Journal of Personality and Social Psychology
2002, Vol. 82, No. 2, 193–207

0022-3514/02/$5.00 DOI: 10.1037//0022-3514.82.2.193

Exhaustive and Heuristic Retrieval Processes in Person Cognition: Further Tests of the TRAP Model

Leonel Garcia-Marques
University of Lisbon

David L. Hamilton
University of California, Santa Barbara

Keith B. Maddox
Tufts University

The twofold retrieval by associative pathways (TRAP) model (L. Garcia-Marques & D. L. Hamilton, 1996) proposes that two distinct modes of retrieval typically underlie recall and frequency estimation. The model accounts for the simultaneous occurrence of greater recall of incongruent information and higher frequency estimation of congruent information. Three experiments provided further tests of the TRAP model. Experiment 1 manipulated cognitive load (at encoding and at retrieval) and the selectivity of the retrieval goal. Under either high load or a selective retrieval goal, incongruent items ceased to be better recalled. Experiment 2 manipulated the accessibility of expectancy-congruent, -incongruent, or -neutral episodes and found corresponding effects in frequency estimates. Finally, Experiment 3 showed that providing part-list retrieval cues inhibits recall but increases frequency estimates. The TRAP model predicted these results.

Social expectancies and memory are two sides of the same cognitive coin. Expectancies capitalize on knowledge accumulated in memory to ease the cognitive load of everyday information processing. Expectancies thus make learning and memory processes more efficient. However, this increased efficiency carries a cost. Many empirical studies have shown that expectancy-congruent information is given disproportional weight relative to incongruent information (for reviews, see Fiske, 1998; Hamilton & Sherman, 1994; Hamilton, Sherman, & Ruvolo, 1990). One example is the expectancy-based illusory correlation effect (Hamilton & Rose, 1980), the overestimation of the frequency of congruent episodes relative to noncongruent episodes. However, a prominent exception to this pattern of findings is the incongruency effect, the advantage in recall of expectancy-incongruent over congruent information in person cognition (Hastie & Kumar, 1979; Srull, 1981; for meta-analytic reviews, see Stangor & McMillan, 1992; Rojahn & Pettigrew, 1992). Both of these effects are well-known and robust phenomena. Yet the idea that expectancy-incongruent items would be better recalled but that expectancy-congruent items would be overestimated appears counterintuitive. Researchers seemingly assumed that both effects could occur, but that the prevalence of expectancy-congruent or -incongruent information would depend on the particular experimental circumstances.

Consistent with this assumption, qualitatively different explanatory accounts have been provided for these two effects. Explanations for the expectancy-based illusory correlation have stressed either the filtering role of expectancies on information processing or the processing advantage of expectancy-congruent information (Hamilton, 1981). In contrast, the advantage for incongruent information in memory is explained as a result of the difficulties of integrating incongruent with congruent information in a unified person impression (Hastie, 1988; Hastie & Kumar, 1979; Srull, 1981; Srull, Lichtenstein, & Rothbart, 1985). These difficulties result in the formation of a large number of associations between incongruent and other types of information. As congruent information does not require such effortful integration, fewer associations are formed as these items are encoded. At free recall, these associations are used as retrieval pathways and, as a consequence, incongruent information benefits from a higher probability of being recalled relative to congruent information.

Recently, however, Garcia-Marques and Hamilton (1996) proposed a different perspective—the twofold retrieval by associative pathways (TRAP) model—in which both effects are incorporated within the same theoretical framework. We adopted the assumptions of the associative network model used to account for the incongruency effect (Hastie & Kumar, 1979; Srull, 1981; Srull et al., 1985) but added new assumptions about retrieval processes (described below). The consequence is that the TRAP model makes predictions not only about processes involved in free recall

---

Leonel Garcia-Marques, Department of Psychology and Educational Sciences, University of Lisbon, Lisbon, Portugal; David L. Hamilton, Department of Psychology, University of California, Santa Barbara; Keith B. Maddox, Department of Psychology, Tufts University.

We thank Jim Sherman for his valuable commentaries on a previous version of this article and gratefully acknowledge the help of Margarida Garrido, Rita Jerónimo, and Sofia Santos in collecting the data. Preparation of this manuscript was supported in part by National Institute of Mental Health Grant 40058 to David L. Hamilton.

Correspondence concerning this article should addressed to Leonel Garcia-Marques, Faculdade de Psicologia e Ciências da Educação, Alameda da Universidade, 1649–013 Lisboa, Portugal. E-mail: gmarques@mail.telepac.pt

but also about how people make frequency estimates on the basis of information stored in memory.

According to the TRAP model, the apparent discrepancy between the two effects occurs because different retrieval procedures are typically involved in free recall and frequency estimation. One retrieval mode was called *exhaustive* and corresponds to standard views of memory search. Exhaustive retrieval is an effortful search process that navigates sequentially through associative memory, using each retrieved item as a cue for the retrieval of other items in the most thorough possible way. Its output consists of specific episodic traces. The other retrieval mode was designated as *heuristic* because it represents a fast and low-resource-demanding indirect way of probing memory. It takes the degree of fit between the available retrieval cues and the stored memory traces as a whole (or the retrieval fluency of this cue) as a clue to some aspect of memory content (e.g., the frequency of a given episodic feature). The retrieval output of these two modes depends on different factors. Whereas the output of exhaustive retrieval will be affected by cognitive load and retrieval strategy, the output of heuristic retrieval will be affected by memory accessibility (i.e., recent and frequent priming) and cue salience (Higgins, 1989, 1996).

The TRAP account of both the incongruency and the illusory correlation effect goes as follows: Free recall typically requires an exhaustive search of memory because its retrieval goal is unselective ("Recall all the items you read.") and because its output consists of specific stored items. To access the greatest number of stored items, previously formed associative links are explored systematically. In the typical person-memory paradigm (Hastie & Kumar, 1979; Srull, 1981), incongruent items have more such associative links with other items, and their probability of being accessed and retrieved is higher relative to congruent items. In contrast, in standard illusory correlation paradigm (Hamilton & Rose, 1980), the frequency estimation task is typically performed in a heuristic manner. It does not entail an exhaustive search of memory because its retrieval goal is selective, focused on accessing certain types of items ("How many X occurred?") and the required output is a generic judgment. Thus retrieval for the purpose of estimation is derived from the ease or the fluency of retrieval of information associated with the target or attribute in question. Because expectancy-congruent items are more strongly associated with the target node than are incongruent items, they are more easily accessed during retrieval, thereby leading to higher frequency estimates. In fact, Garcia-Marques and Hamilton (1996) predicted and found that both effects could occur simultaneously for the same participants.

Although this basic dissociation documented the usefulness of distinguishing between exhaustive and heuristic search, some basic tenets of the TRAP model remained untested. Namely, TRAP assumes that whereas the exhaustive mode is nonselective and low in cognitive efficiency, and its output is elemental, the heuristic mode is accessibility based and efficient, and its output is a composite judgment. This article reports three experiments designed to test these assumptions of the TRAP model.

Experiment 1

When a perceiver tries to form an impression of a target, expectancy-incongruent episodes trigger attempts to integrate the incongruent information with what was previously known about the target (Sherman & Hamilton, 1994). These integrative elaborations affect the resulting cognitive representation such that incongruent items become more densely associated with other items than do congruent items. As free recall is an exhaustive retrieval task, incongruent items are recalled with higher probability than are congruent items (i.e., the incongruency effect).

If, however, other task demands were to interfere with this elaborative processing during encoding, and were therefore to undermine formation of interepisodic associations, the basis for predicting greater recall of incongruent items would no longer be present. Thus, if perceivers operate under a cognitive load during encoding, no difference in recall of congruent and incongruent items occurs (Bargh & Thein, 1985; Macrae, Hewstone, & Griffiths, 1993; Srull, 1981; Srull et al., 1985; Stangor & Duan, 1991).

According to our framework, however, integrative elaboration per se is necessary but not sufficient for incongruency effects to occur. The incongruency effect should occur only when the previously recalled behaviors are used sequentially as cues in the retrieval process—a retrieval strategy that involves navigation through the episodic network. This is a highly demanding retrieval strategy because the travel along interepisodic retrieval pathways implies maintaining at least part of the episodic network in working memory, dealing with continually changing retrieval cues, and keeping track of previously used retrieval cues. Any or all of these aspects of the retrieval process could be easily disrupted by other demands on one's attention and resources. Thus we predicted that the incongruency effect would emerge only when sufficient cognitive resources are available at the time retrieval takes place. Although previous research has suggested that recall is a resource-consuming process (Baddeley, Lewis, Eldridge & Thomson, 1984; Craik, Govoni, Naveh-Benjamim, & Anderson, 1996; Fernandes & Moskovitch, 2000; Hicks & Marsh, 2000; Martin, 1970; Naveh-Benjamim, Craik, Guez, & Dori, 1998; Naveh-Benjamim & Guez, 2000; Trumbo & Milone, 1971), the hypothesis that divided attention at recall would cancel the incongruency effect has not, to the best of our knowledge, been tested. Experiment 1 was designed to provide such a test.

In contrast, we expected no effect of the cognitive-load manipulation on judgment tasks (frequency estimates, impression judgments). In heuristic retrieval, memory is assessed by the ease of retrieval of a small subset of relevant episodes or by a global memory strength response to a given retrieval cue or prompt. It is a less effortful and more efficient retrieval process, in the sense of being less resource demanding and less interference prone. Because of these features, we expected that frequency estimation would not be influenced by the presence of a cognitive-load task either during encoding or during retrieval.

Finally, even if integrative elaboration occurs and sufficient cognitive resources are available, the incongruency effect may not emerge if the recall goal operative at the moment induces a selective or ordered retrieval strategy. Suppose, for example, the recall instructions asked for recall of a certain type of item (e.g., intelligent behaviors), and then later for recall of another type of item (stupid behaviors). We refer to this as an *ordered-recall* strategy (in contrast to the typical *unordered-recall* strategy used in free-recall tasks). In fact, it seems likely that this strategy is rather common in everyday recall (Holtgraves & Srull, 1990).

In our case, resorting to previously recalled behaviors as cues for the retrieval of the remaining items would be a rather cumbersome strategy because of the number of associations (formed at encoding) that link a behavior to opposite types of behaviors. Thus, if one were searching for intelligent behaviors, using specific intelligent behaviors as cues would probably result in retrieving a great number of goal-irrelevant behaviors (i.e., stupid behaviors). To prevent this outcome, it is likely that trait cues are used instead. Appropriate trait cues would help to prevent the retrieval of irrelevant items. As a consequence, interepisodic retrieval pathways would not be explored, and thus no incongruency effect would emerge. In fact, according to the present framework, only under an unordered-recall goal will the incongruency effect emerge in free recall.

In sum, the TRAP model posits that the incongruency effect requires availability of cognitive resources both at encoding and at retrieval and an exhaustive retrieval goal. In Experiment 1 we manipulated these variables and tested their effects on recall, frequency estimation, and impression judgments.

Method

Participants and Design

Participants were 126 students at the University of Lisbon (54 men and 72 women) who volunteered for the study. Participants were randomly assigned to the cells of a 3 (cognitive load: no load, load at encoding, load at retrieval) × 2 (recall goal: ordered, unordered) × 4 (target: mathematician, disco bouncer, childcare professional, traffic policeman) × 2 (trait recall sequence: congruent first, incongruent first) between-subjects factorial design. The last variable was nested within the ordered-recall level of the recall goal variable.

Pretesting of Stimulus Materials

Thirty-eight psychology students not participating in the actual experiment rated several occupational groups on a number of trait scales correlated with either intelligence–stupidity or friendliness–unfriendliness, using 9-point rating scales. From these occupational groups, two pairs were chosen for having generated opposing stereotypic trait expectancies: mathematicians (intelligent), disco bouncers (unintelligent), childcare professionals (friendly), and traffic policemen (unfriendly). In a later session, the same 38 judges rated a large number of behaviors on three trait scales, correlated either with intelligence–unintelligence (i.e., intelligent–unintelligent, ignorant–knowledgeable, and fast thinker–dull) or with friendliness–unfriendliness (i.e., friendly–unfriendly, helpful–unhelpful, and sensitive–insensitive) on similar 9-point rating scales. Twenty-four trait-diagnostic behaviors were selected to represent each dimension (half illustrative of each pole). These behaviors were rated, on average, outside the neutral range of the scale (4.00 to 6.00). In addition, we selected 12 neutral behaviors rated as not being diagnostic of either of those trait dimensions (i.e., their mean rating did not differ from the scale's midpoint by more than 1 point).

Construction of Stimulus Sets

Twenty-four different booklets were constructed, each containing 36 behavior descriptors. The booklets varied in trait-dimension replication (intelligence–unintelligence or friendliness–unfriendliness) and the order of presentation of the behaviors. Each booklet contained either 12 intelligent, 12 unintelligent, and 12 neutral behaviors or 12 friendly, 12 unfriendly, and 12 neutral behaviors.

Procedure

Participants were tested in small groups up to 8 persons at a time. The instructions informed them that the study "concerned with the way in which we form an impression of a person on the basis of his actions." Participants were told that they would be presented with a list of behaviors performed by a given person and were encouraged to try to imagine the type of person he is and to form an overall impression of him. Participants were further informed of the target person's occupation and the kind of impression he has produced in persons who have frequently interacted with him in the past. Thus participants read that João Fonseca was either (a) a mathematician and that people who know him well describe him as very intelligent, a fast thinker, and a knowledgeable person, (b) a disco bouncer, very unintelligent, dull minded, and ignorant, (c) a childcare professional, very friendly, helpful, and sensitive, or (d) a traffic policeman, very unfriendly, unhelpful, and insensitive.

Stimulus booklets were prepared such that the booklets for the mathematician and the disco bouncer target conditions and the booklets for the childcare professional and traffic policeman contained exactly the same behavior sets; only the initial expectancies differed. Consequently, the 12 congruent behaviors for one target condition became incongruent for the matched target condition. Participants read through the booklet, following tape-recorded instructions to turn to the next page every 8 s. After reading the 36 behaviors, participants performed a filler numerical task that took approximately 15 min.

Cognitive load was induced during encoding, during retrieval, or at neither time. In the load-at-encoding conditions, participants read that to simulate "our everyday busy mental life," they would be asked to perform a simultaneous task—the memorization of a nine-digit number. These participants were provided with the nine-digit number and were given 10 s to study it prior to presentation of stimulus behaviors. Participants in the load-at-retrieval conditions were given comparable instructions after having read the stimulus behaviors but immediately prior to assessing the dependent measures. These participants were provided with the nine-digit number, and they were given 10 s to study it prior to completing the dependent measures.

All participants were then asked to either (a) freely recall all the behaviors that had been presented (unordered-recall conditions) or (b) recall one particular type of behavior (e.g., intelligent) and then to recall the other behavior type (e.g., unintelligent; ordered-recall conditions). Within ordered-recall conditions, we manipulated which behavior type was asked for first. After performing the recall task, participants were asked to estimate the frequency with which behaviors illustrative of each relevant trait occurred. Participants had 10 min to perform the free-recall task in unordered-recall conditions and 6 min to perform each of the two ordered-recall tasks in ordered-recall conditions. Participants were given no time limit to perform the frequency estimates.

After completing both tasks, participants were asked to use the three 9-point scales correlated with the relevant trait dimension and a likability scale to convey their impressions of the targets, again with no time limit. Finally, all participants were debriefed and thanked for their participation.

Dependent Measures

Three dependent measures were assessed: (a) free recall of the presented behaviors, (b) estimates of the frequency with which each behavior category (trait) occurred for each target, and (c) trait judgments of the targets on the three relevant 9-point rating scales (intelligent–unintelligent, ignorant–knowledgeable, fast thinker–dull, or friendly–unfriendly, helpful–unhelpful, sensitive–insensitive) and on a fourth, likable–dislikable scale.

Results

To check that participants in the cognitive-load conditions did memorize the nine-digit number, we asked them to reproduce that number. We eliminated from all analyses any participants that committed more than four errors in the reproduction of that number (see Gilbert & Hixon, 1991). This resulted in eliminating 1 participant from the load at encoding–unordered-recall condition, 3 participants from the load at encoding–ordered-recall condition, 2 participants from the load-at-retrieval/unordered-recall condition, and 2 participants from the load-at-retrieval/ordered-recall condition. Also, as preliminary analyses found that the patterns of results generalized across trait–target replications, we collapsed data across those replications.

Behavior Recall

A coder, blind to the experimental conditions and using a lenient gist criterion, categorized the behavioral descriptions recalled by each participant. Reliability of the coding procedure has been documented previously (Garcia-Marques & Hamilton, 1996). Recall intrusions were very infrequent (less than 3%) and were excluded from all analyses. As a manipulation check of the load manipulation, we compared the total number of behaviors recalled in unordered–no-load versus unordered–load-at-encoding conditions. (This is the only contrast that can be meaningfully compared with previous results of the relevant literature, see Bargh & Thein, 1985; Macrae et al., 1993; Srull, 1981; Srull et al., 1985; Stangor & Duan, 1991.) Load participants did recall a lower number of behaviors than no-load participants ($M = 13.46$ vs. $M = 11.82$), $t(44) = 1.74$, $p = .04$, one-tailed, $SD = 3.27$, thus attesting to the effectiveness of the load manipulation.

As neutral items were not contemplated in ordered recall, we performed a one-way analysis of variance (ANOVA) with cognitive load as the independent variable on the number of neutral behaviors remembered under unordered-recall conditions. The effect of load was null, $F(2, 57) = 1.15$, *ns*. We therefore do not consider further the recall of expectancy-neutral behaviors.

Table 1 presents mean recall as a function of the main independent variables of this experiment. A 3 (cognitive load) × 2 (recall goal) × 2 (item type: congruent vs. incongruent) ANOVA, with repeated measures on the last factor, was computed. A significant main effect for item type, $F(1, 112) = 4.60$, $p < .04$, $MSE = 2.15$, showed that congruent behaviors ($M = 4.60$) were better recalled than incongruent behaviors ($M = 4.19$). However, this effect was qualified by 2 two-way interactions.

The first was the predicted Cognitive Load × Item Type interaction, $F(2, 112) = 4.35$, $p < .02$, $MSE = 2.15$, depicted in Figure 1. This interaction is clearly due to the fact that in the no-load conditions, participants recalled fewer congruent ($M = 4.34$) than incongruent ($M = 4.68$) behaviors, whereas the opposite was true in both the load-at-encoding ($M = 4.54$ vs. $M = 3.93$) and the load-at-retrieval ($M = 4.93$ vs. $M = 3.97$) conditions. Another way to describe this interaction is to say that the cognitive-load manipulation produced no effect on the number of congruent behaviors recalled, $t(112) = 1.20$, *ns*, whereas it decreased the number of incongruent behaviors recalled, $t(112) = 2.23$, $p < .03$. This interaction pattern indicates that dividing attention either during encoding or during retrieval makes the incongruency effect disappear. For load-at-encoding conditions, this result replicates previous research (Macrae et al., 1993; Srull, 1981; Srull et al., 1985) and is consistent with the notion that developing interepisodic associations represents a highly demanding encoding process, but one that is critical for the incongruency effect to emerge. For load-at-retrieval conditions, these results represent a novel empirical contribution and support the TRAP-derived predictions we made earlier concerning the low efficiency of the exhaustive recall mode. Thus, even when participants had the necessary cognitive resources during encoding to develop a richly interassociated episodic network, the increased demands on resources at retrieval reversed the incongruency effect.

The second interaction was the Recall Goal × Item Type interaction, $F(1, 112) = 4.58$, $p < .04$, $MSE = 2.15$, shown in Figure 2. This interaction (collapsing across load conditions) reflects the fact that whereas in unordered-recall conditions the number of congruent behaviors recalled ($M = 4.42$) equaled the number of incongruent ($M = 4.42$) behaviors recalled, in ordered-recall conditions congruent behaviors ($M = 4.79$) were much better recalled than incongruent behaviors ($M = 3.97$). These results support our argument that navigating through interepisodic associates requires a nonselective or unordered-recall goal (as is typically the case in person-memory experiments). When a more restricted recall goal is called for, interepisodic associations are no longer exhaustively searched and consequently the incongruency effect disappears.

We also compared recall of congruent and incongruent behaviors for each of the six experimental conditions. It is noteworthy that the incongruency effect appeared only in the no-load/unordered-recall condition, $t(112) = -2.31$, $p < .01$, one-tailed. Moreover, a reversed (i.e., congruency) effect was significant in the load-at-retrieval–ordered-recall condition, $t(112) = 2.29$, $p < .01$, two-tailed. In fact, we contrasted the difference between recall of congruent and incongruent behaviors in the no-load/unordered-recall condition with the same difference in each of the other conditions. In each case, the contrast was significant ($p < .05$ or better). This result supports our argument that both sufficient cognitive resources and a nonselective retrieval strategy are nec-

Table 1
Congruent and Incongruent Behaviors Recalled as a Function of Cognitive Load and Recall Goal

	No load				Load at encoding				Load at retrieval			
	Congruent		Incongruent		Congruent		Incongruent		Congruent		Incongruent	
Recall	*M*	*n*	*M*	*n*	*M*	*n*	*M*	*n*	*M*	*n*	*M*	*n*
Unordered	4.00	21	5.05	21	4.55	20	4.10	20	4.70	20	4.10	20
Ordered	4.68	22	4.32	22	4.53	17	3.76	17	5.17	18	3.83	18

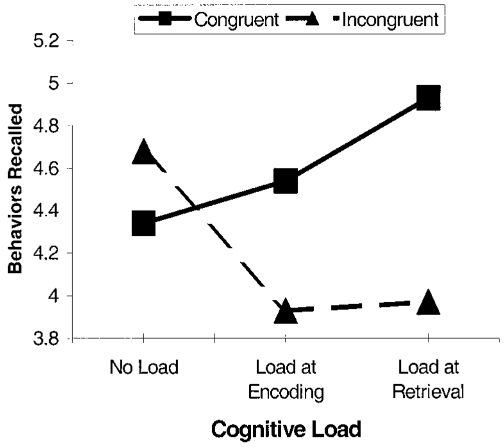

Figure 1. The number of congruent and incongruent behaviors recalled as a function of cognitive load (Experiment 1).

essary conditions for the incongruency effect to emerge; absence of either one is sufficient for its cancellation.

Conditional Probabilities in Unordered Recall

A useful index of the underlying retrieval search process in person memory has been the examination of serial conditional probabilities. In particular, the probability of retrieving a congruent behavior given that another congruent behavior has just been retrieved (a C–C sequence) can be compared with the probability of retrieving an incongruent behavior following recall of a congruent behavior (a C–I sequence). This comparison has been interpreted as an indication of interepisodic retrieval search. Theoretically, the former probability is lower because no C–C episodic associations are formed under standard impression formation conditions (Asuncion & Lam, 1995; Hamilton, Driscoll, & Worth, 1989; Sherman & Hamilton, 1994; Srull, 1981; Srull et al., 1985).

Table 2 presents the relevant conditional probabilities. Several aspects are informative. First, it is clear that the standard pattern of conditional probabilities was obtained in the no-load/unordered-recall condition. Second, and in contrast, a very different pattern was obtained in the load-at-encoding and load-at-retrieval conditions. In particular, the comparison between the no-load and the load-at-retrieval conditions is theoretically important because these conditions shared the same encoding contexts, differing only in retrieval. It is obvious that the differences in C–C and C–I probabilities are dramatically reversed from the no-load to the load-at-retrieval conditions. A planned contrast showed this difference to be highly significant, $t(58) = 3.24$, $p < .001$, one-tailed, $SD = 0.42$.[1] Thus this result supports our argument that depleting resources at retrieval changes the retrieval process in a fundamental way and that sufficient availability of resources at retrieval is a critical condition for the standard pattern of conditional probabilities to be obtained.

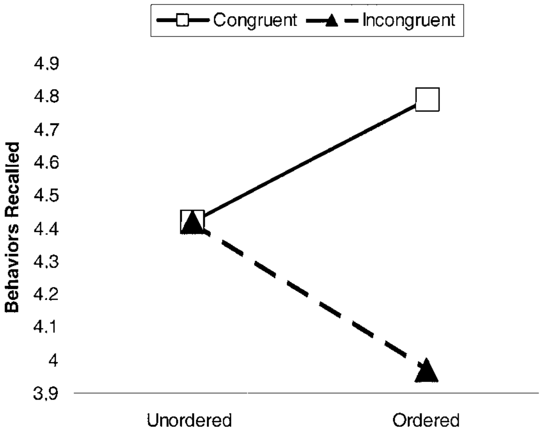

Figure 2. The number of congruent and incongruent behaviors recalled as a function of recall goal.

Table 2
Conditional Recall Probabilities as a Function of Cognitive Load

Recalled behavior sequence	No load	Load at encoding	Load at retrieval
C–C	.25	.37	.42
C–I	.44	.30	.19
I–C	.33	.30	.34
I–I	.35	.21	.36

Note. C = congruent; I = incongruent.

Frequency Estimates

A 3 (cognitive load) × 2 (recall instructions) × 2 (item type) mixed ANOVA, with the last factor being within subjects, was computed on frequency estimates. Only the item-type main effect was significant, $F(1, 112) = 24.16$, $p < .0001$, $MSE = 28.71$, reflecting the strong difference between estimated frequencies for congruent ($M = 14.47$) and incongruent behaviors ($M = 11.09$). No other effects involving the difference between congruent and incongruent behaviors approached significance (all Fs < 1). As predicted, cognitive load did not affect frequency estimates, supporting the idea that the retrieval process underlying frequency estimation (i.e., the heuristic mode) is a highly efficient process (see Bargh, 1996).

---

[1] Recently, Skowronski and Welbourne (1997) voiced the concern that the conditional probability measure may be flawed for its inability to compensate for differences in random expected conditional probabilities. These concerns are especially important when researchers compare conditions that differed in the relative number of congruent and incongruent behaviors recalled (as in our case). In fact, the differences between C–C and C–I found in the no-load conditions disappeared when we used the two indexes developed by Skowronski and Welbourne (conditional probability difference [CPDIFF] and adjusted ratio for individual sequences [ARIS]). The same did not happen in the load-at-retrieval condition. However, the critical contrast under test was significant (or at least marginal in one case) using either CPDIFF and ARIS with standard residuals (i.e., after partialing out recall baseline effects) or even using the relative number of congruent items as a covariate (i.e., congruent behaviors recalled / congruent + incongruent behaviors recalled). (Only with the ARIS index did the critical comparison attained marginal significance [$p < .10$]; all other methods replicated the above presented result [$p <. 05$ or better].)

Also, it is important to note that these results replicate our previous dissociative finding between recall and frequency estimates (Garcia-Marques & Hamilton, 1996) and extend this finding to our present single-target context. Specifically, in the no-load/unordered-recall condition the positive difference between congruent and incongruent frequency estimates was highly significant, $t(112) = 2.68$, $p < .005$, whereas the opposite difference was found in recall for this condition.

Impression Ratings

To simplify reporting the trait-impression data, the scale values for the disco bouncer and traffic policeman were reversed so that all scale values had the same meaning: high scores mean more stereotypic ratings. Scores above 5 are biased toward congruency, and scores below 5 are biased against congruency. A 3 (cognitive load) × 2 (recall instructions) × 3 (trait scales) mixed ANOVA, with the last factor being within subjects, was computed on trait ratings. No significant effects emerged. However, mean ratings on all three scales were above 5 (ranging from 5.60 to 5.90) and differed significantly from the scale's midpoint (i.e., 5; all ps $<$ .001 or better). Thus, as was the case for frequency estimates, impression ratings were impervious to cognitive load, as predicted. Again replicating our previous results, we obtained the predicted dissociation between recall and impression ratings in the no-load/unordered-recall condition, extending this dissociation to the single target case.

Relationships Among Recall, Frequency Estimates, and Impression Ratings

We next examined the relationships among recall, frequency estimates, and impression ratings. For the first two variables we computed difference scores (congruent–incongruent items). We then assessed the reliability of the three rating scales. Because the three scales had an unsatisfactory Cronbach's alpha (i.e., .48), we decided to use only two scales that had the highest correlation ($\alpha > .50$) for each of the four trait replications, and we averaged the two ratings to produce a single trait-scale index. As the pattern of results for the unordered-recall conditions paralleled the results from the corresponding ordered-recall conditions, we collapsed the data across the recall-goal variable. Table 3 presents the obtained correlations among recall, frequency estimates, and impression indexes as a function of cognitive load.

Table 3
Correlations Between Recall, Frequency Estimates, and Impression Trait Judgments as a Function of Cognitive Load

	Recall/ impressions		Recall/ frequency estimates		Frequency estimates/ impressions	
Load	*r*	*n*	*r*	*n*	*r*	*n*
No load	−.12	42	.18	42	.57***	42
Load at encoding	.12	37	.15	37	.28*	37
Load at retrieval	.39**	38	.38**	38	.46***	38

Note. Recall and frequency estimates are as taken as difference scores between congruent and incongruent items. Impressions trait scales are reversed as necessary so that a higher value would correspond to a more stereotyped impression.
* $p < .10$. ** $p < .05$. *** $p < .01$, all two-tailed.

Again, the critical comparison concerns the no-load and the load-at-retrieval conditions (the conditions sharing the same encoding contexts, differing only at retrieval). The results reported in Table 3 show that when resources were depleted at retrieval, recall, frequency estimates, and impression ratings ceased to be dissociated. In our account, this effect was due to the fact that when resources are depleted both retrieval modes explore the same components of the episodic network (namely, the target-node–episode associations). Thus this correlational evidence provides further support for our framework.

One additional result is noteworthy. It is the low correlations between recall and frequency estimates and between recall and trait-impression ratings obtained in the load-at-encoding condition. The load task in this condition divided attention during encoding, and therefore we expected that no interepisodic associations would be formed. If this were the case, then we would expect higher correlations because exhaustive and heuristic modes would have to search through the same components of the episodic network. This clearly was not the case. We return to this point in the *Discussion* section.

Discussion

The main findings of this experiment and their significance can be summarized as follows. First, the basic dissociation between free recall and frequency estimates reported by Garcia-Marques and Hamilton (1996) was replicated. Under standard impression formation conditions (i.e., the no-load/unordered-recall condition), participants recalled more incongruent than congruent behaviors, but estimated a greater frequency of congruent than incongruent behaviors.

Second, cognitive load had a dramatic impact on the recall data. Having to perform a concurrent task either at the time of information encoding or at the time of information retrieval was sufficient to dissolve the incongruency effect. The results for the load-at-retrieval participants were particularly noteworthy because in this condition participants presumably were able to develop interepisodic associations during encoding, yet the incongruency effect disappeared. The TRAP model posits that using each retrieved behavior as a cue for the retrieval of the next one (as typically done in free recall) is a high resource-demanding strategy. However, such exhaustive retrieval is a necessary condition for the incongruency effect to occur. In contrast, cognitive load had no effect on tasks assumed to typically depend on a more efficient retrieval strategy (i.e., the heuristic retrieval mode). Thus both frequency estimates and impression trait ratings were unaffected by the performance of a concurrent task.

Third, the manipulation of recall instructions also had strong effects. Asking participants to use an ordered-recall strategy (first recalling a given type of behavior and then recalling another type of behavior) made the incongruency effect disappear. Again this result fits nicely with the TRAP model's characterization of the exhaustive retrieval mode. As we have argued, the incongruency effect reflects a thorough exploration of behavior-to-behavior associative links (when they exist). By asking participants to retrieve a single type of item (i.e., the ordered-recall-goal conditions), free recall adopts the selective character of the heuristic retrieval process and hence the interitem associations are not pursued during retrieval.

Fourth, the recall results of the no-load/unordered-recall condition differed significantly from those in every other condition of the present experiment. This finding corroborates our notion that resource availability (both at encoding and at retrieval) and an unordered-recall strategy are necessary conditions for the incongruency effect.

Fifth, conditional probabilities for the unordered-recall conditions provided us with an additional empirical basis for theoretical development. In fact, the strong increase of C–C (congruent-to-congruent) transitions observed in the load-at-retrieval condition suggested that, as predicted, the scarcity of resources made participants use a strategy that did not focus on interepisodic associative links. A generic constant cue (e.g., "João Fonseca") with stronger target–item associations may have produced repeated sampling of congruent behaviors. The pattern of conditional probabilities for the load-at-encoding condition was similar, but less clear-cut.

Finally, the correlations among recall, frequency estimates and impression ratings were much stronger in the load-at-retrieval than in the no-load conditions. This pattern of results is also consistent with TRAP model predictions. In the no-load condition, exhaustive and heuristic retrieval modes rely on different features of the network (interitem associative links and target–item links, respectively). In contrast, in the load-at-retrieval condition the two retrieval modes search through the same features (target–item associations). The correlational pattern obtained in the load-at-encoding conditions was, however, much weaker. Theoretically, we expected that depriving participants of cognitive resources during encoding should have prevented the development of interepisodic associations, and consequently both exhaustive and heuristic retrieval modes should have searched through the same retrieval pathways (i.e., target–episode associations). Therefore, as in load-at-retrieval conditions, stronger correlations should have occurred. One possibility is that the cognitive-load manipulation was not sufficiently demanding to completely prevent participants from developing a few interepisodic associations. This would also explain the weaker difference in this condition between the C–C and C–I conditional probabilities. Of course, a more definitive test of this interpretation is needed.

Experiment 2

The research strategy in Experiment 1 concentrated primarily on our claims concerning the variables that affect exhaustive retrieval (cognitive load, recall strategy). In our discussion of Experiment 2 we shift our focus to heuristic retrieval. Experiment 2 was designed to investigate more specifically the TRAP model's contention that the output of heuristic retrieval is highly vulnerable to memory accessibility. Thus, if retrieval fluency is manipulated independently of expectancies, corresponding changes in frequency estimates should follow.

A priming task was used to manipulate the differential accessibility of trait instances (expectancy–congruent, expectancy–incongruent, or neutral). This task was interpolated between the study and the test phase (cf. Gabrielcik & Fazio, 1984). Directly manipulating the accessibility of behavioral trait instances would hopefully provide data tapping the assumed process by which availability produces the expectancy-based illusory correlation effect and other congruency effects. As we argued earlier, the heuristic-search task uses the strength of association between target and item or the retrieval fluency of an item given a referent cue (e.g., a name or general characterization of the target) as a cue to memory content. Thus, according to this line of reasoning, priming neutral instances should not affect the occurrence of the usual illusory correlation effect, whereas priming congruent instances should increase it, and priming incongruent instances should decrease it. In sum, it was predicted that the illusory correlation effect would be greatest after congruent priming, intermediate after neutral priming, and smallest after incongruent priming.

Method

Participants

Participants were 53 University of California, Santa Barbara, undergraduate students (25 men and 28 women), whose participation earned them partial credit for a psychology course.

Pretesting of Stimulus Materials

From previous pretesting (i.e., Garcia-Marques & Hamilton, 1996) two occupational groups were chosen as targets: waitresses and librarians. These groups were chosen because they evoke opposite trait expectancies, namely uncultured versus cultured and fun versus boring. From previous pretesting we also selected 12 trait-diagnostic behaviors for each trait. In addition, we selected 12 neutral behaviors that were not diagnostic of any of those traits.

Construction of Stimulus Sets

All participants were presented with 36 different behavior-descriptive sentences. Each behavior-descriptive sentence referred to the targets, identified by a commonly used first name and occupational group membership. Each behavior description was congruent, incongruent, or neutral with regard to the occupational group stereotype. Each stimulus subset referred to both groups (librarians and waitresses).

There were 16 different stimulus sets. The stimulus sets varied in trait dimension replication (cultured–uncultured or boring–fun), occupational group ascription of the stimulus behaviors (each behavior was ascribed to the two target groups in different stimulus sets), and the order of presentation of the behaviors. As in Experiment 1, occupational group ascription varied orthogonally with each specific behavior, so that for each stimulus set there was a parallel version in which the only change was the group label. This label modification transformed each expectancy-congruent item into an incongruent item and vice versa. Each of these group-label versions presented the behaviors in four different random orders.

Design

This study had a 2 (processing goals: memory versus impression formation) × 3 (nature of prime: congruent, incongruent, or neutral) × 2 (trait replication: boring–fun or cultured–uncultured) × 2 (behavior sets) between-subjects factorial design.

Procedure

Participants were tested individually on an IBM personal computer. Participants read that they were to participate in a study "concerned with the way in which we process and retain verbal descriptions of action" (memory condition) or one dealing "with the way in which we form an impression of a group of persons on the basis of their actions" (impression formation condition). Then they were presented with one of the four stimulus sets of 36 behaviors (2 [trait replications] × 2 [group-label versions]). Each stimulus behavior appeared on the screen for 6 s. After reading all the behaviors, participants had to perform a distracting (lexical-decision) task that took about 10 min.

When the filler task was completed, the priming manipulation was implemented. Participants were informed that their memory for the behaviors they had been presented with earlier in the experiment would be tested. The instructions stated

> You will be asked some questions concerning specific behaviors included in that list. The questions will always have the form: "Who did X?" or "Who said Y?" And we want you to provide the group membership of the person referred to in that particular sentence.

Participants were instructed to press, as quickly as they could, one of the labeled keys (*W* for waitress and *L* for librarian). The need for accuracy was strongly emphasized.

The priming induction consisted of four questions presented to each participant (two pertaining to behaviors ascribed to librarians, two pertaining to waitresses). There were three priming conditions. Four expectancy-relevant and four expectancy-neutral behaviors were randomly chosen as test items. In the congruent-priming condition, participants were presented with four expectancy-relevant behaviors that had originally been ascribed to congruent group labels (two for each group). In the incongruent-priming condition, participants were presented with the same four expectancy-relevant behaviors that had originally been ascribed to incongruent group labels. In the neutral-priming condition, the four expectancy-neutral behaviors were presented as test stimuli.

After completing the priming task, participants were asked to make frequency estimates for each trait–group label match and to evaluate each target group on three 9-point rating scales, two for the relevant trait dimension along with a likability scale. The order of performing these two tasks was counterbalanced, as was the order of performing the frequency judgments (congruent-trait instances first or incongruent-trait instances first). Finally, participants were debriefed, thanked, and dismissed.

Results

Preliminary Analyses

To check the validity of the priming task, it was crucial to determine whether participants based their responses to the four queried behaviors on their memories of the presented information. (Alternatively, participants could have responded to the questions just by guessing.) Thus, the overall error rate was assessed. The overall average number of errors was 1.12, which is significantly lower than the number of errors to be expected had participants responded randomly (i.e., 2), $t(53) = -6.19$, $p < .0001$. In addition, a one-way ANOVA, using priming event as a between-subjects variable, revealed that errors were equally common, regardless of whether the questions included in the priming task were expectancy congruent, incongruent, or neutral, $F(2, 50) = 1.98$, *ns*. In sum, the data showed that participants based their responses to the priming task on their memory for the presented information.

To ensure the effectiveness of the priming task (in terms of design assignment), participants who made more than two misidentifications of the group label of the primed behaviors were excluded from all analyses. Four such participants were therefore deleted from all the analyses.

As in Experiment 1, preliminary analyses of frequency estimates showed that the same pattern of effects was obtained for each target, each trait replication, and each behavior set. Therefore the results were collapsed across these variables and aggregated for item type (expectancy congruent, incongruent, or neutral).

Preliminary analyses of impression ratings also showed identical results across trait replications and behavior sets. Thus all further analyses used data collapsed across these variables. As we used bipolar trait scales and two targets, we created a target variable, a within-subjects variable with two conditions: waitress and librarian targets. The trait dimensions were collapsed in such a way that if the trait judgments followed expectancies, waitresses' ratings would always be higher than librarians' ratings (i.e., boring [1]/fun [9] and cultured [1]/uncultured [9]).

Primary Analyses

The frequency estimates were analyzed by a five-way ANOVA with processing goal (memory vs. impression formation), nature of the priming event (congruent vs. incongruent vs. neutral prime), task order (frequency estimates first vs. impression ratings first), and frequency estimates order (congruent traits first vs. incongruent traits first) as between-subjects factors and item type (congruent vs. incongruent) as a within-subjects factor.[2] An item-type main effect emerged from the analysis, showing that, overall, participants estimated that they had been presented more congruent behaviors ($M = 15.90$) than incongruent behaviors ($M = 13.49$), $F(1, 25) = 6.36$, $p < .02$. Thus the illusory correlation effect was replicated. However, this effect was qualified by the predicted interaction between item type and priming event, $F(2, 25) = 8.35$, $p < .002$, shown in Figure 3.

To further analyze this interaction, we compared the magnitude of the illusory correlations (i.e., the difference between congruent and incongruent frequency estimates) across priming-event conditions. The predictions of the TRAP model were that the illusory correlations would be strongest in congruent-priming conditions, intermediate in neutral-priming conditions, and weakest in incongruent-priming conditions. The results confirmed these predictions. The difference between the congruent and the incongruent frequency estimates was strongest for congruent-priming conditions ($M = 18.62$ vs. $\mathrm{M} = 11.50$, respectively), intermediate in neutral-priming conditions, ($M = 16.56$ vs. $\mathrm{M} = 13.75$, respectively), and lowest in incongruent-priming conditions ($M = 12.71$ vs. $M = 15.12$, respectively). The impact of the priming manipulation is demonstrated by the fact that, in these latter priming conditions, the very robust illusory correlation effect was even reversed: Incongruent items were estimated to have appeared more frequently than congruent items. A priori contrasts showed that these differences between congruent- and incongruent-frequency estimates differed from one another significantly ($p < .05$ or better).

This pattern of results strongly supports the notion that the process of estimating the frequency of a category is carried out by

---

[2] An alternative between-subjects analysis using only the first estimate of each participant produced exactly the same pattern of results.

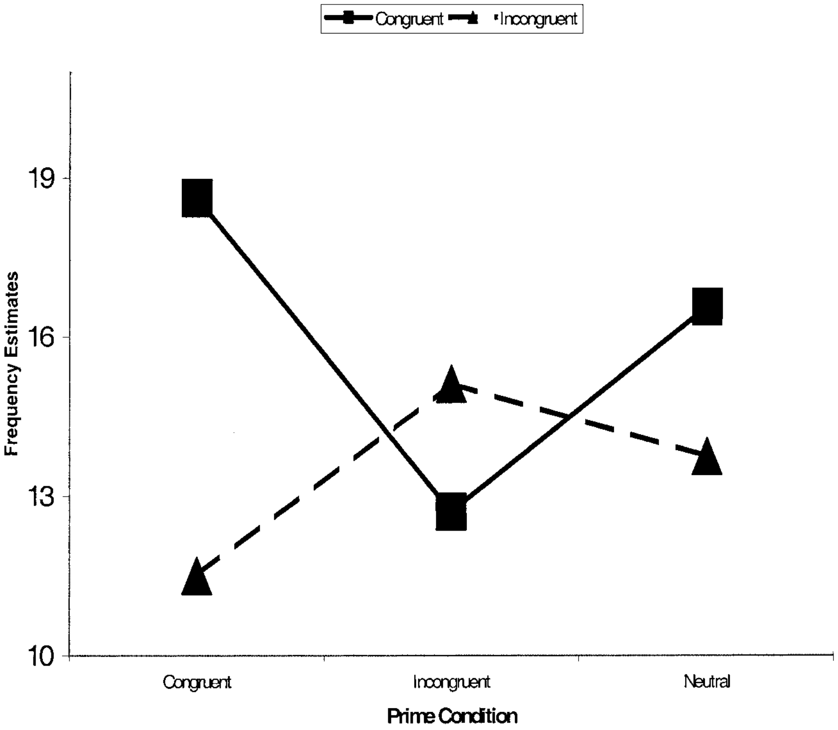

Figure 3. Frequency estimates as a function of prime condition (Experiment 2).

assessing the retrieval fluency with which relevant information can be accessed, given a target referent cue (i.e., the heuristic retrieval). The manipulated accessibility of instances of a given category directly affected estimates of its frequency. This finding generalized across processing goals, as evidenced by the absence of a processing Goal × Priming Event interaction, $F(2, 25) < 1$.

Finally, a nonpredicted three-way interaction of priming event, frequency-judgments order, and item type was also reliable, $F(2, 25) = 4.10$, $p < .03$. The relevant data are presented in Table 4. To simplify the pattern of results, Table 4 depicts the frequency estimation data as difference scores between congruent and incongruent estimates. Inspection of this table reveals that while congruent primes seemed to promote a uniform increase in the magnitude of the illusory correlation effect, this was not the case for the other two priming conditions. In particular, the effect of the incongruent prime seems to be much greater when the incongruent estimate is performed first.

Table 4

Difference Scores Between Congruent and Incongruent Frequency Estimates as a Function of the Priming Event and Frequency Estimation Order

Frequency estimation order	Congruent prime		Neutral prime		Incongruent prime	
	M	*n*	*M*	*n*	*M*	*n*
Congruent first	7.25	8	−.50	8	2.00	8
Incongruent first	7.00	8	5.25	8	−6.33	9

Note. Mean frequency estimates are collapsed across processing goal and task order.

Impression Ratings

The results for impression ratings were quite similar to those for frequency estimates. These data were analyzed by a four-way ANOVA, using processing goal (memory vs. impression formation), priming event (congruent vs. incongruent vs. neutral prime), and task order (frequency estimates first vs. impression ratings first) as between-subjects factors, and target (waitresses vs. librarians) as a within-subjects factor. The predicted two-way interaction between priming event and target was significant, $F(2, 25) = 4.74$, $p < .02$. As Figure 4 shows, the difference between impression ratings of waitresses and librarians was strongest for congruent-priming conditions ($M_{\text{waitress}} = 6.31$ vs. $M_{\text{librarian}} = 4.75$), intermediate in neutral-priming conditions ($M_{\text{waitress}} = 5.62$ vs. $M_{\text{librarian}} = 4.75$) and lowest in incongruent-priming conditions ($M_{\text{waitress}} = 4.71$ vs. $M_{\text{librarian}} = 5.41$). Again the impact of the priming manipulation is apparent in the reversal of the sign of the difference score. A priori contrasts equivalent to those used for frequency estimates were again calculated. However, the differences between priming conditions were not always reliable. In particular, neutral- and congruent-priming conditions did not significantly differ, $t(30) = 0.79$, *ns*. Also the difference between incongruent and neutral-priming conditions only reached marginal significance, $t(31) = 1.60$, $p < .06$. In sum, impression-ratings results were in line with those found for frequency estimates but tended to be weaker. The only other significant result was a nonpredicted and uninterpretable three-way interaction be-

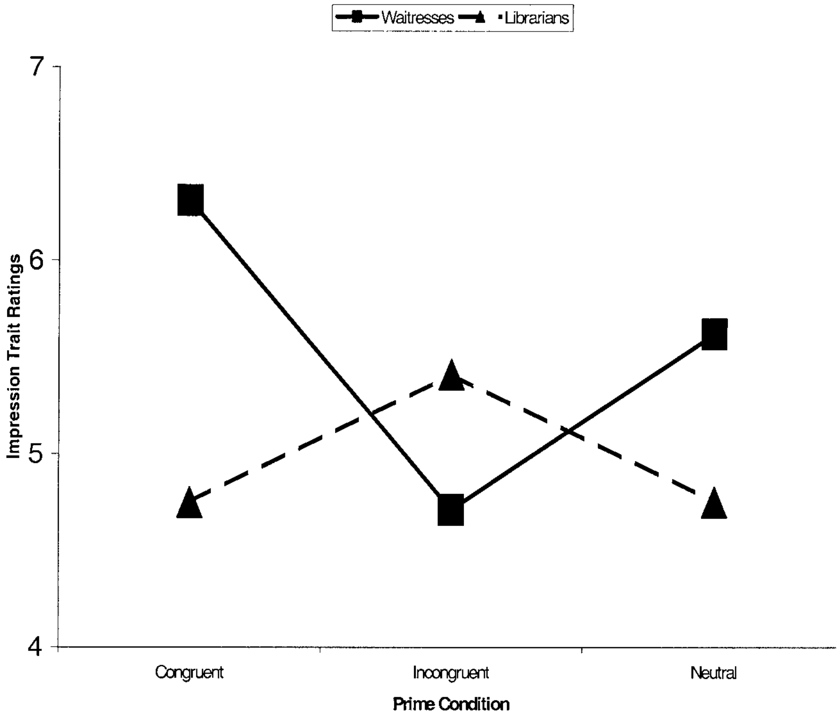

Figure 4. Impression trait ratings of waitresses and librarians as a function of prime condition (Experiment 2).

tween processing goal, task order, and target, $F(1, 25) = 7.85$, $p < .01$.

These results show that manipulating the accessibility of some instances of a given category of behaviors (e.g., congruent behaviors) performed by a given social target (e.g., librarians) has a direct impact on the final impression of that target. These findings are quite consistent with the predictions derived from the TRAP model.

Discussion

The goal of this experiment was to provide a test of the nature of the heuristic retrieval mode hypothesized by the TRAP model. If the expectancy-based illusory correlation effect is mediated by the retrieval fluency of the target information, its magnitude should be positively related to the differential accessibility of expectancy-congruent behaviors (being greatest for congruent-prime, intermediate for neutral-prime, and smallest for incongruent-prime conditions). The pattern of results obtained in the study confirmed this prediction. Specifically, manipulation of the accessibility of small samples of presented behaviors had the predicted impact on both the frequency-estimation judgments and on the impression ratings. These data, therefore, provide strong support for an account of the expectancy-based illusory correlation effect based on the TRAP model.

In this experiment, manipulation of the accessibility of a biased sample of items affected frequency estimates, presumably because heuristic retrieval is highly vulnerable to memory-accessibility effects. But what effects of such priming should be expected in exhaustive retrieval? In fact, the manipulation used in Experiment 2 is very akin to the part-list cued recall paradigm that has been studied in the cognitive literature for some time (Slamecka, 1968; see also S. C. Brown, 1968). In this literature the impact of presenting part-list cues is assessed with a free-recall memory test. The somewhat surprising but robust result is that the probability of recalling noncued list items is decreased.

Taken together, the part-list cue inhibition effect and the results of Experiment 2 suggest that presenting part-list cues at retrieval may inhibit free recall but enhance frequency estimates. Because free recall was not assessed in Experiment 2, Experiment 3 was designed specifically to test the hypothesis that providing part-list cues at retrieval will have dissociative effects in the exhaustive and heuristic retrieval modes.

Experiment 3

The part-list cue effect in recall is commonly explained with reference to associative interference and response competition (e.g., Rundus, 1973). The general rationale is as follows. First, the retrievability of a given stored item is a function of its degree of activation. Second, the total amount of item activation is fixed, and thus relevant stored items must compete for retrieval. Third, items too strongly activated may therefore impede less activated items from being retrieved. These assumptions are, in fact, incorporated in associate network models of person memory (Hastie, 1988; Srull & Wyer, 1989; Wyer & Srull, 1989). More recent approaches have attributed less importance to the associative strength component and have posited selective retrieval processes that inhibit nontarget items to facilitate the retrieval of target items (M. C. Anderson, Bjork, & Bjork, 1994; M. C. Anderson & Neely, 1996; M. C. Anderson & Spellman, 1995). From either perspective, it follows that if a subset of the relevant stored items is made available at retrieval, the probability of noncued items being retrieved is lessened.

Note that, according to these explanations, response competition is a crucial component of this effect. Response competition can only occur, however, when the target for retrieval consists of a set of individual items, as in the case in free recall. According to the TRAP model, however, the output of heuristic retrieval can be a composite (i.e., the overall familiarity with the retrieval cue) because only a generic judgment is called for, not the retrieval of individual items. Empirical support for this assumption comes from findings that participants are able to produce reasonable estimates of the frequency of given categories without being able to retrieve any specific exemplars (Beyth-Marom & Fischhoff, 1977). These results are also akin to research in areas like the tip-of-the-tongue (R. Brown & McNeill, 1966) or feeling-of-knowing effects (Hart, 1965; for a review see Nhouyvanisvong & Reder, 1998), in which people have been shown to be able to make generic judgments about availability in memory of previously learned items without being able to retrieve them.

Thus according to the TRAP model, although the elemental nature of exhaustive retrieval will often result in response competition, the composite nature of heuristic retrieval will often result in response integration. Consequently we expected that part-list cueing would have a dissociative effect in exhaustive and heuristic retrieval tasks. We predicted that presenting part-list cues at retrieval would (a) inhibit retrieval of noncued list items in free recall but (b) increase the corresponding frequency estimates. To the best of our knowledge, this dissociation has never been obtained before and, if found, would provide further support for the need to distinguish between exhaustive and heuristic modes of retrieval.

Method

Participants and Design

Participants were 34 University of Lisbon students (8 men and 26 women) who volunteered for the study. Participants were randomly assigned to the cells of a 2 (part-list cueing conditions: cueing vs. no cueing) $\times$ 4 (target: mathematician, disco bouncer, child care professional, or traffic policeman) $\times$ 2 (cue replications) $\times$ 2 (serial position of cues: list beginning vs. list end) between-subjects factorial, with the last two variables being nested within the cueing levels of the part-list cueing variable.

Construction of Stimulus and Cue Sets

The lists of behaviors used were taken from the pretests reported in Experiment 1. In this experiment, however, we used only stereotype-congruent and neutral behaviors to minimize the formation of interepisodic associations, because these associations might undermine the inhibitory effects of part-list cueing (M. C. Anderson & McCulloch, 1999; Smith & Hunt, 2000). There were eight different booklets, each containing 30 behavior descriptors. The booklets varied in trait dimension replication (intelligence, stupidity, friendliness, or unfriendliness) and the order of presentation of the behaviors. Each booklet contained 30 behaviors: 18 intelligent and 12 neutral behaviors, 18 unintelligent and 12 neutral behaviors, 18 friendly and 12 neutral behaviors, or 18 unfriendly and 12 neutral behaviors. Two different replication sets of behaviors were used as cues for each target group/trait replication.

Procedure

Participants were tested in small groups up to 8 persons at a time. The instructions informed them that the study was "concerned with the way in which we form an impression of a person on the basis of his actions." Participants were told that they would be presented with a list of behaviors performed by a given person and they were to form an overall impression of him. Participants were further informed that to facilitate their task, they would be told about the target occupation and the kind of impression he has produced in persons who have frequently interacted with him in the past. As in Experiment 1, participants read that João Fonseca was either (a) an intelligent mathematician, (b) an unintelligent disco bouncer, (c) a friendly child care professional, or (d) an unfriendly traffic policeman.

All participants then received a booklet. They read through the booklet, following tape-recorded instructions to turn to the next page every 8 s. After reading the 30 behaviors, participants performed a filler numerical task that took about 15 min. All participants were then asked to free recall all the behaviors they had been presented with. In cueing conditions, participants were provided with four of the previously presented behaviors and encouraged to use these behaviors to help them to reproduce the other behaviors included in the booklet they had been presented. Participants had 10 min to perform the free-recall task. After performing the recall task, participants were asked to estimate the frequency with which behaviors illustrative of each relevant trait had occurred. In cueing conditions, participants were again presented with the same four behavior cues and were encouraged to use these cues to help them perform the task. Participants were given no time limit to perform the frequency estimates. After completing both tasks, participants were asked to use 9-point scales to make the three relevant trait-dimension judgments and a fourth likability judgment, again with no time limit. Participants from cueing conditions were again presented with the four behavior cues. Finally, they were debriefed and thanked for their participation.

Results

As preliminary analysis showed no effects for the serial position of the cues and cue replication, we omit further reference to these variables. Also, as the same analyses showed the same pattern of effects across target replications, we collapsed across replications.

Recall

As in Experiment 1, a coder blind to the experimental conditions, using a lenient gist criterion, categorized the behavior descriptions recalled by each participant. Recall intrusions were very infrequent (less than 2%) and were excluded from all analyses.

Neutral items. We first examined recall of neutral behaviors. We performed a one-way ANOVA with part-list cueing (or not) as the independent variable on the number of neutral behaviors recalled. We had no strong expectations regarding the effect of congruent-list cues in neutral-behavior recall. In fact, the effect of cueing on recall of neutral behaviors was null, $F(1, 32) = 0.05$, *ns*.

Congruent items. To enhance comparability, we did not include recalled behaviors that have been used as cues in cueing conditions in the analysis of both cueing and noncueing conditions.

As predicted, the two part-list conditions did differ, $t(32) = 1.76$, $p < .05$ (one-tailed), $SD = 2.27$. Participants in the no-cueing condition remembered a greater number of congruent behaviors ($M = 6.38$) than did participants in the cueing condition ($M = 5.11$). This result shows that the effects of part-list cueing found in the cognitive literature generalize to impression formation contexts.

Frequency Estimates

As predicted, participants in cueing conditions provided significantly, $t(30) = 2.14$, $p = .02$, one-tailed, $SD = 5.48$, higher frequency estimates of number of congruent behaviors presented ($M = 14.18$) than did no-cueing–condition participants ($M = 10.20$). This result replicated the effects found in Experiment 2 and, together with the free-recall results, represents the predicted dissociation.

Impression Ratings

To simplify reporting the trait-impression data, we reversed the scale values for the disco bouncer and traffic policeman so that all scale values had the same meaning—high scores mean more stereotypic ratings. As the three scales exhibited a reasonable coefficient alpha (.70), we collapsed the ratings across the three scales. A 2 (part-list cueing) × 3 (trait scales) repeated-measure ANOVA was performed. No significant effects emerged in the analyses, all *F*s < 1. Thus part-list cueing did not affect trait estimates.

Discussion

The results of Experiment 3 showed that the provision at retrieval of four behaviors included in the originally presented list of behaviors had a dissociative impact on free recall versus frequency estimates. Whereas such cueing inhibited free recall, it increased the corresponding frequency estimates. As discussed earlier, the common explanations of the part-list cue effect stress the role of item competition for retrieval. Like other associative models of person memory (Hastie, 1988; Srull & Wyer, 1989), our depiction of the retrieval process underlying free recall—the exhaustive mode—attributes an important role to retrieval competition. As we noted earlier, response competition can only occur when retrieval is elemental, that is, when the retrieval target consists of specific individual traces. However, the TRAP model also asserts that in heuristic retrieval mode, the output of retrieval can consist of a composite because the goal of retrieval is a generic judgment. In this case, as response competition does not occur, the inhibition effect does not emerge. Thus, part-list cues can only enhance the retrieval accessibility of the relevant category as a whole. Thus the TRAP model predicted this new dissociation of the effects of part-list cues in free recall (the prototypical exhaustive task) and frequency estimation (a task that in most circumstances depends on heuristic retrieval).

Surprisingly, the effects of part-list cues on frequency estimates did not generalize to trait impressions. The effects of priming obtained in Experiment 2 were also more prominent in frequency estimates than in trait ratings. Although we would have predicted otherwise, this null finding could result from ceiling effects, from the use of a proportional integration judgment rule (e.g., an averaging rule, see N. H. Anderson, 1965), and/or from the dissipation of the inhibitory effect of the retrieval cues. Further research that manipulates both set size and task order might be useful for testing these alternatives. Finally, it is important to note that the frequency-estimate results from this experiment replicate the findings of Experiment 2 concerning the role of accessibility in heuristic retrieval. The fact that the two experiments used different paradigms attests to the reliability of this outcome derived from the TRAP model.

General Discussion

Developing the TRAP Model

The TRAP model focuses on the contention that two different search processes underlie different memory measures. Specifically, free recall involves an exhaustive and nonselective search of memory, reporting specific items in memory. In contrast, frequency estimation involves a heuristic and selective search of memory, reported in summary estimation form. The three experiments reported in this article extend, in several theoretically significant ways, the findings reported by Garcia-Marques and Hamilton (1996). In particular, Experiment 1 was focused on properties of the exhaustive retrieval process hypothesized to underlie free recall and included experimental interventions that were predicted to undermine the incongruency effect. Because the TRAP model holds that exhaustive retrieval is a cognitively demanding and resource-consuming process, we predicted that experimental load manipulations would interfere with that process. In contrast, the heuristic retrieval process is not assumed to be as cognitively demanding, and therefore should be less disrupted by cognitive-load manipulations. We manipulated cognitive load during encoding or during retrieval to test parallel, but distinct, aspects of the TRAP model. In both cases cognitive load eliminated the incongruency effect. Participants in the load conditions did not recall more incongruent items than congruent items.

Although we made similar predictions for the load during encoding and load during retrieval conditions, and in fact the results for those two conditions were quite similar, it is important to note that theoretically the reasons for these outcomes are quite different in the two cases. Load during encoding presumably disrupts the formation of interitem associations, meaning that the network of associations consists primarily of the "vertical" links connecting specific items with the target node. In that case the network resembles that assumed to exist in the absence of impression formation goals, such as the memory condition used in some past research (Garcia-Marques & Hamilton, 1996; Srull, 1981; Stern, Marrs, Millar, & Cole, 1984), and indeed the results were similar to those. In contrast, when the load manipulation is not introduced until after the encoding phase, one can assume that the associative network formed in this condition is comparable to that formed under standard impression formation conditions (our control condition), and hence does include the interitem linkages. However, the cognitive-load task in this case interferes with the search through that network at the time of retrieval. Our results show that this search process can be disrupted such that participants did not take advantage of the interitem associations that presumably existed as a result of the encoding process.

Experiment 1 also included a manipulation of retrieval goal in free recall, in which the typical exhaustive and nonselective recall was compared with an ordered-recall strategy involving recall of one type of item at a time. In the TRAP model, the ordered-recall strategy resembles the retrieval process of the heuristic search process, attempting to access only one type of item in performing the required retrieval task. As such, it would not use any interitem associations, which could easily lead to recall of a different type of item. Instead, the ordered-recall task would rely heavily on useful search procedures dependent on, for example, traits as retrieval cues. The consequence was that in all cases, ordered recall resulted in better recall of congruent than incongruent items.

Whereas Experiment 1 included manipulations intended to influence the nature of the resulting network representation (load during encoding) and the exhaustive search process (load during retrieval), Experiment 2 tested theoretical hypotheses about heuristic retrieval. Specifically, our results showed that a priming manipulation included to experimentally influence the accessibility of information in memory affected the magnitude of the illusory correlation. Priming congruent items enhanced the typical illusory correlation effect, whereas priming incongruent items diminished it. The results support the TRAP model's assumption that frequency estimates are affected by the ease of accessibility of information during retrieval.

Experiment 3 extended these findings by using a part-list cued-recall paradigm. Consistent with past research using this paradigm, providing a partial list of previously presented items as retrieval cues actually depressed recall of noncued items. However, and consistent with our predictions, this cueing manipulation actually had an augmenting effect on frequency estimates, compared with no-cue conditions. Thus, Experiment 3 focused on the TRAP-derived difference between exhaustive and heuristic retrieval modes—their response-competition versus response-integration nature—and the dissociation found Experiment 3 in the effects of part-list cueing support this distinction.

Understanding Expectancy-Based Illusory Correlations

The TRAP model originated in an attempt to understand the relationship between two replicable, but seemingly incompatible, findings—the incongruency effect in free recall and the expectancy-based overestimation of congruent items in frequency estimates. In the present work the TRAP model has been extended beyond that initial goal, and our findings have provided empirical support for new theoretical hypotheses. In addition, however, the results of these studies increase our understanding of those two effects in meaningful ways.

It is clear that the expectancy-based illusory correlation is a highly robust effect. In all three studies, and in virtually all conditions of each study, participants made higher frequency estimates for congruent items than for expectancy-incongruent items. The fact that this effect appeared consistently across manipulations of cognitive load (introduced during either encoding or retrieval), priming, and part-list cueing adds further documentation to the generality of this effect. Moreover, the augmentation of this effect following our priming manipulation (Experiment 2) and the part-list cueing (Experiment 3) support the interpretation that this enhanced estimation of expectancy-congruent information is due to the greater accessibility of this information during retrieval.

Understanding the Incongruency Effect

Our results shed new understanding on the incongruency effect in free recall that has been observed in many person-memory experiments. First, the incongruency effect is strongly influenced by load during encoding, with the effect virtually disappearing under this condition. This result presumably is due to the interference of the load with the formation of interitem associations during the encoding process, and according to most theoretical accounts, those associations are crucial for the incongruency effect to occur. Second, the incongruency effect is also strongly influenced by load during retrieval, such that the effect disappeared when a cognitive load was introduced at the time of recall. According to the TRAP model, this outcome reflects the fact that free recall is a demanding process that requires that cognitive resources be available. Third, the incongruency effect is strongly influenced by recall instructions, such that it disappeared when participants were instructed to recall items of one type (intelligent behaviors) and then another type (unintelligent behaviors), rather than the typical recall task in which items can be recalled in any order. The ordered-recall instructions focused retrieval search on certain item content and therefore altered the usual retrieval process of searching through the network representation for any item not previously recalled, regardless of content. Thus, even though the interitem associations existed in this condition, they were not used in the retrieval process, and hence the incongruency effect did not occur.

As shown in Experiment 1, as in many previous studies, the incongruency effect did occur when participants were (a) given a goal of forming an impression, (b) presented information pertaining to one trait dimension, (c) not hindered by cognitive load during either encoding or retrieval, and (d) given recall instructions that encouraged exhaustive, unordered retrieval. These are exactly the conditions under which most person-memory experiments have investigated this effect. However, given the design of Experiment 1, these conditions existed, and hence the incongruency effect occurred, in only one of the six conditions of the experiment. Thus our research documents several important constraints on the conditions under which this outcome will occur.

Implications for Understanding Impression Formation

What are the implications of these findings for our more general thinking about the incongruency effect in the impression formation process? A clear answer to that question will require additional research. However, our results, along with those of others (e.g., Hamilton et al., 1989), may suggest that the enhanced recall of expectancy-incongruent information is not as widespread a phenomenon as our literature might suggest. It may be, of course, that that one condition under which the incongruency effect was observed is one that occurs quite commonly in everyday life, and that the other conditions (under which the effect disappeared) represent unusual circumstances that are useful for theoretical analysis but are not representative of everyday circumstances. In contrast, we propose exactly the opposite. That is, one might wonder how often the perceiver seeks to form an impression of a stranger (a) in which all information acquired pertains to only one trait dimension, (b) with no distractions on attentional and processing resources, and (c) later tries to recall any and all information about the person, in any order, with no concern for content themes. Rather, it seems much more likely that impression formation typically occurs when (a) the perceiver learns an array of information pertaining to several facets of the person's life, (b) the perceiver is simultaneously coping with multiple demands on cognitive resources, and (c) subsequent retrieval of specific information about the person will be for a given purpose that will guide the retrieval process to search for particular content. As shown in the present studies and in previous research (Garcia-Marques & Hamilton, 1996; Hamilton et al., 1989), it is under those conditions that the incongruency effect is much less likely to occur.

Avenues for Future Research

Although the present results provide useful support for the TRAP model, many research avenues remain yet to be explored. As a final note, we briefly offer ideas regarding two possibilities for future development.

We have argued that exhaustive retrieval is a less efficient mode than heuristic retrieval, but we have not yet provided direct evidence of this point. We have presented data indicating that exhaustive retrieval is more interference prone than is heuristic retrieval, but we lack a direct measure of the cognitive resources that are demanded during the operation of each of the two retrieval modes. Research investigating this question would provide a useful extension of our work. Finally, our research and theorizing have basically identified and focused on two memory tasks (i.e., free recall and frequency estimation) that typically depend on the operation of the two retrieval modes. There are, however, a host of memory tasks (e.g., recognition, implicit tests, and memory judgments) by which information is retrieved from memory and used for various purposes. These other retrieval and memory processes provide new opportunities and challenges for evaluating the fruitfulness of TRAP's distinction between exhaustive and heuristic retrieval modes.

In closing, we would like to express our optimism about the usefulness of the TRAP and similar integrative models to social cognition. The TRAP model explores the somewhat neglected topic of retrieval in person cognition and, more importantly, it builds on decades of previous excellent but disparate research (from the illusory correlations and person-memory fields). TRAP is able to offer an integrative and cumulative common framework for these separate lines of inquiry. Social cognition has perhaps overemphasized fascinating new findings while devoting too little time to developing integrative theoretical approaches. We believe that a new balance between these important scientific facets is in order, and hope that integrative models like TRAP will become more common.

References

Anderson, N. H. (1965). Averaging versus adding as a stimulus-combination rule in impression formation. *Journal of Experimental Psychology, 70*, 394–400.

Anderson, M. C., Bjork, R. A., & Bjork, E. L. (1994). Remembering can cause forgetting: Retrieval dynamics in long term memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 20*, 1063–1087.

Anderson, M. C., & McCulloch, K. C. (1999). Integration as a general boundary condition on retrieval-induced forgetting. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 25*, 608–629.

Anderson, M. C., & Neely, J. H. (1996). Interference and inhibition in memory retrieval. In E. L. Bjork & R. A. Bjork (Eds.), *Memory* (pp. 237–313). San Diego, CA: Academic Press.

Anderson, M. C., & Spellman, B. A. (1995). On the status of inhibitory mechanisms in cognition: Memory retrieval as a model case. *Psychological Review, 102*, 68–100.

Asuncion, A. G., & Lam, W. F. (1995). Affect and impression formation: Influence of mood on person memory. *Journal of Experimental Social Psychology, 31*, 437–464.

Baddeley, A. D., Lewis, V., Eldridge, M., & Thomson, N. (1984). Attention and retrieval from long-term memory. *Journal of Experimental Psychology: General, 113*, 518–540.

Bargh, J. A. (1996). Automaticity in social psychology. In E. T. Higgins & A. W. Kruglanski (Eds.), *Social psychology: Handbook of basic principles* (pp. 169–183). New York: Guilford Press.

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.

Beyth-Marom, R., & Fischoff, B. (1977). Direct measures of availability and judgments of category frequency. *Bulletin of the Psychonomic Society, 9*, 236–238.

Brown, R., & McNeill, D. (1966). The "tip of the tongue" phenomenon. *Journal of Verbal Learning and Verbal Behavior, 5*, 325–337.

Brown, S. C. (1968). Role of stimulus-term and serial-position cues in constant-order paired-associates. *Journal of Experimental Psychology, 77*, 269–274.

Craik, F. I., Govoni, R., Naveh-Benjamim, M., & Anderson, N. D. (1996). The effects of divided attention on encoding and retrieval processes in human memory. *Journal of Experimental Psychology: General, 125*, 159–180.

Fernandes, M. A., & Moskovitch, M. (2000). Divided attention and memory: Evidence of substantial interference effects at retrieval and encoding. *Journal of Experimental Psychology: General, 129*, 155–176.

Fiske, S. T. (1998). Stereotyping, prejudice, and discrimination. In D. T. Gilbert, S. T. Fiske, & G. Lindzey (Eds.), *The handbook of social psychology: Vol. 2* (4th ed., pp. 357–411). Boston: McGraw-Hill.

Gabrielcik, A., & Fazio, R. H. (1984). Priming and frequency estimation: A strict test of the availability heuristic. *Personality and Social Psychology Bulletin, 10*, 85–89.

Garcia-Marques, L., & Hamilton, D. L. (1996). Resolving the apparent discrepancy between the incongruency effect and the expectancy-based illusory correlation effect: The TRAP model. *Journal of Personality and Social Psychology, 71*, 845–860.

Gilbert, D. T., & Hixon, J. G. (1991). The trouble of thinking: Activation and application of stereotypic beliefs. *Journal of Personality and Social Psychology, 60*, 509–517.

Hamilton, D. L. (1981). Stereotyping and intergroup behavior: Some thoughts on the cognitive approach. In D. L. Hamilton (Ed.), *Cognitive processes in stereotyping and intergroup behavior* (pp. 333–354). Hillsdale, NJ: Erlbaum.

Hamilton, D. L., Driscoll, D. M., & Worth, L. T. (1989). Cognitive organization of impressions: Effects of incongruency in complex representations. *Journal of Personality and Social Psychology, 57*, 925–939.

Hamilton, D. L., & Rose, T. L. (1980). Illusory correlation and the maintenance of stereotypic beliefs. *Journal of Personality and Social Psychology, 39*, 832–845.

Hamilton, D. L., & Sherman, J. W. (1994). Stereotypes. In R. S. Wyer, Jr., & T. K. Srull (Ed), *Handbook of social cognition: Vol. 2. Applications* (2nd ed., pp. 1–68). Hillsdale, NJ: Erlbaum.

Hamilton, D. L., Sherman, S. J., & Ruvolo, C. M. (1990). Stereotype-based expectancies: Effects on information processing and social behavior. *Journal of Social Issues, 46*, 35–60.

Hart, J. T. (1965). Memory and the feeling-of-knowing experience. *Journal of Educational Psychology, 56*, 208–216.

Hastie, R. (1988). A computer simulation model of person memory. *Journal of Experimental Social Psychology, 24*, 423–447.

Hastie, R., & Kumar, P. A. (1979). Person memory: Personality traits as organizing principles in memory for behaviors. *Journal of Personality and Social Psychology, 37*, 25–38.

Hicks, J. L., & Marsh, R. L. (2000). Toward specifying the attentional demands of recognition memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 26*, 1483–1498.

Higgins, E. T. (1989). Knowledge accessibility and activation: Subjectivity and suffering from unconscious sources. In J. S. Uleman & J. A. Bargh (Eds.), *Unintended thought* (pp. 75–123). New York: Guilford Press.

Higgins, E. T. (1996). Knowledge activation: Accessibility, applicability, and salience. In E. T. Higgins & A. W. Kruglanski (Eds.), *Social psychology: Handbook of basic principles* (pp. 133–168). New York: Guilford Press.

Holtgraves, T., & Srull, T. K. (1990). Ordered and unordered retrieval strategies in person memory, *Journal of Experimental Social Psychology, 26*, 63–81.

Macrae, C. N., Hewstone, M., & Griffiths, R. J. (1993). Processing load and memory for stereotype-based information. *European Journal of Social Psychology, 23*, 77–87.

Martin, D. W. (1970). Residual processing capacity during verbal organization in memory. *Journal of Verbal Learning and Verbal Behavior, 9*, 391–397.

Naveh-Benjamim, M., Craik, F. I. M., Guez, J., & Dori, H. (1998). Effects of divided attention on encoding and retrieval processes in human memory: Further support for an asymmetry. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 24*, 1091–1104.

Naveh-Benjamim, M., & Guez, J. (2000). Effects of divided attention on encoding and retrieval processes: Assessment of attentional costs and component analysis. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 26*, 1461–1482.

Nhouyvanisvong, A., & Reder, L. M. (1998). Rapid feeling-of-knowing: A strategy selection mechanism. In V. Y. Yzerbyt & G. Lories (Eds.), *Metacognition: Cognitive and social dimensions* (pp. 35–52). Thousand Oaks, CA: Sage.

Rojahn, K., & Pettigrew, T. F. (1992). Memory for schema-relevant information: A meta-analytic resolution. *British Journal of Social Psychology, 31*, 81–109.

Rundus, D. (1973). Negative effects of using list items as recall cues. *Journal of Verbal Learning and Verbal Behavior, 12*, 43–50.

Sherman, J. W., & Hamilton, D. L. (1994). On the formation of interitem associative links in person memory. *Journal of Experimental Social Psychology, 30*, 203–217.

Skowronski, J., & Welbourne, J. (1997). Conditional probability may be a flawed measure of associative strength. *Social Cognition, 15*, 1–12.

Slamecka, N. J. (1968). An examination of trace storage in free recall. *Journal of Experimental Psychology, 76*, 504–513.

Smith, R. E., & Hunt, R. R. (2000). The influence of distinctive processing on retrieval-induced forgetting. *Memory & Cognition*, 503–508.

Srull, T., K. (1981). Person memory: Some tests of associative storage and retrieval models. *Journal of Experimental Psychology: Human Learning and Memory, 7*, 440–463.

Srull, T. K., Lichtenstein, M., & Rothbart, M. (1985). Associative storage and retrieval processes in person memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 11*, 316–345.

Srull, T. K., & Wyer, R. S. (1989). Person memory and judgment. *Psychological Review, 96*, 58–83.

Stangor, C., & Duan, C. (1991). Effects of multiple task demands upon memory for information about social groups. *Journal of Experimental Social Psychology, 27*, 357–378.

Stangor, C., & McMillan, D. (1992). Memory for expectancy-congruent and expectancy-incongruent information: A review of the social and social developmental literatures. *Psychological Bulletin, 111*, 42–61.

Stern, L., Marrs, S., Millar, M., & Cole, E. (1984). Processing time and the recall of inconsistent and consistent behaviors of individuals and groups. *Journal of Personality and Social Psychology, 47*, 253–262.

Trumbo, D., & Milone, F. (1971). Primary task performance as a function of encoding, retention and recall on a secondary task. *Journal of Experimental Psychology, 91*, 273–279.

Wyer, R. S., Jr., & Srull, T. K. (1989). *Memory and cognition in its social context*. Hillsdale, NJ: Erlbaum.

Received October 10, 2000
Revision received July 6, 2001
Accepted July 14, 2001 ■