (1982) studies support this conclusion.

Kelly's (1955) more provocative notion of "personal viewpoints" corresponds to the Gestalt perspective that psychological phenomena occur as part of a system of mutually interdependent elements that has a meaning and importance *as a whole*. An example of such a system would be the chronically accessible "actual self-desired self" linkages that constitute self-discrepancies. As described briefly earlier, Higgins and Tykocinski (1992) examined whether persons who possess different types of self-discrepancies are sensitive to different types of psychological situations and thus would remember better events related to these psychological situations. All subjects read the same essay about the life of a target person in which events reflecting different types of psychological situations occurred (e.g., finding money on the street [presence of a positive]; escaping an unpleasant school day because of an election [absence of a negative]). As predicted, a target person's experiences that reflected the presence and absence of positive outcomes were remembered better by predominant actual-ideal than by predominant actual-ought discrepancy persons, whereas the reverse was true for experiences that reflected the absence and presence of negative outcomes.

Contextually priming self-discrepancies can also affect memory. In one study, for example, Higgins et al. (1994) contextually primed either ideal or ought self-guides by having subjects either think about their hopes and wishes, or think about their duties and obligations, respectively. All subjects read the same essay about the life of a target person in which events reflecting either an approach or an avoidance predilection occurred (e.g., "Because I wanted to be at school for the beginning of my 8:30 psychology class which is usually excellent, I woke up early this morning [approaching a match to a desired end-state]"; "I wanted to take a class in photography at the community center, so I didn't register for a class in Spanish that was scheduled at the same time [avoiding a mismatch to a

desired end-state]"). As predicted, a target person's experiences that reflected an approach predilection was remembered better when the ideal self-guide was primed than when the ought self-guide was primed, whereas the reverse was true for experiences that reflected an avoidance predilection.

There is also evidence suggesting that interrelations among *accessibility, salience, and judged usability* can influence memory. Higgins and King (1981) describe a study that manipulated the sex composition of the groups in which subjects participated and found that sex-stereotypic reproduction of target person information was reduced in groups where a subject was the only member of his or her sex (sole) compared with groups where a subject's sex was in the majority. Higgins and King (1981) suggested that the undergraduate subjects were likely to have both traditional stereotypic beliefs and modern anti-stereotypic beliefs available, but that the former beliefs were likely to be more chronically accessible. In the absence of contextual influences on knowledge activation, the stereotypic beliefs were more likely to be activated and influence encoding and memory. But sex distinctiveness or salience was likely to increase the accessibility of modern beliefs, thus increasing the likelihood that these modern beliefs would be activated and influence judged usability of the stereotypic beliefs.

Feelings

The higher the accessibility of a stored knowledge unit, the more likely that a stimulus event will produce feelings related to that knowledge. For example, a study by Strauman and Higgins (1987) mentioned earlier contextually primed desired-self attributes by having subjects complete sentences describing other people that began with a desirable attribute (e.g., "A friendly person is _____"). The study found that priming desired-self attributes from which subjects' actual-self attributes were discrepant automatically produced dejection-related emotions such as "sad" or agitation-related emotions such as "worried" depending on whether the desired-self attributes were part of a primed ideal self or a primed ought self, respectively.

The literature on feelings also illustrate how accessibility effects depend on interrelations between accessibility and other variables. In a study by Sinclair, Hoffman, Mark, Martin, and Pickering (in press), for example, the subjects were contextually primed with either emotionally positive scrambled sentences, emotionally negative sentences, or neutral sentences. The positive and negative primes were selected to have high arousal content (e.g., "knife slashed the stabbed"). After being primed, the subjects either exercised or sat still to vary physiological arousal. They then rated their emotional state either immediately or after a delay. The study found that subjects in the exercise plus delay condition made extreme positive or negative emotional self-ratings when they were positively or negatively primed, respectively.

That this priming effect on emotional self-ratings was stronger in the exercise than in the no-exercise (delay) condition demonstrates the importance of *applicability* because the arousal-related constructs that were primed

would be more applicable to the subjects' feeling state in the exercise condition than in the no-exercise condition. In addition, that this priming effect on emotional self-ratings was stronger in the delay than in the immediate (exercise) condition demonstrates the importance of *judged usability* because the accessibility of arousal-related constructs would be more likely to be attributed to the exercise, instead of an emotional state, in the immediate than in the delay condition.

The role of applicability in accessibility effects on feelings is also evident in Houston's (1990) studies of empathy. Subjects were selected who varied in their chronically accessible "actual self-desired self" linkages. Some subjects possessed predominantly actual-ideal discrepancies, whereas other subjects possessed predominantly actual-ought discrepancies. Houston (1990) found that the subjects showed greater empathy for the emotional suffering of other people when the emotional significance of their chronically accessible discrepancies was applicable to the specific feelings suffered by the target person.

Additional evidence of the role of judged usability in accessibility effects on feelings is provided by Wegner, Erber, and Zanakos's (in press) research on ironic effects of mood control. In one study, some of the subjects were instructed to try to change their mood in a positive direction, such as trying not to be sad after sad personal episodes had been made accessible. Compared with subjects who were given no instruction to change the mood produced by the accessible episodes, the subjects who tried to change their mood positively actually became *less* happy when also given a cognitive load (i.e., remembering a 9-digit number). One explanation of this ironic effect of mood control is that subjects consciously monitor their sad thoughts to control them, which increases their accessibility. With sufficient mental resources, subjects' judgment that sad thoughts are inappropriate (i.e., have low usability given the instructions) is effective in inhibiting the activation of the highly accessible sad thoughts. But when the mental resources are insufficient, as in the cognitive load condition, then the high accessibility of the sad thoughts produces sad feelings despite judged usability.

Judgment

The higher the accessibility of a stored knowledge unit, the more likely that stimulus information will be judged in relation to that knowledge. Of all the consequences of knowledge activation, the effects of accessibility on evaluative, descriptive, and explanatory judgments have been investigated most extensively in the literature. Most of this chapter's earlier review of the principles of knowledge activation have described accessibility effects on judgments. Moreover, the social cognition literature has described how judgments themselves influence subsequent memory, feelings, decisions, and so on (e.g., Fiske & Taylor, 1991; Wyer & Srull, 1989). Thus, it is unnecessary for me to describe accessibility effects on judgments or the consequences of these judgments. Instead, I will use the literature on judgments to illustrate some additional effects of the interrelations among accessibility, applicability, and judged usability.

Higgins and Brendl (1995) directly tested the following "activation rule" concerning the interrelation between *accessibility and applicability*: "Strong accessibility can compensate for weak applicability." This activation rule has a long history. Wertheimer, for example, proposed that past experience "fills in" information missing in external stimuli (see Heidbreder, 1933). And, more directly, Bruner (1957a, 1957b) proposed that the more easily or quickly a given stimulus can be coded in terms of a particular category, the less stimulus information is necessary for categorization to occur in terms of the category. This activation rule is based on the notion, discussed earlier, that both accessibility and applicability contribute to the likelihood that the excitation level of a construct will reach activation threshold. The greater the accessibility of a construct, the higher its level of excitation. The greater the applicability of a construct to the attended features of a stimulus, the more the excitation level of the construct increases on exposure to the stimulus. Finally, the higher a construct's excitation level, the greater the likelihood that the construct will be activated and used.

Together these propositions imply that a higher excitation level from stronger accessibility can compensate for a decreased contribution to excitation level from lower applicability. Higgins and Brendl (1995) found that even when a stimulus description was extremely vague, that is, very weakly related to "conceited," the target person's behaviors yielded conceited-related spontaneous impressions when the accessibility of the construct "conceited" was maximized—contextual priming of conceited (without awareness), short priming-to-stimulus delay, and relatively high levels of chronic accessibility for conceited.

This assimilation effect of accessibility on judgment, like all such accessibility effects on judgment, occurs because applicability under normal circumstances contributes substantially to knowledge activation. People naturally assume, for example, that a category comes to mind in the presence of a stimulus because the stimulus is a member of the category. This is a reasonable and adaptive assumption. Indeed, a loss of this assumption would create an existential crisis. But sometimes, as in the Higgins and Brendl (1995) study, the contribution of the stimulus to knowledge activation, compared with the contribution of prestimulus accessibility, is rather minimal. Still, people function as if the stimulus contributed substantially to knowledge activation and judge the stimulus to be related to this knowledge. Thus, assumed applicability lies at the heart of accessibility effects on judgment.

The extremely vague stimulus description used in the Higgins and Brendl (1995) study was no more related to the construct "conceited" than a baseline, "no information" description. Still, when accessibility was high enough, spontaneous judgments of conceited were made. As an investigator, I wonder what the subjects without high accessibility for conceited would have thought of the high accessibility subjects' judgment that the target person was conceited. From their perspective, this judgment is not supported by the extremely vague stimulus information. This situation is similar to that of a patient in a mental hospital who judges an event in a way that makes little sense to others who observed the same event. These

other people may conclude that the person making the judgment is "delusional." In intergroup relations, a member of one group may judge the behavior of an outgroup member in a way that makes little sense to others who observed the event, and the others may conclude that the person making the judgment is "prejudiced." In such cases, the person making the judgment is considered to be *oversensitive* in some way, and is thus labeled negatively. Understanding the principles of knowledge activation could shed light on the conditions that produce such judgments.

Higgins and Brendl (1995) also found evidence for another "activation rule" concerning the interrelation between accessibility and applicability: "The relation between higher accessibility and stronger judgments is constrained when the applicability of a competing alternative construct is both strong and stronger than the target construct's applicability" (p. 240). For a contrary stimulus in their study, that is, a stimulus both strongly related to "self-confident" and more so than to the target construct "conceited," they found that neither higher levels of chronic accessibility for conceited nor contextual priming of conceited influenced subjects' conceited-related judgments. Such constraints on accessibility effects are adaptive; for example, the high chronic accessibility of the construct "stick" that occurs from high exposure to sticks does not overwhelm people's ability to identify a snake that is right in front of them.

Higgins and Brendl (1995) also tested the following "activation rule" concerning the interrelation between *accessibility, accessibility experiences, and judged usability*: "Higher accessibility can yield stronger judgments even when perceivers are aware of priming events if the additional contribution to activation from applicability and chronic accessibility is sufficiently great" (p. 239). For an ambiguous stimulus description, that is, a stimulus strongly and equally related to both conceited and self-confident, Higgins and Brendl (1995) found that subjects with higher levels of chronic accessibility for conceited made stronger conceited-related judgments of the target's behaviors even when they remembered the priming events (i.e., within the contextual priming with recall condition).

As discussed earlier, it has been suggested that subjects who are aware of the priming events when making a later judgment consider it inappropriate to use the primed construct in their judgment and thus suppress it (see Martin & Achee, 1992). Active suppression of the prime-related judgment may be more likely when the priming events are perceived as a sufficient explanation for the construct being activated in the presence of the stimulus. But when the stimulus is ambiguous, and thus contributes to the construct's excitation level, and when high chronic accessibility also contributes to the speed and ease of construct activation, that is, to the accessibility experience, the priming events alone may no longer be perceived as a sufficient explanation for the experience (cf. Strack, 1992). Active suppression of the judgment related to the highly accessible construct should be less likely in this case.

The preceding activation rule might also underlie some findings in Banaji, Hardin, and Rothman's (1993) research

on implicit stereotyping in person judgment. Subjects judged behaviors that were moderately related to a target construct. In one study, the target construct was "dependent" and it was either contextually primed or not primed. The social category of the target person, either male or female, was also manipulated. The subjects made stronger dependent judgments when dependent was primed *and* the target person was female, that is, was a member of a social category stereotypically associated with the trait "dependent." Most important, this accessibility effect was found even for subjects who recalled the priming events. The chronically accessible stereotype, when applicable to the target person, might have contributed so much additional speed and ease of construct activation that the priming events were not perceived as causing the construct activation. Thus, the judgment related to the highly accessible construct was not suppressed. In general, the results of the Higgins and Brendl (1995) and Banaji et al. (1993) studies suggest that chronically accessible constructs can contribute to a level of excitation so much higher than priming could have produced that an assimilation effect of accessibility is obtained even for perceivers who remember the priming events.

A different kind of interrelation between accessibility and judged usability has been examined recently in relation to Kruglanski's (1989) theory of motivated cognitions (see also Kruglanski, Chapter 17, this volume). Both Ford and Kruglanski (in press) and Thompson et al. (1994) found that the usual accessibility effect on judgment from contextual priming was *not* obtained when subjects had a high accuracy motivation or, more generally, a high need to avoid closure. Similarly, Schuette and Fazio (in press) found that frequent contextual priming of an "attitude object-evaluation" linkage did *not* influence subsequent responses when subjects had high "fear of invalidity." The results of these studies suggest that a high accuracy motivation influences people's judgment of the appropriateness of using the "first thing that comes to mind." Knowledge use depends on the knowledge meeting some accuracy criteria rather than its speed or ease of activation or its simple primacy. If accessible knowledge does not meet the accuracy criteria, it is judged inappropriate to use and is suppressed.

Research by Gilbert and Hixon (1991) exemplifies how the same variable, "cognitive busyness," can influence the relation between accessibility and both applicability and judged usability. They found that when subjects engaged in an additional task requiring mental effort, a chronically accessible stereotype was less likely to be activated, but if activated, it was more likely to be used in making judgments of a member of the stereotype category. First, cognitive busyness might have taken resources away from processing the stimulus information, thereby decreasing the impact of the stimulus as stimulant, which would reduce the contribution of applicability to knowledge activation. Apparently, the chronic accessibility of the stereotype (about Asian American students) was not sufficient to produce activation when applicability was reduced. Second, when activation was allowed to occur by delaying the cognitive busyness until after the

activation phase, the cognitive busyness might have taken resources away from judging whether using the stereotype was appropriate. Thus, stereotype suppression from judged inappropriateness would be reduced, which would increase the likelihood that the stereotype would be used to judge the target.

Evidence also suggests that interrelations among *accessibility*, *salience*, and *judged usability* can influence judgments. Similar to the results of the Higgins and King (1981) study described earlier, Biernat and Vescio (1993) found that subjects were less likely to make racial stereotypic judgments of target women when a target was the sole member of her race in the depicted group than when a target was part of a race-balanced group. Following the reasoning of Higgins and King (1981), Biernat and Vescio (1993) suggested that the traditional race stereotypic beliefs, having higher chronic accessibility, dominate judgment in the balanced groups, but the greater salience of race in the solo groups increases the relative accessibility of modern race beliefs, thus making it more likely that these modern beliefs will be activated and influence judged usability of the stereotypic beliefs.

Behavior

In an early test of accessibility effects on performance, Higgins and Chaires (1980) gave subjects the Duncker (1945) candle problem. To solve this problem, subjects must recognize that a box filled with tacks can be used as a platform for the candle rather than just being a container for the tacks. Higgins and Chaires (1980) found that subjects were much more likely to solve this problem if, as part of a prior "unrelated study," they had been contextually primed with differentiated linguistic constructions (e.g., "tray and tomatoes") instead of undifferentiated linguistic constructions (e.g., "tray of tomatoes"). The contextual priming increased the likelihood that the box would be differentiated from the tacks. Applicability also mattered because priming with "and" as a general grammatical connective (e.g., "seashells and rocks glistened in the sunshine"), rather than specifically differentiating a container from its content, had no effect on performance.

The combination of accessibility and judged usability has also been shown to influence behavior. There is evidence that aggression-related knowledge made accessible through priming is not sufficient alone to produce aggressive action, but it can produce aggression when judged usability is increased by a person or an object modeling or representing aggression in the immediate environment (see Berkowitz, 1984; Carver et al., 1983). This research suggests that, typically, the use of accessible aggression-related procedures is judged to be inappropriate and is thus inhibited, but aggression-related environmental cues can function as a disinhibitor, which produces aggression.

There is also research suggesting that chronic accessibility can influence behavior. In a study by Lau (1989), for example, individuals' candidate choices were measured after the 1984 presidential election. Measures were also taken of the individuals' political orientations, such as

their party identification and political positions on various issues, and of the chronic accessibility of politically related constructs, such as party constructs and issue constructs. Lau (1989) found that each political orientation was related to candidate choice, but especially for individuals who possessed chronically accessible constructs matching a political orientation. Fazio and Williams (1986) also studied the 1984 presidential election. Months before the election, individuals' attitudes toward the presidential candidates were measured as well as the chronic accessibility of the "candidate-evaluation" linkages. They found that individuals' attitudes predicted their voting behavior, but especially for individuals whose "candidate-evaluation" linkages were highly accessible.

Behavior can also be altered by contextually priming different goals or self-guides. Research on self-discrepancies, for example, has found that individuals' performance on simple tasks can be influenced by priming either ideal or ought self-guides, which produces motor retardation or motor agitation, respectively (e.g., Higgins, Bond, Klein, & Strauman, 1986; Strauman & Higgins, 1987). Research by Bargh and colleagues suggests that goals can be contextually primed and influence subsequent actions and performance (see Bargh & Barn-dollar, 1996; Bargh, Chapter 6, this volume).

In one study, for example, either a competitive goal or an affiliative goal was primed by exposing subjects to goal-related words (see Bargh, 1994). There was either a brief delay or a long delay between this priming and subjects' engagement in a supposedly "unrelated" achievement task. In this task, they could either perform as well as possible to do better than their partner or perform below their ability to protect the feelings of their partner. An impression measure was also obtained. The results of the study suggest that priming-to-stimulus delay as a condition of accessibility effects might function differently when priming increases the accessibility of goals than when it increases the accessibility of categorical knowledge. Consistent with the literature, the study found that the effect of priming on category-related judgments decreased as the priming-to-stimulus delay increased. But the study also found that the effect of priming on goal-related performance *increased* as the priming-to-stimulus delay increased; that is, the performance advantage from priming the competitive goal versus the affiliative goal increased over time. Apparently, when the accessibility of goals is increased, a kind of Zeigarnik effect occurs in which a goal becomes increasingly accessible the longer it remains unattained (see Lewin, 1951; Zeigarnik, 1938). Thus, the effects of increased accessibility on behavior can actually increase over time.

CONCLUSION

Even when behaviorism dominated most other areas of psychology, knowledge activation was a principle of social psychology discussed in various terms, especially salience and accessibility. Beginning in the 1970s, social psychologists paid greater attention not only to the effects of

knowledge activation but to its determinants and conditions of use. This chapter began with the question, "When people identify, interpret or, more generally, respond to a stimulus, what knowledge will be activated and used in their response?" In answering this question, accessibility and applicability have been described as two distinct sources of knowledge activation, with salience contributing to knowledge activation through its influence on applicability. Knowledge activation was distinguished from knowledge use, with judged usability (i.e., judged relevance and judged appropriateness) contributing to knowledge use.

Many variables determine what knowledge is used to respond to a stimulus. To begin with, only knowledge that is available to the perceiver can be used. Among available knowledge units, those knowledge units that are more chronically accessible are more likely to be used. But if a knowledge unit has been recently or frequently primed in the current context, then it might be used instead of a chronically accessible unit. Which knowledge unit is used also depends, however, on the feature overlap between different knowledge units and the stimulus. High applicability to the stimulus could result in a knowledge unit being used that was not as accessible as alternatives. Because applicability depends on the overlap between knowledge features and the *attended* features of the stimulus, knowledge units can vary in their applicability depending on which stimulus features are salient in the momentary situation. And even though a particular knowledge unit is activated first, it might not be used because its judged usability is low. Or, despite being considered inappropriate as the assigned meaning of the stimulus, it might be perceived to be a relevant standard of comparison.

Evidently, there is no simple answer to the question with which the chapter began. Considerable progress has been made, however, in identifying the different variables and conditions that influence knowledge activation and use. As illustrated, often a combination of these variables underlie psychological phenomena. Understanding these variables and their interrelations can provide insight into complex social problems, such as stereotyping, aggression, and emotional health.

Our current understanding of the principles of knowledge activation and use also raises additional questions for study. For example, if perceivers remember that it was the momentary salience of an object that grabbed their attention, will this influence their use of the activated knowledge in a way that is comparable to the effects of perceivers remembering priming events? Can different sources of excitation be experienced as arising from other sources, such as salience being experienced as accessibility and then producing the same effects as accessibility experiences? Could the chronic accessibility of one construct be increased by connecting it to another construct that is higher in chronic accessibility? We have only begun to explore the possible interrelations among the different variables underlying knowledge activation and use. But it is already clear that a fuller understanding of many important social psychological phenomena requires that such exploration continue.

REFERENCES

- Ajzen, I. (1977). Intuitive theories of events and the effects of base rate information on prediction. *Journal of Personality and Social Psychology*, 35, 303-314.
- Alexander, M. J., & Higgins, E. T. (1993). Emotional trade-offs of becoming a parent: How social roles influence self-discrepancy effects. *Journal of Personality and Social Psychology*, 65, 1259-1269.
- Anderson, J. R. (1983). *The architecture of cognition*. Cambridge, MA: Harvard University Press.
- Anderson, J. R. (1987). Skill acquisition: Compilation of weak-method problem solutions. *Psychological Review*, 94, 192-210.
- Arkin, R. M., & Duval, S. (1975). Focus of attention and causal attributions of actors and observers. *Journal of Experimental Social Psychology*, 11, 427-438.
- Baldwin, M. W., Fehr, B., Keedian, E., Seidel, M., & Thomson, D. W. (1993). An exploration of the relational schemata underlying attachment styles: Self-report and lexical decision approaches. *Personality and Social Psychology Bulletin*, 19, 746-754.
- Banaji, M. R., Hardin, C., & Rothman, A. J. (1993). Implicit stereotyping in person judgment. *Journal of Personality and Social Psychology*, 65, 272-281.
- Bargh, J. A. (1982). Attention and automaticity in the processing of self-relevant information. *Journal of Personality and Social Psychology*, 43, 425-436.
- Bargh, J. A., & Barndollar, K. (1996). Automaticity in action: The unconscious as repository of chronic goals and motives. In P. M. Gollwitzer and J. A. Bargh (Eds.), *The psychology of action: Linking cognition and motivation to behavior* (pp. 457-481). New York: Guilford.
- Bargh, J. A., Bond, R. N., Lombardi, W. J., & Tota, M. E. (1986). The additive nature of chronic and temporary sources of construct accessibility. *Journal of Personality and Social Psychology*, 50, 869-878.
- Bargh, J. A., Lombardi, W. J., & Higgins, E. T. (1988). Automaticity of chronically accessible constructs in person x situation effects on person perception: It's just a matter of time. *Journal of Personality and Social Psychology*, 55, 599-605.
- Bargh, J. A. & Pietromonaco, P. (1982). Automatic information processing and social perception: The influence of trait information presented outside of conscious awareness on impression formation. *Journal of Personality and Social Psychology*, 43, 437-449.
- Bargh, J. A., & Pratto, F. (1986). Individual construct accessibility and perceptual selection. *Journal of Experimental Social Psychology*, 22, 293-311.
- Bargh, J. A., & Thein, R. D. (1985). Individual construct accessibility, person memory, and the recall-judgment link: The case of information overload. *Journal of Personality and Social Psychology*, 49, 1129-1146.
- Bar-Hillel, M. (1980). The base-rate fallacy in probability judgments. *Acta Psychologica*, 44, 211-233.
- Berkowitz, L. (1984). Some effects of thoughts on anti- and prosocial influences of media events: A cognitive-neoassociation analysis. *Psychological Bulletin*, 95, 410-427.
- Berlyne, D. E. (1958). The influence of complexity and novelty in visual figures on orienting responses. *Journal of Experimental Psychology*, 55, 289-296.
- Biernat, M., & Vescio, T. K. (1993). Categorization and stereotyping: Effects of group context on memory and social judgment. *Journal of Experimental Social Psychology*, 29, 166-202.
- Bobrow, D. G., & Norman, D. A. (1975). Some principles of memory schemata. In D. G. Bobrow & A. Collins (Eds.), *Representation and understanding: Studies in cognitive science* (pp. 131-149). New York: Academic Press.
- Bornstein, M. H. (1979). Perceptual development: Stability and change in feature perception. In M. H. Bornstein & W. Kessen (Eds.), *Psychological development from infancy: Image to intention*. Hillsdale, NJ: Erlbaum.
- Brendl, C. M., Higgins, E. T., & Lemm, K. M. (1995). Sensitivity to varying gains and losses: The role of self-discrepancies and event

- framing. *Journal of Personality and Social Psychology*, 69, 1028-1051.
- Bruner, J. S. (1951). Personality dynamics and the process of perceiving. In R. R. Blake & G. V. Ransey (Eds.), *Perception: An approach to personality*. New York: Ronald Press.
- Bruner, J. S. (1957a). Going beyond the information given. In H. Gruber et al. (Eds.), *Contemporary approaches to cognition*. Cambridge, MA: Harvard University Press.
- Bruner, J. S. (1957h). On perceptual readiness. *Psychological Review*, 64, 123-152.
- Cacioppo, J. T., Petty, R. E., & Sidera, J. A. (1982). The effects of a salient self-schema on the evaluation of pro-attitudinal editorials: Top-down versus bottom-up message processing. *Journal of Experimental Social Psychology*, 18, 324-338.
- Cameron, N. (1963). *Personality development and psychopathology*. Boston: Houghton Mifflin Company.
- Carver, C. S., Cannelen, R. J., Froming, W. J., & Chambers, W. (1983). Modeling: An analysis in terms of category accessibility. *Journal of Experimental Social Psychology*, 19, 403-421.
- Collins, A. M., & Loftus, E. F. (1975). A spreading-activation theory of semantic processing. *Psychological Review*, 82, 407-428.
- Cramer, P. (1968). *Word association*. New York: Academic Press.
- Devine, P. G. (1989). Stereotypes and prejudice: Their automatic and controlled components. *Journal of Personality and Social Psychology*, 56, 5-18.
- Duncker, K. (1945). On problem solving. *Psychological Monographs*, 58, 5(Whole No. 270).
- Duval, S., & Wicklund, R. A. (1972). *A theory of objective self-awareness*. New York: Academic Press.
- Erdley, C. A., & D'Agostino, P. R. (1988). Cognitive and affective components of automatic priming affects. *Journal of Personality and Social Psychology*, 54, 741-747.
- Fazio, R. H. (1986). How do attitudes guide behavior? In R. M. Sorrentino & E. T. Higgins (Eds.), *Handbook of motivation and cognition: Foundations of social behavior* (pp. 204-243). New York: Guilford.
- Fazio, R. H., Chen, J.-M., McDonel, E. C., & Sherman, S. J. (1982). Attitude accessibility, attitude-behavior consistency, and the strength of the object-evaluation association. *Journal of Experimental Social Psychology*, 18, 339-357.
- Fazio, R. H., Powell, M. C., & Herr, P. M. (1983). Toward a process model of the attitude-behavior relation: Accessing one's attitude upon mere observation of the attitude object. *Journal of Personality and Social Psychology*, 44, 723-735.
- Fazio, R. H., & Williams, C. J. (1986). Attitude accessibility as a moderator of the attitude-perception and attitude-behavior relations: An investigation of the 1984 presidential election. *Journal of Personality and Social Psychology*, 51, 505-514.
- Fazio, R. H., & Zanna, M. P. (1978). Attitudinal qualities relating to the strength of the attitude-behavior relationship. *Journal of Experimental Social Psychology*, 14, 398-408.
- Ferguson, T. J., & Wells, G. L. (1980). Priming of mediators in causal attribution. *Journal of Personality and Social Psychology*, 38, 461-470.
- Fiske, S. T. (1980). Attention and weight in person perception: The impact of negative and extreme behavior. *Journal of Personality and Social Psychology*, 38, 889-906.
- Fiske, S. T., Kenny, D. A., & Taylor, S. E. (1982). Structural models for the mediation of salience effects on attribution. *Journal of Experimental Social Psychology*, 18, 105-127.
- Fiske, S. T., & Taylor, S. E. (1991). *Social cognition* (2nd ed.). New York: McGraw-Hill.
- Forbach, G. B., Stanners, R. F., & Hochhaus, L. (1974). Repetition and practice effects in a lexical decision task. *Memory and Cognition*, 2, 337-339.
- Ford, T. E., & Kruglanski, A. W. (in press). Effects of epistemic motivations on the use of accessible constructs in social judgment. *Personality and Social Psychology Bulletin*.
- Gabriele, A., & Fazio, R. H. (1984). Priming and frequency estimation: A strict test of the availability heuristic. *Personality and Social Psychology Bulletin*, 10, 85-89.
- Gilbert, D. T., & Hixon, J. C. (1991). The trouble of thinking: Activation and application of stereotypic beliefs. *Journal of Personality and Social Psychology*, 60, 509-517.
- Graham, S., & Hudley, C. (1994). Attributions of aggressive and nonaggressive African-American male early adolescents: A study of construct accessibility. *Developmental Psychology*, 30, 365-373.
- Hamilton, D. L., & Gifford, R. K. (1976). Illusory correlation in interpersonal perception: A cognitive basis of stereotypic judgment. *Journal of Experimental Social Psychology*, 13, 392-407.
- Heidbreder, E. (1933). *Seven psychologies*. New York: Appleton-Century-Crofts.
- Herr, P. M. (1986). Consequences of priming: Judgment and behavior. *Journal of Personality and Social Psychology*, 51, 1106-1115.
- Herr, P. M., Sherman, S. J., & Fazio, R. H. (1983). On the consequences of priming: Assimilation and contrast effects. *Journal of Experimental Social Psychology*, 19, 323-340.
- Higgins, E. T. (1981). The "communication game": Implications for social cognition and persuasion. In E. T. Higgins, C. P. Herman, & M. P. Zanna (Eds.), *Social cognition: The Ontario symposium* (pp. 343-392). Hillsdale, NJ: Erlbaum.
- Higgins, E. T. (1989a). Knowledge accessibility and activation: Subjectivity and suffering from unconscious sources. In J. S. Uleman & J. A. Bargh (Eds.), *Unintended thought: The limits of awareness, intention and control* (pp. 75-123). New York: Guilford.
- Higgins, E. T. (1989b). Self-discrepancy theory: What patterns of self-beliefs cause people to suffer? In L. Berkowitz (Ed.), *Advances in experimental social psychology*, Vol. 22 (pp. 93-136). New York: Academic Press.
- Higgins, E. T., & Bargh, J. A. (1987). Social cognition and social perception. *Annual Review of Psychology*, 38, 369-425.
- Higgins, E. T., Bargh, J. A., & Lombardi, W. (1985). The nature of priming effects on categorization. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 11, 59-69.
- Higgins, E. T., Bond, R. N., Klein, R., & Strauman, T. (1986). Self-discrepancies and emotional vulnerability: How magnitude, accessibility, and type of discrepancy influence affect. *Journal of Personality and Social Psychology*, 51, 5-15.
- Higgins, E. T., & Brendl, M. (1995). Accessibility and applicability: Some "activation rules" influencing judgment. *Journal of Experimental Social Psychology*, 31, 218-243.
- Higgins, E. T., & Chaires, W. M. (1980). Accessibility of interrelational constructs: Implications for stimulus encoding and creativity. *Journal of Experimental Social Psychology*, 16, 348-361.
- Higgins, E. T., & King, G. (1981). Accessibility of social constructs: Information processing consequences of individual and contextual variability. In N. Cantor & J. Kihlstrom (Eds.), *Personality, cognition, and social interaction* (pp. 69-121). Hillsdale, NJ: Erlbaum.
- Higgins, E. T., King, G. A., & Mavin, G. H. (1982). Individual construct accessibility and subjective impressions and recall. *Journal of Personality and Social Psychology*, 43, 35-47.
- Higgins, E. T., Rholes, W. S., & Jones, C. R. (1977). Category accessibility and impression formation. *Journal of Experimental Social Psychology*, 13, 141-154.
- Higgins, E. T., Roney, C., Crowe, E., & Hymes, C. (1994). Ideal versus ought predilections for approach and avoidance: Distinct self-regulatory systems. *Journal of Personality and Social Psychology*, 66, 276-286.
- Higgins, E. T., & Stangor, C. (1988). Context-driven social judgment and memory: When "behavior engulfs the field" in reconstructive memory. In D. Bar-Tal & A. W. Kruglanski (Eds.), *The social psychology of knowledge* (pp. 262-298). New York: Cambridge University Press.
- Higgins, E. T., & Tykocinski, O. (1992). Self-discrepancies and biographical memory: Personality and cognition at the level of psychological situation. *Personality and Social Psychology Bulletin*, 18, 527-535.
- Higgins, E. T., Van Hook, E., & Dorfman, D. (1988). Do self attributes form a cognitive structure? *Social Cognition*, 6, 177-207.
- Higgins, E. T., & Wells, R. S. (1986). Social construct availability and accessibility as a function of social life phase: Emphasizing

- the "how" versus the "can" of social cognition. *Social Cognition*, 4, 201-226.
- Houston, D. A. (1990). Empathy and the self: Cognitive and emotional influences on the evaluation of negative affect in others. *Journal of Personality and Social Psychology*, 59, 859-868.
- Houston, D. A., & Fazio, R. H. (1989). Biased processing as a function of attitude accessibility: Making objective judgments subjectively. *Social Cognition*, 7, 51-66.
- Huttenlocher, J., & Higgins, E. T. (1971). Adjectives, comparatives, and syllogisms. *Psychological Review*, 78, 487-504.
- Jacoby, L. L., & Kelley, C. M. (1990). An episodic view of motivation: Unconscious influences of memory. In E. T. Higgins & R. M. Sorrentino (Eds.), *Handbook of motivation and cognition: Foundations of social behavior* (Vol. 2, pp. 451-481). New York: Guilford.
- Jacoby, L. L., Kelley, C., Brown, J., & Jasechko, J. (1989). Becoming famous overnight: Limits on the ability to avoid unconscious influences of the past. *Journal of Personality and Social Psychology*, 56, 326-338.
- James, W. (1948). *Psychology*. New York: The World Publishing Company. (Original publication, 1890)
- James, K. (1986). Priming and social categorizational factors: Impact on awareness of emergency situations. *Personality and Social Psychology Bulletin*, 12, 462-467.
- Kahneman, D., & Tversky, A. (1972). Subjective probability: A judgment of representativeness. *Cognitive Psychology*, 3, 430-454.
- Kahneman, D., & Tversky, A. (1973). On the psychology of prediction. *Psychological Review*, 80, 237-251.
- Kassin, S. M. (1979). Consensus information, prediction, and causal attribution: A review of the literature and issues. *Journal of Personality and Social Psychology*, 37, 1966-1981.
- Kelley, H. H. (1950). The warm-cold variable in first impressions of persons. *Journal of Personality*, 18, 431-439.
- Kelly, G. A. (1955). *The psychology of personal constructs*. New York: Norton.
- Koffka, K. (1935). *Principles of Gestalt psychology*. New York: Harcourt, Brace.
- Krech, D., & Crutchfield, R. S. (1948). *Theory and problems of social psychology*. New York: McGraw-Hill.
- Kruglanski, A. W. (1989). *Lay epistemics and human knowledge: Cognitive and motivational bases*. New York: Plenum.
- Kruglanski, A. W., Friedland, N., & Farkash, E. (1984). Laypersons' sensitivity to statistical information: The case of high perceived applicability. *Journal of Personality and Social Psychology*, 46, 503-518.
- Lau, R. R. (1989). Construct accessibility and electoral choice. *Political Behavior*, 11, 5-32.
- Lewin, K. (1951). *Field theory in social science*. New York: Harper.
- Lichtenstein, S. C., Slovic, P., Fischhoff, B., Layman, M., & Combs, B. (1978). Judged frequency of lethal events. *Journal of Experimental Psychology: Human Learning and Memory*, 4, 551-578.
- Logan, G. D. (1980). Attention and automaticity in Stroop and priming tasks: Theory and data. *Cognitive Psychology*, 12, 523-553.
- Lombardi, W. J., Higgins, E. T., & Bargh, J. A. (1987). The role of consciousness in priming effects on categorization. *Personality and Social Psychology Bulletin*, 13, 411-429.
- Macrae, C. N., Bodenhausen, G. V., Milne, A. B., & Jetten, J. (1994). Out of mind but back in sight: Stereotypes on the rebound. *Journal of Personality and Social Psychology*, 67, 808-817.
- Marcel, T., & Forrin, B. (1974). Naming latency and the repetition of stimulus categories. *Journal of Experimental Psychology*, 103, 450-460.
- Markus, H. (1977). Self-schemata and processing information about the self. *Journal of Personality and Social Psychology*, 35, 63-78.
- Martin, L. L. (1986). Set/reset: Use and disuse of concepts in impression formation. *Journal of Personality and Social Psychology*, 51, 493-504.
- Martin, L. L., & Achee, J. W. (1992). Beyond accessibility: The role of processing objectives in judgment. In L. L. Martin & A. Tesser (Eds.), *The construction of social judgments* (pp. 195-216). Hillsdale, NJ: Erlbaum.
- Martin, L. L., Seta, J. J., & Crelia, R. A. (1990). Assimilation and contrast as a function of people's willingness and ability to expend effort in forming an impression. *Journal of Personality and Social Psychology*, 59, 27-37.
- McArthur, L. Z. (1980). Illusory causation and illusory correlation: Two epistemological accounts. *Personality and Social Psychology Bulletin*, 6, 507-519.
- McArthur, L. Z. (1981). What grabs you? The role of attention in impression formation and causal attribution. In E. T. Higgins, C. P. Herman, & M. P. Zanna (Eds.), *Social cognition: The Ontario symposium, Vol. 1* (pp. 201-246). Hillsdale, NJ: Erlbaum.
- McArthur, L. Z., & Ginsberg, E. (1981). Causal attribution to salient stimuli: An investigation of visual fixation mediators. *Personality and Social Psychology Bulletin*, 7, 547-553.
- McArthur, L. Z., & Post, D. L. (1977). Figural emphasis and person perception. *Journal of Experimental and Social Psychology*, 13, 520-535.
- McCall, G., & Simmons, J. L. (1978). *Identities and interaction* (2nd ed.). New York: Free Press.
- McCann, C. D., Ostrom, T. M., Tyner, L. K., & Mitchell, M. L. (1985). Person perception in heterogeneous groups. *Journal of Personality and Social Psychology*, 49, 1449-1459.
- McClelland, D. C., & Atkinson, J. W. (1948). The projective expression of needs: I. The effect of different intensities of the hunger drive on perception. *Journal of Psychology*, 25, 205-232.
- McGuire, W. J., & Padawer-Singer, A. (1976). Trait salience in the spontaneous self-concept. *Journal of Personality and Social Psychology*, 33, 743-754.
- Medin, D. L. (1989). Concepts and conceptual structure. *American Psychologist*, 44, 1469-1481.
- Merriam-Webster's Ninth New Collegiate Dictionary*. (1989). Springfield, MA: Merriam-Webster.
- Milner, P. M. (1957). The cell assembly: Mark II. *Psychological Review*, 64, 242-252.
- Moskowitz, G. B., & Roman, R. J. (1992). Spontaneous trait inferences as self-generated primes: Implications for conscious social judgment. *Journal of Personality and Social Psychology*, 62, 728-738.
- Murray, H. A. (1938). *Exploration in personality*. New York: Oxford University Press.
- Neely, J. H. (1977). Semantic priming and retrieval from lexical memory: Roles of inhibitionless spreading activation and limited-capacity attention. *Journal of Experimental Social Psychology: General*, 106, 226-254.
- Neisser, U. (1967). *Cognitive psychology*. New York: Appleton-Century-Crofts.
- Newcomb, T. M., Turner, R. H., & Converse, P. E. (1965). *Social psychology: The study of human interaction*. New York: Holt, Rinehart & Winston.
- Newman, L. S., & Uleman, J. S. (1990). Assimilation and contrast effects in spontaneous trait inference. *Personality and Social Psychology Bulletin*, 16, 224-240.
- Oakes, P., & Turner, J. C. (1986). Distinctiveness and the salience of social category memberships: Is there an automatic perceptual bias towards novelty? *European Journal of Social Psychology*, 16, 325-344.
- Philippot, P., Schwarz, N., Carrera, P., De Vries, N., & Van Yperen, N. W. (1991). Differential effects of priming at the encoding and judgment stage. *European Journal of Social Psychology*, 21, 293-302.
- Posner, M. I. (1978). *Chronometric explorations of the mind*. Hillsdale, NJ: Erlbaum.
- Posner, M. I., & Warren, R. E. (1972). Traces, concepts, and conscious constructions. In A. W. Melton & E. Martin (Eds.), *Coding Processes in Human Memory*. Washington, DC: Winston.
- Powell, M. C., & Fazio, R. H. (1984). Attitude accessibility as a function of repeated attitudinal expression. *Personality and Social Psychology Bulletin*, 10, 139-148.

- Pryor, J. B., & Kriss, M. (1977). The cognitive dynamics of salience in the attribution process. *Journal of Personality and Social Psychology*, 35, 49-55.
- Reder, L. M. (1983). What kind of pitcher can a catcher fill? Effects of priming in sentence comprehension. *Journal of Verbal Learning and Verbal Behavior*, 22, 189-202.
- Rholes, W. S., & Pryor, J. B. (1982). Cognitive accessibility and causal attributions. *Personality and Social Psychology Bulletin*, 8, 719-727.
- Rogers, C. R. (1961). *On becoming a person*. Boston: Houghton Mifflin Company.
- Roskos-Ewoldsen, D. R., & Fazio, R. H. (1992). On the orienting value of attitudes: Attitude accessibility as a determinant of an object's attraction of visual attention. *Journal of Personality and Social Psychology*, 63, 198-211.
- Ross, G. (1980). Concept categorization in 1 to 2 year olds. *Developmental Psychology*, 16, 391-396.
- Salancik, G. R., & Conway, M. (1975). Attitude inferences from salient and relevant cognitive content about behavior. *Journal of Personality and Social Psychology*, 32, 829-840.
- Sarbin, T. R., Taft, R., & Bailey, D. E. (1960). *Clinical inference and cognitive theory*. New York: Holt, Rinehart & Winston.
- Schuetter, R. A., & Fazio, R. H. (in press). Attitude accessibility and motivation as determinants of biased processing: A test of the MODE model. *Personality and Social Psychology Bulletin*.
- Schwarz, N., & Bless, H. (1992). Constructing reality and its alternatives: An inclusion/exclusion model of assimilation and contrast effects in social judgment. In L. L. Martin & A. Tesser (Eds.), *The construction of social judgments* (pp. 217-245). Hillsdale, NJ: Erlbaum.
- Schwarz, N., Bless, H., Strack, F., Klumpp, G., Rittenauer-Schatka, H., & Simons, A. (1991). Ease of retrieval as information: Another look at the availability heuristic. *Journal of Personality and Social Psychology*, 61, 195-202.
- Schwarz, N., Strack, F., & Mai, H. P. (1991). Assimilation and contrast effects in part-whole question sequences: A conversational-logic analysis. *Public Opinion Quarterly*, 55, 3-23.
- Sedikides, C. (1990). Effects of fortuitously activated constructs versus activated communication goals on person impressions. *Journal of Personality and Social Psychology*, 58, 397-408.
- Sedikides, C., & Skowronski, J. J. (1991). The law of cognitive structure activation. *Psychological Inquiry*, 2, 169-184.
- Shallice, T. (1972). Dual functions of consciousness. *Psychological Review*, 79, 383-393.
- Sherif, M., & Hovland, C. I. (1961). *Social judgment: Assimilation and contrast effects in communication*. New Haven, CT: Yale University Press.
- Sherman, S. J. (1987). Cognitive processes in the formation, change, and expression of attitudes. In M. P. Zanna, J. M. Olson, & C. P. Herman (Eds.), *Social Influence: The Ontario symposium* (Vol. 5, pp. 75-106). Hillsdale, NJ: Erlbaum.
- Sherman, S. J., & Corty, E. (1984). Cognitive heuristics. In R. S. Wyer, Jr. & T. K. Srull (Eds.), *Handbook of social cognition*, Vol. 1 (pp. 189-286). Hillsdale, NJ: Erlbaum.
- Sherman, S. J., Mackie, D. M., & Driscoll, D. M. (1990). Priming and the differential use of dimensions in evaluation. *Personality and Social Psychology Bulletin*, 16, 405-418.
- Sinclair, R. C., Hoffman, C., Mark, M. M., Martin, L. L., & Pickering, T. L. (in press). Construct accessibility and the misattribution of arousal: Schachter and Singer revisited. *Psychological Science*.
- Sinclair, R. C., Mark, M. M., & Shotland, R. L. (1987). Construct accessibility and generalizability across response categories. *Personality and Social Psychology Bulletin*, 13, 239-252.
- Skowronski, J. J., Carlston, D. E., & Isham, J. T. (1993). Implicit versus explicit impression formation: The differing effects of overt labeling and covert priming on memory and impressions. *Journal of Experimental Social Psychology*, 29, 17-41.
- Smith, E. E., & Medin, D. L. (1981). *Categories and concepts*. Cambridge, MA: Harvard University Press.
- Smith, E. R. (1989). Procedural efficiency: General and specific components and effects on social judgment. *Journal of Experimental Social Psychology*, 25, 500-523.
- Smith, E. R. (1990). Content and process specificity in the effects of prior experiences. Target article in T. K. Srull & R. S. Wyer (Eds.), *Advances in social cognition* (Vol. 3, pp. 1-60). Hillsdale, NJ: Erlbaum.
- Smith, E. R. (1993). Procedural knowledge and processing strategies in social cognition. In R. S. Wyer & T. K. Srull (Eds.), *Handbook of social cognition* (2nd ed.). Hillsdale, NJ: Erlbaum.
- Smith, E. R., & Branscombe, N. R. (1987). Procedurally mediated social inferences: The case of category accessibility effects. *Journal of Experimental Social Psychology*, 23, 361-382.
- Smith, E. R., Branscombe, N. R., & Bormann, C. (1988). Generality of the effects of practice on social judgement tasks. *Journal of Personality and Social Psychology*, 54, 385-395.
- Smith, E. R., Stewart, T. L., & Buttram, R. T. (1992). Inferring a trait from a behavior has long-term, highly specific effects. *Journal of Personality and Social Psychology*, 62, 753-759.
- Smith, E. R., & Zarate, M. A. (1992). Exemplar-based model of social judgment. *Psychological Review*, 99, 3-21.
- Sorrentino, R. M., & Higgins, E. T. (1986). Motivation and cognition: Warming up to synergism. In R. M. Sorrentino & E. T. Higgins (Eds.), *Handbook of motivation and cognition: Foundations of social behavior* (pp. 3-19). New York: Guilford.
- Srull, T. K., & Wyer, R. S., Jr. (1979). The role of category accessibility in the interpretation of information about persons: Some determinants and implications. *Journal of Personality and Social Psychology*, 37, 1660-1672.
- Srull, T. K., & Wyer, R. S., Jr. (1980). Category accessibility and social perception: Some implications for the study of person memory and interpersonal judgments. *Journal of Personality and Social Psychology*, 38, 841-856.
- Stepper, S., & Strack, F. (1993). Proprioceptive determinants of emotional and nonemotional feelings. *Journal of Personality and Social Psychology*, 64, 211-220.
- Strack, F. (1992). The different routes to social judgments: Experiential versus informational strategies. In L. L. Martin & A. Tesser (Eds.), *The construction of social judgments* (pp. 249-275). Hillsdale, NJ: Erlbaum.
- Strack, F., Erber, R., & Wicklund, R. (1982). Effects of salience and time pressure on ratings of social causality. *Journal of Experimental Social Psychology*, 18, 581-594.
- Strack, F., Martin, L. L., & Schwarz, N. (1988). Priming and communication: Social determinants of information use in judgments of life satisfaction. *European Journal of Social Psychology*, 18, 429-442.
- Strack, F., Schwarz, N., Bless, H., Kubler, A., & Wanke, M. (1993). Awareness of the influence as a determinant of assimilation versus contrast. *European Journal of Social Psychology*, 23, 53-62.
- Strauman, T. J., & Higgins, E. T. (1987). Automatic activation of self-discrepancies and emotional syndromes: When cognitive structures influence affect. *Journal of Personality and Social Psychology*, 53, 1004-1014.
- Stryker, S. (1980). *Symbolic interactionism*. Menlo Park, CA: Benjamin/Cummings.
- Stryker, S., & Statham, A. (1985). Symbolic interaction and role theory. In G. Lindzey & E. Aronson (Eds.), *Handbook of social psychology*, Vol. 1 (pp. 311-378). New York: Random House.
- Tagiuri, R. (1969). Person perception. In G. Lindzey & E. Aronson (Eds.), *The handbook of social psychology*, Vol. 3 (2nd ed., pp. 395-449). Reading, MA: Addison-Wesley.
- Tajfel, H. (1969). Social and cultural factors in perception. In G. Lindzey & E. Aronson (Eds.), *The handbook of social psychology* (Vol. 3, pp. 315-394). Reading, MA: Addison-Wesley.
- Taylor, S. E., Crocker, J., Fiske, S. T., Sprinzen, M., & Winkler, J. (1979). The generalizability of salience effects. *Journal of Personality and Social Psychology*, 37, 357-368.
- Taylor, S. E., & Fiske, S. T. (1975). Point of view and perceptions of causality. *Journal of Personality and Social Psychology*, 32, 439-445.

- Taylor, S. E., & Fiske, S. T. (1978). Salience, attention, and attribution: Top of the head phenomena. In L. Berkowitz (Ed.), *Advances in experimental social psychology*, Vol. 11 (pp. 249-288). New York: Academic Press.
- Taylor, S. E., Fiske, S. T., Close, M., Anderson, C., & Ruderman, A. (1977). *Solo status as a psychological variable: The power of being distinctive*. Unpublished manuscript, Harvard University.
- Taylor, S. E., & Thompson, S. C. (1982). Stalking the elusive "vividness" effect. *Psychological Review*, 89, 155-181.
- Thompson, E. P., Roman, R. J., Moskowitz, G. B., Chaiken, S., & Bargh, J. A. (1994). Accuracy motivation attenuates covert priming effects: The systematic reprocessing of social information. *Journal of Personality and Social Psychology*, 66, 474-489.
- Thorndike, E. L. (1898). Animal intelligence: An experimental study of the associative processes in animals. *Psychological Review*, 5(2, No. 8).
- Titchener, E. B. (1966). Attention as sensory clearness. In P. Bakan (Ed.), *Attention: An enduring problem in psychology*. Princeton, NJ: Van Nostrand. (Reprinted from *Lectures on the elementary psychology of feeling and attention*, pp. 171-206, 1908, New York: Macmillan)
- Trope, Y., & Ginosar, Z. (1988). On the use of statistical and nonstatistical knowledge: A problem-solving approach. In D. Bar-Tal & A. W. Kruglanski (Eds.), *The social psychology of knowledge* (pp. 209-230). New York: Cambridge Univ. Press.
- Trzebinski, J., & Richards, K. (1986). The role of goal categories in person impression. *Journal of Experimental Social Psychology*, 22, 216-227.
- Tulving, E. (1983). *Elements of episodic memory*. New York: Oxford University Press.
- Tulving, E., & Pearlstone, Z. (1966). Availability versus accessibility of information in memory for words. *Journal of Verbal Learning and Verbal Behavior*, 5, 381-391.
- Tversky, A., & Gati, I. (1978). Studies of similarity. In E. Rosch & B. B. Lloyd (Eds.), *Cognition and categorization* (pp. 79-98). Hillsdale, NJ: Erlbaum.
- Tversky, A., & Kahneman, D. (1973). Availability: A heuristic for judging frequency and probability. *Cognitive Psychology*, 5, 207-232.
- Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, 85, 1124-1131.
- Warren, R. E. (1972). Stimulus encoding and memory. *Journal of Experimental Psychology*, 94, 90-100.
- Wegner, D. M. (1989). *White bears and other unwanted thoughts*. New York: Viking/Penguin.
- Wegner, D. M., & Erber, R. (1992). The hyperaccessibility of suppressed thoughts. *Journal of Personality and Social Psychology*, 63, 903-912.
- Wegner, D. M., Erber, R., & Zanakos, S. (in press). Ironic processes in the mental control of mood and mood-related thought. *Journal of Personality and Social Psychology*.
- Wegner, D. M., & Schneider, D. J. (1989). Mental control: The war of the ghosts in the machine. In J. S. Uleman & J. A. Bargh (Eds.), *Unintended thought* (pp. 287-305). New York: Guilford.
- Wickens, D. D. (1972). Characteristics of word encoding. In A. W. Melton & E. Martin (Eds.), *Coding processes in human memory* (pp. 191-215). New York: V. H. Winston & Sons.
- Williams, C. W. (1993). The effect of priming causal dimensional categories on social judgments. *Social Cognition*, 11, 223-242.
- Winter, L., & Uleman, J. S. (1984). When are social judgments made?: Evidence for the spontaneousness of trait inferences. *Journal of Personality and Social Psychology*, 47, 237-252.
- Wyer, R. S., Jr., & Carlston, D. E. (1979). *Social cognition, inference, and attribution*. Hillsdale, NJ: Erlbaum.
- Wyer, R. S., Jr., & Hartwick, J. (1980). The role of information retrieval and conditional inference processes in belief formation and change. In L. Berkowitz (Ed.), *Advances in experimental social psychology*, Vol. 13 (pp. 241-284). New York: Academic Press.
- Wyer, R. S., & Srull, T. K. (1980). The processing of social stimulus information: A conceptual integration. In R. Hastie, E. B. Ebbesen, T. M. Ostrom, R. S. Wyer, D. L. Hamilton, & D. E. Carlston (Eds.), *Person memory: The cognitive basis of social perception* (pp. 227-300). Hillsdale, NJ: Erlbaum.
- Wyer, R. S., & Srull, T. K. (1981). Category accessibility: Some theoretical and empirical issues concerning the processing of social stimulus information. In E. T. Higgins, C. P. Herman, & M. P. Zanna (Eds.), *Social cognition: The Ontario symposium*, Vol. 1 (pp. 161-197). Hillsdale, NJ: Erlbaum.
- Wyer, R. S., & Srull, T. K. (1986). Human cognition in its social context. *Psychological Review*, 93, 322-359.
- Wyer, R. S., & Srull, T. K. (1989). *Memory and cognition in its social context*. Hillsdale, NJ: Erlbaum.
- Zajonc, R. B. (1960). The process of cognitive tuning and communication. *Journal of Abnormal and Social Psychology*, 61, 159-167.
- Zanna, M. P., & Fazio, R. H. (1982). The attitude-behavior relation: Moving toward a third generation of research. In M. P. Zanna, E. T. Higgins, & C. P. Herman (Eds.), *Consistency in social behavior: The Ontario symposium*, Vol. 2. Hillsdale, NJ: Erlbaum.
- Zeigarnik, B. (1938). On finished and unfinished tasks. In W. D. Ellis (Ed.), *A source book of gestalt psychology* (pp. 300-314). New York: Harcourt, Brace, & World.
- Zukier, H. (1986). The paradigmatic and narrative modes in goal-guided inference. In R. M. Sorrentino & E. T. Higgins (Eds.), *Handbook of motivation and cognition: Foundations of social behavior* (pp. 465-502). New York: Guilford.
- Zukier, H., & Pepitone, A. (1984). Social roles and strategies in prediction: Some determinants of the use of base rate information. *Journal of Personality and Social Psychology*, 47, 349-360.