

Chapter 2
Background

This research program investigates the reasonable possibility that differential information processing strategies can be manifestations of racist prejudgments. In this context, the word prejudgment is used to suggest expectancies — as in cognitive entities or mechanisms preceding yet supporting observable or self-reported judgments. Ultimately, this research program will attempt to use its accumulated technology to assess one desired effect of Air Force race relations education programs. Specifically, we will assess the ability of such education programs to minimize harmful (i.e., biasing) person-memory effects associated with one facet of racism — prejudicial expectation. The defining pillar of this research program leans on a technique developed and reported in the social cognition literature. This technique captures evidence of a rather habit-driven (though not instantaneous) decision made by subjects rapidly presented with information about (typically fictional) characters. These target characters are associated (in the context of the experiment) with some form of generalized expectancy (i.e., they are typically presented as a likable or unlikable person). This is accomplished either by creating the expectancy artificially, or by using targets that are members of a conspicuously or notably-stereotyped group. The rather non-conscious decision involved is one of *either* bolstering one's pre-conceived notions *or* engaging in inconsistency resolution (e.g., *either* marshaling evidence to bolster your prejudicial expectation *or* pondering more earnestly those pieces of information which are inconsistent with your expectancy or well-known prejudicial stereotypes).

Typically, the likelihood of pursuing one *strategy* or the other is manipulated in experimental settings by first providing an artificial expectation, then altering the structure of the person-memory task or adjusting the rate of information flow to the subjects. Early phases of our research program will hopefully reveal how a non-artificial *pre-existing* race-based prejudicial expectancy (of a largely **non-evaluative** as in **non-likable/dislikable** nature) may effect the pursuit of one strategy or the other. Because of the central importance of the *bolstering versus inconsistency-resolution* decision to this research program, what follows will focus on these two information processing strategies. However, in the course of building a logical bridge between these two strategies, we must also discuss the recall of consistent items in the absence of inconsistencies, the individual attribute *need for cognition*, and bolstering compared to *selective attention* and *selective memory*.

Inconsistency Resolution: This strategy is thought to be behind the oft-repeated finding of enhanced recall for expectancy-inconsistent items in memory — at least within the narrowly-defined settings where such recall patterns emerge. Hastie and Kumar (1979) developed the approach which has repeatedly generated such findings. In their pivotal study, their subjects were tasked to form an impression of a target person. First, their subjects would be presented with a list of essentially synonymous adjectives describing some reasonably enduring characteristic of the target (usually of an *evaluative* nature such that the characteristic makes the target person more or less admirable or likable). This list would provide a strong initial impression or expectancy of the target with regard to this characteristic. Their subjects were then presented with a series of behavior-describing sentences — some of these were consistent (i.e., congruent) with expectancies naturally following from the initial impression. Others were either inconsistent or neutral/irrelevant with respect to such expectancies. Hastie and Kumar found that “the single striking effect in the data was the superior recall of the incongruent behavior descriptions in comparison to recall of congruent and neutral descriptions” (p. 27). Numerous studies have replicated these basic findings (Hastie, 1980, 1984; Srull, 1981; Wyer and Gordon, 1982; Stern, et al, 1984; Bargh and Thein, 1985; Belmore and Hubbard, 1987; Driscoll, Hamilton, and Sorrentino, 1989). Typically, these studies find that neutral or irrelevant behaviors are the least likely to be recalled.

The Associative Network and the Person-concept Node. Theoretically, these findings have been explained via the associative network model briefly described in our introduction (i.e., the H-SAN model in Figure 1). As applied in these studies, the experimental subject's evaluative concept of the target person (including the initial expectancy) is represented by a *person-concept node* (Hamilton, et al, 1989). Items of information about the target person (i.e., behavioral items) become attached to this node as they are encoded. As such, links between these accumulating items in memory and the associated person-concept node are the minimum structural components in such a person-memory schema (assuming subjects have an evaluative concept or impression of the target person — see below).

Direct Inter-item Links By-passing the Person-concept Node. This model further asserts that subsequent comparison or integration of behavioral items (within the thought processes of the experimental subjects) will result in direct inter-item links which do not pass through the person-concept node. Principal investigators in this area (e.g., Hastie, 1980; Srull, 1981) have postulated that the formation of direct inter-item links (resulting from such integration or comparison efforts) is much more likely to occur to or from inconsistent behavioral items (as opposed to forming between two consistent items). This presumption is based on the notion that inconsistent items are more difficult to comprehend and are likely to remain in working memory for a longer period. During this time, other behavioral items are presumably retrieved from long-term memory for comparison or for increasing the understanding of the inconsistent items (Srull, et al, 1985). They have further postulated that irrelevant/neutral items are likely to have connections only with the person-concept node (sometimes referred to as a type of super-ordinate node). Presumably, this occurs because irrelevant/neutral items have no other logical connection to the other behavioral items except by association with the person performing the behaviors.

The Retrieval Process in the Associative Network Model. This model also assumes the existence of an information retrieval process whereby the above described cognitive structure is manifested in differential probabilities of recall for consistent, inconsistent, and neutral/irrelevant behavioral items. Presumably, when such a retrieval process is called upon to summon up these items, the process begins at the person-concept node and proceeds along one of the vertical paths (see Figure 1 amended below) to the first item to be recalled. Then the process presumably moves along one of the horizontal direct inter-item links to the next item to be recalled. If there are no such links (or these links lead only to items already recalled), the retrieval process presumably returns to the person-concept node and travels back down another vertical path, and so on. This process “continues until the subject can no longer recall additional items” (Hamilton, 1989, p. 926). Since there are more pathways leading to inconsistent items, the H-SAN model predicts an increased probability of recall for these items (as compared to consistent and neutral/irrelevant items).

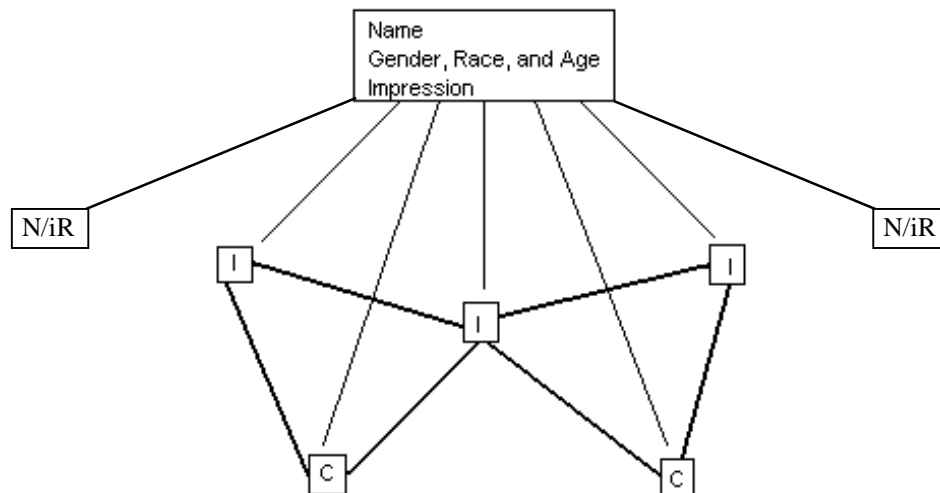


Figure 1 Amended. Hastie-Srull Associative Network (H-SAN) including Irrelevancies

H-SAN and Task (Tasked to Form Impression vs. Remember). Srull, et al (1985) conducted a series of experiments designed to test numerous predictions which follow from such a model. Their first experiment replicated previous findings. Following an initial impression of the target person, they gave their subjects ten seconds to view each behavioral item. Half of these subjects were tested under *memory set* conditions (i.e., “remember the behavioral information as well as possible.....[you will] later be asked to recall its basic content” p. 322). The other half were given *impression formation* instructions where subjects are instructed to understand the target person as well as they can. The predicted recall advantage for inconsistent items was observed — and this advantage was greatest under the impression formation conditions.

The H-SAN Model and Need for Cognition. In their second experiment, Srull, et al, applied Cohen, Stotland and Wolfe’s (1955) individual difference variable *need for cognition* (nCog). This variable was initially described as “a need to structure relevant situations in meaningful, integrated ways — it is a need to understand and make reasonable the experiential world” (p. 291, Cohen, et al, 1955 quoted in Srull, et al, 1985). The slightly expanded conceptual definition applied more recently was relayed by Srull, et al (1985) as “an individual’s tendency to think about things in relationship to one another and attempt to resolve any ambiguity that exists in a situation” (p. 324).

Applying nCog to H-SAN Experimentally. Srull, et al, applied nCog in their second experiment by using Cacioppo and Petty’s (1982) empirically-based assessment device to generate a median split of their subjects (high and low nCog). They predicted that subjects who characteristically tend to think more about items in relation to other items in memory will form more associative inter-item links. Specifically, they predicted that the high and low nCog groups would differ most in the number of inconsistent items recalled (since inter-item links are more likely to occur with such items). They predicted that neutral/irrelevant items would be essentially unaffected by nCog since they are connected

only to the superordinate person-concept node. Recall of consistent items was predicted to be moderately affected since some would be used to compare to inconsistent items during inconsistency resolution. These were exactly the results obtained. High nCog subjects recalled 80 percent more inconsistent items, 40 percent more consistent items, and only six percent more irrelevant items ($N=100$, $F=31.88$, $p<.001$). Looking at individual nCog scores and individual recall patterns, the correlation between the number of inconsistent items recalled and nCog was fairly high ($r=.73$), but moderate for consistent items ($r=.48$) and very low for irrelevant items ($r=.05$).

Attitude and nCog in Person-Memory Research. As a research construct, nCog is often seen in the context of attitude and attitude-change studies. It is not thought to operate as a unchanging individual difference, nor is its effect on information processing entirely predictable. However, as nCog increases, the likelihood of critical, integrated consideration of incoming information increases. Low nCog subjects might only apply such critical consideration to information from dubious sources. Individuals high in nCog may process information systematically (and critically) regardless of whether this information comes from minority or majority (i.e., mainstream) sources (Eagly and Chaiken, 1993). Conversely, subjects low in nCog may only think critically about information from sources they view as non-conformist, suspect, or in the minority. In other words, low nCog subjects are likely to avoid the greater efforts of critical systematic processing — and instead rely on simple heuristics yielding snap judgments (e.g., following rules like “consensus implies correctness” p. 327, Eagly and Chaiken, 1993).

We believe that including nCog in person-memory studies becomes especially important when facets of racism are applied as independent variables. Snap judgments (and assessing behaviors based on superficial evidence) could logically be a contributing factor in developing and maintaining high levels of racism. Therefore, nCog would be predicted to be moderately (and negatively) correlated with racism — especially prejudicial expectation (PE). As shown below, this negative correlation was found in the Phase I results using 397 Caucasians as subjects. High nCog scores are associated with systematic information processors who are viewed as scrutinizing information “by thinking about this information in relation to other information they may possess about the object or issue discussed” (p. 326, Eagly and Chaiken, 1993). This is precisely the sort of consideration thought to yield direct inter-item links in person-memory schemata. As such, when using racism as an independent variable and recall behaviors sensitive to inter-item links as the dependent variable, nCog needs to be accounted for in the model. The means by which we plan to do this will be explained in the methodology section of Phase II. At this juncture, we should note that controlling for nCog levels does not hold constant the reliability low nCog subjects attach to the behavioral items experienced in person-memory research. The degree to which these subjects suspect the truthfulness of an item’s source may still affect the likelihood of systematic/comparative/critical (versus heuristic) processing (as explained in the previous paragraph). In our design, these assessments and their effects are treated as random error (since we have not theorized — nor attempted to control — any systematic determinants for each subject’s assessment of an item source’s reliability).

Processing Time for Inconsistent Items. Two separate studies (Hemsley and Marmurek, 1982; Crocker and Vitkus, 1983 cited in Srull, et al, 1985) supported the H-SAN model’s explanation for more inter-item associative links to and from inconsistent items — namely, that “once an item is identified as incongruent with a prior expectancy, it is held in working memory for a longer period of time” (p. 326, Srull, et al, 1985). These studies used unobtrusive recordings of self-controlled stimulus

(behavioral item) presentation times. These presentation times were much longer for items inconsistent (rather than consistent) with a prior expectancy — as predicted.

Recall of Consistent Items in the Absence of Inconsistencies. Srull, et al, (1985) added a direct test of the H-SAN model's assumption that the inter-item links result from encoding strategies. Srull (1981) had demonstrated that adding consistent behavior items to the stimulus list had no effect on the recall of inconsistent items, but adding inconsistent items improved the recall of consistent items. Srull, et al (1985) attempted to demonstrate that this property of the model is limited to links between inconsistent and consistent items. Their results support their assertion that adding inconsistent items does *not* simply improve recall in general. They found that there is no significant difference between the proportion of consistent and neutral/irrelevant items recalled when there are no inconsistent items in the stimulus list (of course, they were using an artificially-generated expectancy). However, when inconsistent items are added, consistent recalls improve as before, but there is no change in the recall of neutral/irrelevant items. These items apparently remain isolated in the network as shown above in Figure 1 Amended.

Inter-recall Response Intervals. The H-SAN model also predicts a difference in the average time lag between recollections depending on the preceding and following behavioral item type (i.e., the metaphorical or corresponding location in the network — again, see Figure 1 Amended). In other words, since the model predicts more inter-item links involving inconsistent items in the network, it is reasonable to expect that inter-recall response times (i.e., response intervals) will be shorter when the item previously recalled was inconsistent. Since neutral/irrelevant items are allegedly isolated in the network, the response interval following their recollection should be the longest. The H-SAN model would predict the average time interval following consistent items to be somewhere in-between. Such results were obtained by Srull, et al (1985, Experiment 4).

Inconsistency Resolution in the Absence of an Initial Impression. The research from this area having the most direct relevance to our study deletes the initial (i.e., artificial) expectancy generation. By omitting the list of synonymous adjectives from the front end of the experiment, the subjects are left to form an expectancy of the target person based only on behavioral items. Srull, et al (1985) conducted studies “designed to determine whether subjects will spontaneously abstract out generalized expectancies by using completely bottom-up processes” (p. 331).

In their fifth experiment, Srull, et al presented their subjects with fifty behavioral items (actually ten blocks of five sentences each) all pertaining to the same target person. Within each block of five, there were one friendly, two unfriendly, and two irrelevant sentences. Sentence presentation orders were randomized for each subject. Srull, et al, predicted that at the beginning of the sentence presentation, their subjects would be “unaware of which items cluster together and which do not; however, at some point during the acquisition period, subjects should begin to conceptualize the target as an unfriendly person and generally expect unfriendly acts. If one plots the probability of recall as a function of block position...the [inconsistency] effect described earlier should begin to appear at the point the expectancy is formed” (p. 331). The sort of results expected in such a study (under impression formation conditions) are shown in Figure 2. Srull, et al found exactly this sort of recall profile in their fifth experiment.

For their sixth experiment, they reversed the relative proportion of friendly and unfriendly items. Here, they found that “the pattern of results replicated that obtained in Experiment 5 in virtually every detail” (p. 333). In other words, the inconsistency effect survived this reversal.

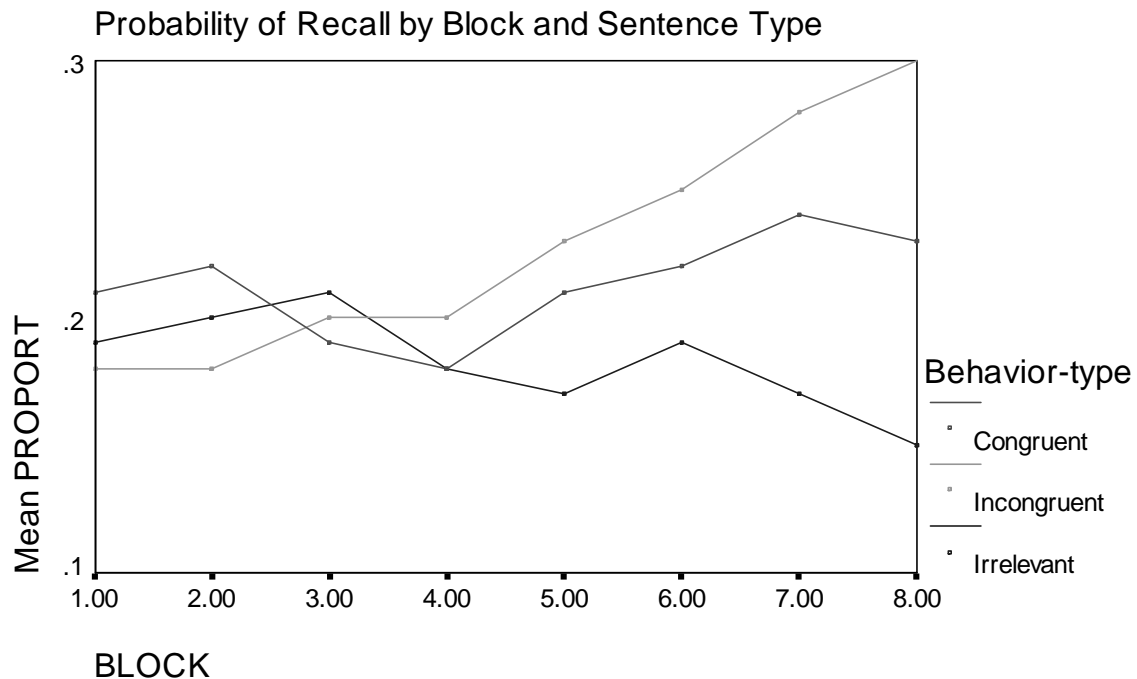


Figure 2. A Srull-type Expectancy Emergence Chart

The results from these last two experiments suggest that the H-SAN model generalizes to circumstances where prior expectancies are not artificially fed to subjects. They also suggest that the model may apply equally well to “both predominately positive and predominately negative behavioral information sets” (p. 334).

Selective Attention and Selective Memory. Obviously, theories which suggest that holders of expectancies are more likely to recall incongruent behaviors become counter-intuitive in the context of racist expectancies exposed to incongruent behaviors (performed by racial minority targets). In other words, one would not generally expect highly racist Caucasoid raters to habitually make a point of remembering the most noble acts of Mongoloid or Negroid rates. A more traditional view would suggest very different recall patterns. William James’ arguments (1890 cited in Eagly and Chaiken, 1993) would suggest that “people tend to approach and attend to information that upholds their attitudes and beliefs but avoid or pay little attention to conflicting information” (p. 591). Put in terms of

Festinger's cognitive dissonance theory (1957, cited in Eagly and Chaiken, 1993) subjects would be expected to seek out information that supports their beliefs, thereby ensuring consonance and avoiding dissonance.

Selectivity Favoring Attitudinally-Congenial Information. Historically, these assertions have been referred to by social psychologists as an *attitude-congenial selectivity* or *attitude-driven selectivity*, and they were widely accepted until Freedman and Sears' (1965) review of relevant research. This review concluded that a selectivity for "congenial" information in natural settings was more likely to be a *de facto selectivity* driven by convenience or accessibility. For instance, conservatives were seen as more likely to live in conservative towns with conservative newspapers in which they would read conservative views. Consistent with this assertion, their review also found that in laboratory settings where non-attitudinal causes of selectivity were controlled, there was no evidence for selective attention.

The Freedman and Sears review diminished interest in attitudinal effects on selective attention, but some later findings called its conclusions into question. Research conducted from the late 1960's to the early 1980's provided evidence for important moderators of attitude-driven selectivity. Frey's (1986) integrative review of this evidence built a case contrary to at least one conclusion of the Freedman and Sears review — specifically, Frey concluded that the view that selectivity is more likely under dissonance-producing circumstances was supported.

Selectivity Moderators. Important moderators of selectivity manifestation include behavioral free-choice and commitment (Cotton and Hieser, 1980, cited in Eagly and Chaiken, 1993). In other words, when subjects have freely chosen to perform an attitude-reflective behavior, afterwards they are more likely to seek out supportive (i.e., attitude-congenial) information. There is some evidence that selective attention is more likely when one's attitude-relevant behavior is irreversible. Conversely, when such behavior is reversible, there may even be a weak tendency to attend to non-supportive information (Frey and Rosch, 1984). However, the "weakness" here may be related to the generally weaker effects associated with non-supportive information. In Eagly and Chaiken's (1993) summary of this area, they found that "reviewers have judged that evidence for selective avoidance of non-supportive information is, in general, weaker than for selective approach of supportive information" (p. 594). They point out that the moderators just described (e.g., free choice/commitment, decision reversibility) show stronger and more consistent effects on information seeking behaviors than on information avoidance behaviors. Again, this is consistent with cognitive dissonance theory, since "only information seeking can actually reduce dissonance — whereas information avoidance can only prevent new dissonance from arising" (p. 594). They also cite evidence for information preferences somewhat independent of a particular attitude. For instance, subjects tend to prefer information with a high utility in relation to future decisions or assessments (e.g., Lowe and Steiner, 1968, and Freedman, 1965a, both cited in Eagly and Chaiken, 1993).

Selectivity Manifestation Given Artificiality or Unimportance. Eagly and Chaiken concluded their summary of this area by pointing out some limitations that may also affect our proposed research. They believe that most of this research has been performed under conditions likely to dampen selectivity effects. Specifically, they argue that the behavioral commitments and attitudes examined in many of these studies are relatively unimportant to the participating subjects. Eagly and Chaiken

express concerns about the artificiality of the supportive and non-supportive information presented in the lab studies which dominate this area, and they believe that the participating subjects will appear less biased in favor of their attitudes (due to impression management) than “they would if their informational preferences were not under scrutiny” (p. 595).

Hypothesized Mechanisms for Memory Selectivity. There are several hypothesized mechanisms by which attitudes might affect memory storage and retrieval (Eagly and Chaiken, 1993). Increased attention is presumed by most researchers to increase retention. Holding attention constant, other mechanisms might make attitudinally-incongruent information more difficult to retrieve. For instance, over time subjects may forget details or subtle/tangential qualifiers of events in memory. If so, they might “cognitively distort attitude-discrepant information over time as a means of reducing the dissonance that exposure to such information presumably arouses” (p. 599). Though unlikely in the context of our proposed research, such forgetfulness can be viewed in Freudian terms — repression causing the information to “be suppressed from consciousness due to its inherent unpleasantness” (p. 599). And, of course, attitudes may bias memory by serving as retrieval cues. In other words, inconsistent information may not fit well with previously-conceived notions of the attitude-object, so the links to the memory of the attitude-object may not be as strong (while within the context of a recall task, a salient attitude might guide the retrieval process to closely associated evidence supporting or consistent with the attitude).

Bolstering in the Context of the H-SAN Model. In the context of the H-SAN related research previously described, increased attention could come into play if behavioral-item/stimulus exposure times tended to be longer than needed for initial comprehension. In other words, subjects could spend these extra moments attending more selectively to those behavioral items congenial to (or consistent with) their expectations or attitudes about the target person (while not giving the others a second glance or second thought). Conversely, they could spend these extra moments focusing on non-supportive/inconsistent information, since their newly forming impression is reversible and they seek knowledge-expanding data as they consider such a reversal. The selective memory mechanisms just mentioned might come into play if an experimental design encourages post-exposure reconsideration of the encoded behavioral items (i.e., subjects could fill in the gaps in their memory with abridgments or amendments consistent with their impression of the target). In either case, their impression could serve as a retrieval cue for consistent information.

Research investigating how subjects use such opportunities (or moments) has been conducted by Wyer and Martin (1986) and by Wyer, Budesheim, Lambert, and Martin (1987). This research has been conducted within the context of *bolstering* — which was defined by Srull and Wyer (1989) as a participating subject’s “review [of] the behaviors that are consistent with their concept of the person — to confirm [this concept’s] validity” (p. 15). Srull and Wyer assert that it is provoked when subjects are confronted with behavioral items (concerning the target person) that are evaluatively inconsistent with their prior impression of the person. The bolstering response to such a provocation also presumes that the subjects are given the time or opportunity to engage in an impression-bolstering review. By design, this may be accomplished either by slowing the rate of behavioral information presentation or setting aside a post-presentation period for review and consideration of the target.

Theoretically, an impression-bolstering review should strengthen the links connecting consistent or impression-congenial items to the super-ordinate node representing the evaluative impression of the target person (Srull and Wyer, 1989). This should increase the likelihood of later recall for these items. It seems reasonable to conclude that those experimental designs which produce a recall advantage for inconsistent items are apparently not allowing sufficient opportunity for bolstering — unless one assumes that a potential reversal of the impression is under consideration so that inconsistent information is being selectively attended.

In one experiment intended to compare (i.e., compete) the effects of bolstering versus the effects of inconsistency resolution, Wyer and Martin (1986, Experiment 1) provided subjects with a trait description of a target person, then a list of behaviors — some consistent with the trait — others contrary. The opportunity to bolster was manipulated by tasking some subjects to perform a distractor task (i.e., drawing a map of the United States and labeling the states in five minutes) while the other subjects were asked to spend the same five minutes thinking about the target person in an overall fashion. Specifically, under this latter condition, the subjects were asked to consider what the target person would really be like if they were to meet him — and then to write down a general description of the person. In the condition applying the distractor task (with presumably very little opportunity for post-information processing), subjects remembered more inconsistent items in the subsequent recall test. The other condition — referred to as the *thought* condition — encouraged substantial post-information processing. As expected, these subjects recalled more consistent items — apparently the result of bolstering.

On-line Bolstering Opportunities. Wyer, Budesheim, Lambert and Martin (1987) took this a step further by comparing the effects of post-information bolstering opportunity to the effects of *on-line* bolstering opportunity (where *on-line* refers to delays in the information presentation rate — presumably allowing impression-bolstering reviews to occur simultaneously with item presentation). Their principal premise was that both bolstering and inconsistency resolution are important components of person impression formation, but that due to limited information processing resources, subjects cannot always do both. In other words, even when subjects want to do both, task demands may often force subjects to “decide upon information-processing strategies that emphasize some cognitive objectives over others” (p.3). The goal of their experiment was to see what sort of task demand characteristics would force a change in strategy.

Consistent with prior research in this area, Wyer, et al assumed that recall profiles would suggest which strategy or strategies were applied by the subjects. This, in turn, was based on the underlying premise that the recall probability of a given behavior is a positive function of both:

- (a) the strength of its links with the central evaluative person-concept, and
- (b) the strength and number of its cross-links with other behaviors in memory.

As mentioned above, the latter of these is assumed to increase recall probability by increasing the number of ways the behavior can be accessed from other components of the target’s representation in memory. If bolstering increases the strength of associations between expectancy-consistent items and the person-concept node, this strategy (in and of itself) should not adversely affect the recall of inconsistent behaviors (improved by cross-links formed during inconsistency resolution) unless, of course, the

work load *makes it seem impossible* to do both. In this case, if bolstering is the chosen strategy, inconsistency resolution occurs less or not at all, and inconsistent item recall falls as recollections of consistent items improve.

Automatic versus Conscious Decisions (among alternative strategies). The phrase *makes it seem impossible* implies that we are studying a fully conscious decision. That is not entirely certain. The “specific processes underlying the attainment of each of these cognitive objectives may be performed automatically” (Wyer, et al, 1987, p. 7). One distinction that must be made is between decisions and processes. All the same, limiting the issue to decisions, the distinction between conscious versus automatic is not always clear. In terms of Wyer and Srull’s (1986) model of human cognition in social contexts, the decision can be thought of as a higher-order encoding objective (residing in a learned goal schema) which is transmitted via the executor to an encoder/organizer. The encoder/organizer is theorized to act as a higher-level interpreter of information transmitted into working memory from a more basic comprehension processor which, in turn, is fed by both long-term memory and the externally-fed sensory stores. This encoder/organizer “has a library of alternative encoding and organization procedures that are performed automatically once they are called” (p. 325). It’s output “depends heavily on the nature of the goal schema....for example, a target’s behavior may be encoded quite differently if the goal is to form a coherent impression [versus] a goal of simply memorizing the [behavioral] information” (p. 325).

The problem with the Wyer and Srull model is the ambiguous use of the phrase *consciously controlled*. On one hand, their model suggests that the sort of higher-order processing central to our study is typically not performed “unless some higher-order processing objective requires it.....[such] encoding would theoretically be performed by the encoder/organizer upon instructions transmitted to it by the executor” (p. 324). On the other hand, they add the following footnote: “the executor is not the intelligence of the system. It has little decision-making capability. Rather, it is primarily an *information-transmission* device that is told what to do by a goal schema” (p. 324). They go on to say that the intelligence of the system resides in the *learned* goal schemata and the libraries of hard-wired routines associated with the processing units. Ultimately, they argue that the procedures in the encoder/organizer’s library are performed without “conscious awareness of the cognitive steps involved.....in contrast, the activities of the executor are consciously directed” (p. 326).

Limited to the context of this study, we suspect a somewhat simpler reality. Our hypotheses treat these information-processing strategy decisions as *rather habitual* or *learned* responses approaching automaticity, but very likely subject to conscious *awareness* and probably memorable as behavioral events in and of themselves after the completed task. We think the notion that these strategy decisions might be reliably manipulated by altering behavioral item presentation rates is in line with this view, and after one more brief comment, we will turn to Wyer, et al’s (1987) use of this manipulation.

The Confound of Habit or Automaticity. Our proposed laboratory exercise may find evidence suggesting that bolstering is by and large consciously controlled while inconsistency resolution is more automatic. For instance, subjects high in socially desirable/impression management (IM) may avoid bolstering stereotype-consistent information. On the other hand, engaging in inconsistency resolution may reflect nCog levels, but may be largely unrelated to impression management. Devine’s conclusions (1989) suggest that the consequences of automatic or non-conscious cognitive processes will indicate a lack of significant difference (between high and low racism scores) in the accessibility or

application of racial stereotypes. If her conclusions generalize to the H-SAN model, then under non-bolstering (i.e., presumably less-controlled) conditions, racism measures should *not* significantly predict recall patterns (once nCog is controlled). In other words, once the influence of the likely (but moderate) negative correlation between nCog and prejudice is minimized, the results of automatic cognitive processes should not differentiate high racism scores from low ones. Conversely, racism scores *should* predict recall patterns when time for consciously-controlled bolstering is allowed. Some of this will depend on the degree to which the early-on developed *need* for cognition has stabilized into cognitive processing *habits*. It is very reasonable to expect that all sorts of on-line processing decisions are more habit-based than decisions made during a five-minute period of reflection. It may be that even when subjects are given much more time than needed to read a behavioral item (i.e., stimulus), their information processing-strategy decisions remain largely habit-driven. More than differing recall probabilities will be required, however, since a recall preference for stereotype-consistent items in highly prejudiced subjects suffers from reasonable alternative explanations (e.g., dissonance-reducing filter action possibly at the incoming end of Wyer and Srull's comprehender/workspace/encoder/organizer connection *or* personal agenda-supporting filter action at the outgoing response end of their executor — presumably applying processing procedures associated with a goal schema transmitted by the executor). A pronounced anti-racism agenda could produce the opposite recall advantage for the same reasons.

A Potential Anti-Racism Effect. The recall pattern differences resulting from an anti-racism trait versus racism could also be couched in terms of bolstering effects — where subjects measured at very low levels of racial prejudice might be motivated to bolster the behavioral items supporting a more laudable image of a Black target. However, if the previously-mentioned notions of Devine (1989) are correct, a very low racism score should not change what the subjects rather automatically view as stereotype-inconsistent. Again, this would suggest that racism scores will have less predictive power for recall after non-bolstering conditions — though they may be powerfully predictive of recall patterns when bolstering is allowed (especially in the almost completely unencumbered provisions of a post-presentation bolstering condition).

Manipulations to Affect the Bolstering/Inconsistency Resolution Decision. The specific manipulations used by Wyer, et al (Experiment 1, 1987) to force information-processing strategy decisions were two-fold. First, they manipulated behavioral item presentation times. Behaviors were presented serially for either 8 seconds (moderate presentation time) or 12 seconds (long presentation time). These researchers consider a short presentation time to be on the order of 4 seconds — a condition applied in their second experiment — which presumably allows little to no on-line processing. Secondly, Wyer, et al allowed half of their subjects five minutes to think about the information they had just seen. The other half were given a five minute distractor task. After this five minute period, all subjects were given a short judgment form on which they were asked to “rate the target person with respect both to *likability* and several specific traits to which the presented behaviors were relevant” (p. 16). Wyer, et al then gave their subjects a surprise recall test. They asked them to write down all of the behaviors they could remember — using their own wording if necessary. They were allowed ten minutes for their recall task. These manipulations are similar to those used by Wyer and Martin (1986) as described earlier. Wyer, et al added the 12-second presentation condition to allow for on-line bolstering.

We discuss this Wyer, et al study in such detail because it is the person-memory design closest to our own. The key differences between their design and ours being our use of the subjects' own

racial prejudice in place of an experimenter-provided expectancy, as well as our use of a less personal and less *social* dimension for evaluation (*academic potential* versus Wyer, et al's *likableness*). Compared to experimenter-provided impressions, expectations based on racial prejudice might act more like what Stangor and Ruble (1989) would call well-entrenched expectancies. Our use of an evaluative dimension that is more task-related may affect the likelihood that a salient person-concept node is even *formed*. We consider the possible impact of all of these issues below. However, even with these differences, our hypotheses have been heavily influenced by the Wyer, et al results — to which we now turn.

Descriptive versus Evaluative Inconsistencies. Wyer, et al's expected results were couched in terms of two forms of inconsistency: descriptive versus the merely evaluative. Descriptive inconsistencies are those behavioral items which are directly related (but counter to) the current impression or concept of the target person. In their words: “although the concept is primarily *evaluative* in nature (i.e., either favorable or unfavorable), it nevertheless has *descriptive features* that reflect the basis for its formation” (p. 5). For instance, if a favorable concept has been formed of the target person based on reports of his conscientious behaviors on the job, then behavioral items detailing sloppy, careless job-related actions would be *descriptive inconsistencies*. Merely *evaluative inconsistencies* might be behavioral items describing sinister actions (on his part) in his private relationships. The original concept was favorable, but reports of sinister actions (even confined to one's personal life) are unfavorable. In this example, the vague middle-ground between descriptive and merely evaluative would be associated with non-*conscientious* behaviors that are *not* job-related. It is this vague middle-ground that may seem to apply in our design.

However, in our sentence set, we have attempted to use only those behavioral items which are either negative, positive, or neutral/irrelevant in relation to our intended appraisal dimension *academic potential*. For those items which are not a direct reflection of apparent academic ability, the intent has been to present items which, at a minimum, reflect the target's apparent dedication to intellectual self-improvement or lawful academic success. We've attempted to keep all items that are not related to academic potential (in *any* reasonable way) evaluatively neutral. Our efforts in this regard (coupled with Wyer, et al's results — see below) hopefully protected us from effects associated with the distinction between descriptive and evaluative consistency.

The Effects of Manipulating the Opportunity to Bolster. Wyer, et al replicated the findings of several previous studies: when sentence presentation time was moderate (as in 8 seconds) *and* there was no opportunity for post-presentation bolstering (due to the five-minute distractor task), sentences describing inconsistent behaviors were more likely to be recalled. Presumably, these subjects are afforded just enough time for inconsistency resolution (yielding a recall advantage for inconsistent behaviors). Of more particular interest, however, is the finding that when subjects are given the opportunity to reflect on all of the information for several minutes after its presentation, the recall advantage swings to consistent behaviors — even though recall of inconsistent items is not “appreciably affected” (p. 18). And then when Wyer, et al moved the opportunity to bolster from the post-presentation period to the on-line processing times (by applying the distractor task and expanding the sentence presentation times to 12 seconds), the recall advantage stays with consistent behaviors — again presumably due to bolstering — but now the bolstering strategy apparently occurred in place of the inconsistency resolution. This was manifested by a drop in the recall probabilities for inconsistent items (though this begs the

question of what the subjects do with 12 seconds of looking at an inconsistent item — presumably they coast or resume bolstering previous consistencies). In both bolstering conditions, the differences in recall probabilities (inconsistent versus consistent) were significantly greater than this same difference in the non-bolstering condition. However, there was no significant effect due to the type of inconsistency — descriptive versus evaluative (cf. preceding paragraph).

Inconsistencies' Selective Effects on Key Consistencies. Srull, et al (1985) had demonstrated that adding inconsistent behavioral items to the list improved recall for the consistent items — probably by way of inter-item links generated during inconsistency resolution efforts. Based on their results, Wyer, et al (1987) would temper that inference somewhat. Inconsistency resolution in their study may have only increased recall for a few consistent items — presumably the ones used to help understand the occurrence of an unexpected behavior. Subjects trying to explain such a behavior “may rely on their general world knowledge about general types of situational factors that might have led to it” (Hastie, 1984, cited in Wyer, et al, 1987, p. 21). Our proposed design may be especially prone to such an effect — inasmuch as our target’s behaviors occur largely in university settings where situational factors are well-known to our subjects. They are likely to use this knowledge (as much as or even more than our provided sentences) in their inconsistency-resolution efforts. Interpretation of our principal dependent variables, however, should not be seriously hampered by this. It may, in fact, increase the differences across conditions — since given adequate opportunity to bolster, our subjects may engage in “a recall and review of all of the consistent behaviors they have previously encountered and a reassessment of the implications of each behavior individually for the central person concept” (Wyer, et al, 1987, p. 21). Compared to this, the few consistent items probably re-considered (during inconsistency resolutions) should not be enough to alter the average recall probability for all consistent items to the degree it will under bolstering conditions.

The Effects of Anticipated Processing Load. In their second experiment, Wyer, et al eliminated the post-presentation bolstering opportunity, and instead manipulated the likelihood of a particular information processing decision by adjusting anticipated processing load. Their findings are of special interest to us, since our proposed design uses 56 sentences — a greater number than usually applied in similar person-memory studies. In reality, they did not adjust the number of sentences — they just adjusted the subjects’ expectations of this number to one of two levels: low = 20 sentences and high = 90 sentences. When subjects were expecting a large number of sentences (i.e., apparently more than they thought they could remember at one time), they increased the amount of both on-line bolstering *and* inconsistency resolution.

To see if subjects might be changing “the relative priority they give to different types of processing over the course of the sequence of behaviors presented” (p. 33), Wyer, et al computed recall proportions separately for each third of the stimulus series (i.e., first third of the sentence list, then second third, etc.). Unlike the primacy and recency effects often seen in list learning tasks, in this study, the middle sentences were recalled better than the other two “thirds” (p. 33). Since our design looks for a strategy-shift in mid-stream (with the location of the shift hopefully influenced by racial prejudice levels), strategy shifts related only to cognitive load perceptions are problematic. To minimize this sort of contamination, we will emphasize the magnitude of the forthcoming task in our instructions (as accurately and saliently as possible). Beyond this, comparisons of changing recall proportions block by block must take this problem into account. Fortunately, Wyer, et al did not find this to be a critical

difficulty, and concluded that “whatever information processing strategies subjects adopted at the outset did not change appreciably as more information was received” (p. 34). As long as these potential sequential effects do not affect performance differentially by sentence type, we should be able to distinguish them from the effects of our theorized processes.

Enhanced Recall for Inconsistencies as a Transient Phenomenon. A different sort of sequential effect must also be considered. Stangor and Ruble (1989) present an alternative way of explaining the somewhat counter-intuitive findings associated with recall of inconsistencies. They contrast studies providing evidence of a recall bias favoring consistencies (e.g., Bodenhausen and Lichtenstein, 1987) with those yielding contrary evidence (e.g., Hastie & Kumar, 1979; Srull, 1981). They believe that “one possible reason for mixed findings in this area is the failure, in general, to recognize that expectations are not static; they may be based on a little or a lot of information and the knowledge base underlying the expectations may change with time” (p. 19). In other words, mature well-entrenched expectancies, by definition, develop over time.

Developmental Nature of Expectancies. Stangor and Ruble (1989) argue that the developmental nature of expectancies has received very little empirical attention. They suggest that traditional studies finding a recall advantage for inconsistencies imply (by design) that an experimenter-provided expectancy may be similar to a stereotype-based prejudice. Stangor and Ruble believe it is more reasonable (based on cognitive-developmental literature — they cite Ruble, 1987; Miller & Turnbull, 1986) to argue that during the process of expectancy formation (when subjects have little knowledge) they try to form an accurate representation via elaborative processing of inconsistencies. They suggest that the recall patterns evident in the work of Hastie, Srull and Wyer are the product of this developmental activity. Conversely, “well-established expectations may lead incongruent information to be discounted and attributed to situational causes.....leading perceivers to filter or ignore inconsistent information, in an attempt to maintain the established expectancy” (p. 20). They add that perceivers may be so sure of their impressions that they do not seriously attend to new information. Finally, they add that well-established expectancies act as effective retrieval cues for recall of consistencies.

Testing for Mature Expectancy Effects. Stangor and Ruble argue that (prior to their 1989 study) there were no direct tests of their notion that well-established expectancy effects overshadow incongruency effects. To look at these issues from a slightly different perspective, they modified the traditional design in three ways. First, their study used an impression formation task where subjects read about members of two fraternities — one extroverted and the other introverted (arguing that people naturally form stereotype-based expectancies about social clubs or guilds). Secondly, they eliminated the traditional experimenter-provided expectancy and added thirty additional behavioral items at the front end — so that subjects could develop their expectancies naturally or “actively” (p. 23). The extremity of this stimulus material was manipulated. The assessment dimension was extroversion/introversion (an at least moderate departure from likable/dislikeable). Finally, they used a separate block of subjects to rate each behavioral item on an eleven-point scale (-5 = very introverted; +5 = very extroverted) so that they could assess recall advantages for consistent or inconsistent items in a more “absolute” (p. 21) sense.

Stangor and Ruble presented half of their subjects (termed the “experienced” condition) with the above-mentioned thirty extra sentences (for developing an expectancy). They then applied a five-minute distractor task to allow the expectancy to mature. This was followed by presentation of sixty more sentences (at the varying degrees of extroversion/introversion just described). The other subjects (in the “inexperienced” condition) only saw this second list of sixty sentences. Another five-minute distractor task was applied to both groups, followed by a ten-minute recall test. Recollections from the initial thirty items (by subjects in the “experienced” group) were ignored. Their overall judgments about the two fraternities were collected *after* the recall test.

Using their eleven-point rating system, Stangor and Ruble found that the experienced group developed far more extreme impressions of both the introverted fraternity and the extroverted fraternity. The extremity of the congruent items recalled was greater than the extremity of the incongruent items recalled — and the experienced group recalled more extreme items than the inexperienced group — with both target sets. This was despite the fact that the incongruent items were written so as to be more extreme than the congruent items. Stangor and Ruble believe this “confirms how strongly recall was in the direction of congruency” (p. 29).

Though Stangor and Ruble did not measure the pre-existing expectations of their subjects — they later asked a new group of subjects to describe typical expectations. This post-hoc exercise confirmed their fears that fraternity members are usually seen as extroverted. As such, they suspect that pre-existing expectations in their original subjects affected their recall patterns — which appeared to reflect some difficulty in their subjects’ ability to form a coherent impression of an introverted fraternity.

They also admit that their findings are not entirely contrary to those favoring incongruencies. Their study tested memory for targeted social groups (as opposed to a target-person) and “consistency effects are more prevalent in memory for groups because perceivers find it less important to.....reconcile incongruent information about groups than about individuals” (p. 32). The simultaneous use of two target groups for all subjects may have encouraged a more streamlined cognitive representation for each (due to perceived task load in a cognitive resource-depletion sense). This might have favored storage and/or retrieval of congruencies.

Stangor and Ruble compare their results with those from the Srull, et al (1985) study that initially inspired our design. Srull, et al reported that the incongruency advantage increased over time (diametrically contrary to Stangor and Ruble’s arguments). They explain this discrepancy in two ways: Srull’s subjects did not have five minutes in which to “consolidate their expectancies” and (at least in their experienced condition) Stangor and Ruble used nearly twice as many sentences as Srull, et al. They suggest that these differences and the associated results point to a possible inverted “U” shaped function for the recall of inconsistencies. In other words, in the beginning when expectancies are very weak, inconsistencies do not really appear incongruent — so richer processing does not yet occur and their recall advantage is nil. As the expectancy begins to form, its reversal is contemplated and elaborate consideration of incongruencies is likely. Once the expectancy is strong, inconsistencies may be ignored, discounted or distorted, so their recall advantage disappears — and, in fact, reverses. Stangor and Ruble believe that racial, ethnic, or gender stereotype-based expectancies will act as well-entrenched expectations.

Formation of the Evaluative Person-Concept. Despite the fact that Stangor and Ruble (1989) studied memory for groups, simultaneously applied two targets, and did not control for very likely pre-existing expectations, their findings are still relevant to our study because they applied an assessment dimension not directly related to the sure-fire evaluation-provoking dimensions often used in these designs (e.g., kind/hostile, honest/dishonest, etc.). Person-memory research typically involves material related to the likability of the target (Srull, Lichtenstein, and Rothbart, 1985; Srull and Wyer, 1989; Wyer and Martin, 1986). Unfortunately, Stangor and Ruble still used an assessment dimension (extroversion) that is not really a performance dimension — at least not in the traditional sense. What still remains to be tested is the way measures of pre-existing expectations about task performance affect the use of impression-bolstering opportunities.

Applying performance-related behavioral descriptions in a person-memory design with a student as the target and other students as the subjects, it is not clear that the bolstering versus inconsistency resolution research will apply in any case. Its applicability will likely depend on the formation of an evaluative person-concept node which is more probable with behaviors that are idiotic, unkind, dishonest, etc. We believe job performance behaviors would also yield this evaluative node within the cognitive representations of supervisors whose livelihood can depend on subordinate performance. Students evaluating another student may be more akin to team appraisal systems than traditional performance appraisal scenarios. These circumstances add an admitted exploratory quality to our study.