

Time versus items in judgment of recency

DOUGLAS L. HINTZMAN
University of Oregon, Eugene, Oregon

In the numerical judgment of recency (JOR) task, subjects judge how many items have intervened since the test item was previously presented. Two experiments were conducted to determine whether the basis of JOR is the age of the memory (time) or the number of intervening items. Subjects went through a long list that was made up of alternating fast blocks and slow blocks, but the block structure was disguised by probabilistic selection of a short or long intertrial interval. In both experiments, JOR was found to be a simple function of time, with no added contribution from the number of items.

It is essential to everything remembered that it be something which is past; and we cannot conceive a thing to be past, without conceiving some duration, more or less, between it and the present.

(Thomas Reid, 1785/2002, pp. 258–259)

In the numerical judgment of recency (JOR) task, experimental subjects are asked to judge how many items have intervened since the test item was previously presented. The number of items is, of course, highly correlated with time. There are pragmatic reasons for asking subjects to judge recency in terms of items. Time is experienced as continuous and is divided only by convention into arbitrary measures such as seconds and minutes. By contrast, the number of discrete items is a concept that everyone immediately grasps. Regardless of the way the question is asked, a numerical JOR might ultimately be based on an intuitive sense of how much time has passed since the test item was presented before. To transform a sense of the memory's age into a numerical JOR, one would just need an idea of the overall presentation rate.

Several theories predict that JOR should be primarily determined by time. These include a model in which JOR is based on decaying trace strength (Hinrichs, 1970) and one in which it is based on the degree to which the memory trace has consolidated (Wickelgren, 1972, 1974). Others have proposed that each memory incorporates a "time tag" (Yntema & Trask, 1963) or "temporal address" (Gallistel, 1990). One version of the latter hypothesis holds that time of encoding is recorded by reference to the states of coupled oscillators of different frequencies (G. D. A. Brown, Preece, & Hulme, 2000). To formulate a JOR, one would compare the present oscillator states with those that the test item retrieves from memory and use an estimate of overall presentation rate to transform time into items.

Other theories predict that JOR should be determined by the number of intervening items. In a 1682 lecture,

Robert Hooke (1705/1969) proposed that memory traces were anatomically arranged in a chain, according to serial order, and that recency could be inferred from a physical distance in the brain (see Hintzman, 2003b). Similar to Hooke's model, but limited to JOR over very short lags, is the idea that the subject counts or scans back to the test item, in a short-term buffer store (Atkinson & Shiffrin, 1968; Hacker, 1980; Murdock, 1972).

A different kind of hypothesis holds that JOR is based on changes in cognitive context. The test item retrieves contextual information from memory, and the retrieved context is compared with the current context—the greater the difference, the greater the judged lag (Hintzman, 2002). If one adopts the assumption that changes in cognitive context are driven primarily by task-related processing of stimuli (see Block, 1990; Howard & Kahana, 2002), then this view predicts that JOR will be determined mainly by the number of intervening items.

The present experiments were done to determine whether JOR is more closely related to items or to time. Subjects went through a long, continuous list of stimuli, making old–new recognition judgments to each and a numerical JOR to each item they judged to be old (see Hintzman, 2001, 2002, 2003a). The list was organized into alternating fast blocks and slow blocks, but the block structure was disguised by making selection of the intertrial interval (ITI) probabilistic. In both types of block, some trials were followed by a short (500-msec) ITI and others by a long (2,500-msec) ITI. The only difference between fast blocks and slow blocks was in the probabilities of the two ITIs.

Two subjective observations may be worth mentioning at the outset, based solely on the investigator's experience in piloting the experiments. The first is that although a subject is aware of variation in ITI, there are no apparent clues to the block structure of the list. The second is that the extra 2,000 msec in a long ITI does not encourage rehearsal or other strategies. One's typical mental state is simply that of anticipation of the next stimulus—fairly close to the "mental vacuum" desired to experimentally contrast time versus items (Reitman, 1974).

Correspondence concerning this article should be addressed to D. L. Hintzman, Department of Psychology, University of Oregon, Eugene, OR 97403 (e-mail: hintzman@uoregon.edu).

EXPERIMENT 1

Subjects went through a long list of nouns that was organized into alternating fast and slow 25-item blocks. In fast blocks, the probability of a short ITI was .8, and that of a long ITI was .2. In slow blocks, the two probabilities were reversed. Experimental items were repeated within a block at lags of 5, 11, and 20 items. The temporal durations of these lags (the sums of the ITIs and the subject's response times) were also recorded.

Method

Subjects. Twenty-two University of Oregon undergraduates participated for course credit. They were tested individually in sessions lasting approximately 45 min.

Materials. The stimuli were drawn from a pool of 436 common English nouns, 4–10 letters in length. They were selected at random, without replacement, and assigned to a fixed order of conditions in a 550-item list.

The list consisted of 22 blocks of 25 items each. Each block followed the same pattern, which was constructed to meet these constraints: (1) The first and second trials for 9 items occurred within the same block. Three of these items were repeated at each of the lags of 5, 11, and 20, where *lag* is defined as the difference between the serial positions of Trial 1 and Trial 2. (2) No words that were adjacent on Trial 1 were also adjacent on Trial 2. (3) There were never more than 3 new items (Trial 1) or old items (Trial 2) in succession. (4) So that block boundaries were not demarcated by long runs of old items followed by new items, seven positions were occupied by filler items whose repetition patterns spanned successive blocks. The repetition lags of these items were 5, 13, 16, and 25. Either 3 or 4 of these fillers were new (Trial 1) in each block after the first.

Procedure. The procedure followed that of Hintzman (2002, 2003a). The subjects were seated before a Macintosh IICI computer controlling a full-page monitor and told that they would see a long list of words in which many words would be repeated. Their first task was to decide whether each word was new or old in the list, by

responding on the “Z” or “?” key of the keyboard. If they decided that a word was old, they would also be asked to judge how many items back in the list the word was when it had been previously encountered. They were told that a word could be repeated after 5, 10, 15, 20, 25, or 30 intervening items, and that the JORs, or lag judgments, were to be made using the 1–6 keys of the numerical keypad, which had been relabeled with the appropriate digits. The subjects were also warned that the computer would sometimes hesitate before presenting the next word and sometimes present the next word immediately. They were told to ignore this and always judge how many items back the word was when they had seen it before.

Each test word was displayed in 48-point, lowercase Helvetica font, in the center of the screen. When the word appeared, the subject first was asked to indicate whether it was old or new. The response alternatives (“Z” = *new* and “?” = *old*) appeared beneath the word on the screen. If the subject responded “old,” this instruction was immediately replaced by the two lines: “How many items back?” and “5 10 15 20 25 30.” Illegal keypresses resulted in a warning and a repeated request for a response. Each display remained on the screen until the subject responded.

Each trial was followed by a blank ITI of either 500 or 2,500 msec, assigned probabilistically. In fast blocks, $p(500) = .8$, and $p(2,500) = .2$. In slow blocks, the two probabilities were reversed. Half the subjects started with a fast block, and half with a slow block, with the two block types alternating thereafter.

Results

The data of 1 subject were lost because of equipment failure, and those of another subject were dropped because of a low recognition hit rate (under 80%). The following analyses are based on the data from the 20 remaining subjects. The first two blocks (50 trials) of each subject's data were treated as practice and were not analyzed. This resulted in 30 trials per condition per subject, or a total of 600 trials per condition.

Recognition. The mean false alarm rate was 4.1%, and the mean hit rate was 95.5%. Possibly because per-

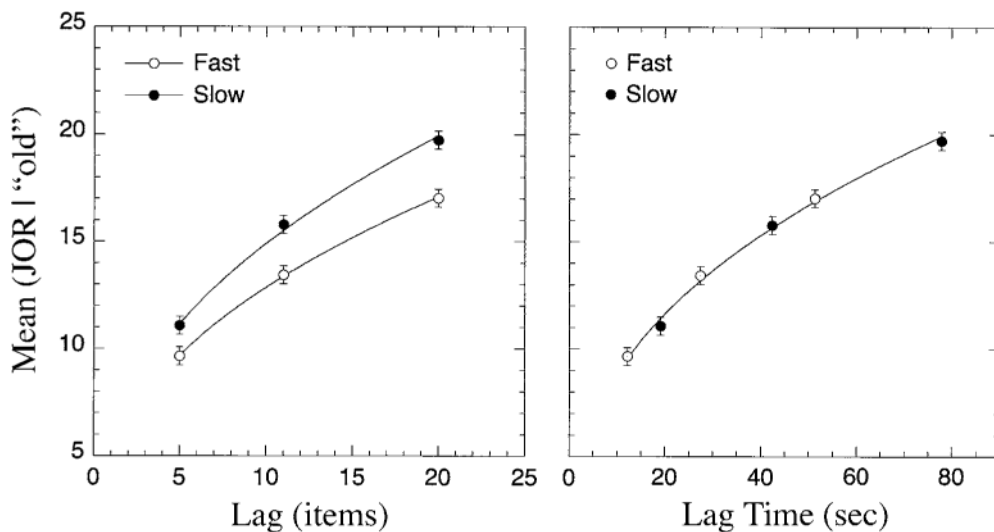


Figure 1. Mean JOR from Experiment 1, plotted as a function of test lag, measured in items (left panel) and time (right panel). Best-fitting power function in the right panel is $JOR = 3.56t^{0.40}$. Error bars show *SE* on the basis of the subject \times condition mean square (see Loftus & Masson, 1994).

formance was so close to ceiling, the hit rates showed no systematic effect of rate or lag or rate \times lag interaction (all $ps > .10$).

JOR. The left panel of Figure 1 presents mean JOR plotted against lag, measured in items. These data were analyzed using planned contrasts—dichotomous for rate and linear for lag—following the method outlined by Rosenthal and Rosnow (1985). Mean JOR (the judged number of items, or lag) increased with lag and was longer in slow blocks than in fast blocks ($t > 8.0$ for both contrasts). The lag \times rate interaction was also significant [$t(19) = 2.350, p = .0297$].

In the right panel of Figure 1, the same data are plotted against lag, measured as the average time from the offset of Trial 1 to the onset of Trial 2.¹ As can be seen, a single power function describes the data from both block types quite well, accounting for 99.6% of the variance. These data are consistent with the hypothesis that JOR subjects use the retrieved memory's age to estimate the number of intervening items.

EXPERIMENT 2

This experiment was a conceptual replication of Experiment 1, with several parameters changed. The main changes were that: (1) the blocks were twice as long, (2) there were four experimental lags, ranging from 9–43, (3) the six JOR alternatives ranged from 10–60, instead of from 5–30, (4) the stimuli were male and female first names, and (5) to prevent the subjects from noticing rate changes in these longer blocks, the probabilities of a short ITI in fast blocks and slow blocks were made more alike (.7 and .3, respectively).

Method

Subjects. Eighteen subjects were recruited in the same way as in Experiment 1.

Stimuli. The stimulus pool was a set of 400 male and female first (given) names, taken from a U.S. Census Web site. The pool had been edited beforehand to remove close variants and similar-appearing names.

Materials and Procedure. The experimental list of 550 items was divided into 11 blocks of 50 trials each. Again, a fixed structure defined each block. For 16 of the items, Trial 1 and Trial 2 occurred within the same block, at lags of 9, 18, 28, and 43 items. Four names were assigned to each of these lags. There were also two filler names that occurred only once, and eight repeated fillers for which Trial 1 and Trial 2 fell in successive blocks. Repetition lags for these fillers ranged from 31–43.

The names appeared in 48-point uppercase Times font, and the subjects were told to respond on keys 1–6 of the numerical keypad, by multiplying by 10—thus the JORs ranged from 10–60 in steps of 10. Blank ITIs were again 500 and 2,500 msec. In the first block, $p(500) = .5$, but thereafter it alternated between .7 in fast blocks and .3 in slow blocks. In other respects, the procedure was identical to that of Experiment 1.

Results

The data of 2 subjects, who had recognition hit rates below 80%, were deleted from further analysis, leaving $n = 16$. Again, the first 50 trials were treated as practice and ignored. The following analyses are thus based on 20 observations per condition per subject, or a total of 320 observations per condition.

Recognition. The false alarm rate on Trial 1 was 8.0%, and the overall hit rate on Trial 2 was 93.3%. In contrast to Experiment 1, analysis of the hit rates yielded an effect of lag [$t(15) = 8.48$] and a lag \times rate interaction [$t(15) = 3.02, p = .009$]. The means are shown in Figure 2, fitted with straight lines to emphasize the form of

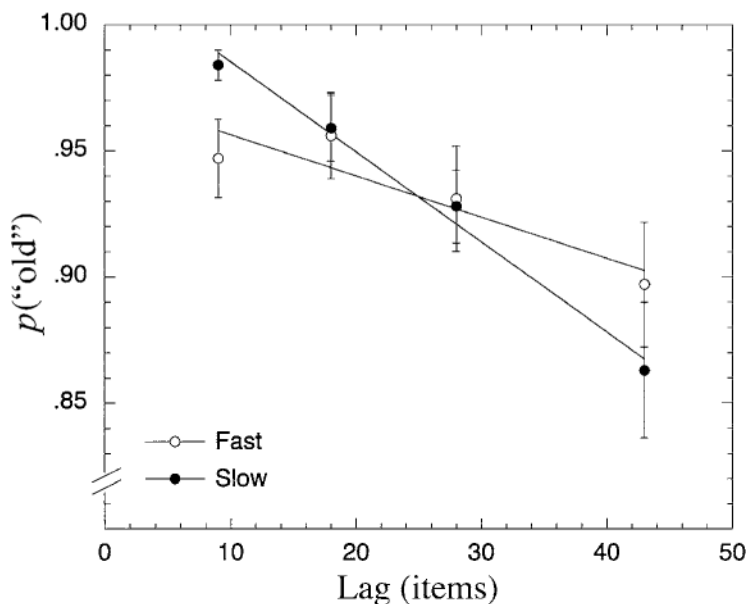


Figure 2. Probability of a recognition hit, from Experiment 2, as a function of rate and test lag. Error bars show SEM.

the interaction. At lag 9, hit rates were lower in fast blocks than in slow blocks [$t(15) = 2.67, p = .018$], while at lag 43 the difference was in the opposite direction [$t(15) = .111$].

This interaction, which was not anticipated, can be understood as the result of two different effects of presentation rate. First, the ITI immediately following Trial 1 in slow blocks was usually long, allowing the item more time to be encoded into memory. Second, retention intervals in slow blocks were longer, allowing more time for forgetting, especially over the longer lags. It may be worth mentioning that the form of this interaction is similar to that seen in the classic probe-digit experiment of Waugh and Norman (1965), which compared the effects of fast and slow presentation rates on cued recall.² Taken together, the two results suggest that it might be fruitful for future research to reexamine the issue of rate effects on forgetting.

JOR. In Figure 3, mean JOR is plotted as a function of lag measured in items (left panel) and in time (right panel). The difference between fast- and slow-block means was smaller than it was in Experiment 1, almost certainly because the difference in ITI probabilities was smaller. In this experiment, slow blocks were on average 28% longer in duration than fast blocks, whereas in the previous experiment, slow blocks averaged 55% longer than fast blocks. Otherwise, the pattern of JOR means was the same: As in Experiment 1, mean JOR plotted against items showed effects of lag and rate ($t > 5.5$ for both) and a rate \times lag interaction [$t(15) = 2.476, p = .028$]. When mean JOR is plotted against time, a single power function accounts for 99.0% of the variance.

DISCUSSION

These experiments show that when subjects are asked to judge how many items have intervened since a test

item was previously studied, they base their judgments on the apparent age of the retrieved memory. To map a trace's perceived age onto the number of intervening items requires that it be scaled according to the rate at which the items are being presented. Under the present experimental conditions, the subject's assessment of rate apparently was not local, but global—that is, it spanned multiple blocks.

Undoubtedly, there are boundary conditions on this finding. For example, subjects surely know whether they spent the last 2–3 sec processing items or passively waiting, so one would not expect the strict dependence on time to extend down to lags of 1 or 2 items. Also, there is little reason to doubt that subjects could recalibrate JOR mappings on the basis of local variation in rate if the rate changes were obvious, but here they were not. Subjects were aware of variation in ITI, but not that there were alternating fast blocks and slow blocks.

The conclusion that subjects judge recency on the basis of time rather than items may appear to conflict with the literature on retrospective judgments of duration. In those experiments, subjects are asked to judge the length of a remembered interval that was filled by a distinctive task (e.g., presentation of a list). Early findings led Ornstein (1969) to conclude that remembered duration increases with the number and complexity of stimuli that can be retrieved from the judged interval. Later results suggest instead that the determining factor is not the amount of information recalled, but the complexity of the sequence of cognitive operations that were performed in the interval (for a review, see Block, 1990). Under either interpretation, one would not anticipate the present outcome, where the number of items—and accordingly the amount and complexity of processing—had no discernable effect beyond that attributable to time.

This apparent discrepancy can be resolved by considering differences between the tasks. A retrospective




Figure 3. Mean JOR from Experiment 2, plotted as a function of test lag, measured in items (left panel) and time (right panel). Best-fitting power function in the right panel is $JOR = 4.33t^{0.41}$. Error bars show SE on the basis of the subject \times condition mean square (see Loftus & Masson, 1994).

judgment of duration concerns a past interval marked by a distinct beginning and end. It is likely to be stored as a “chunk” that refers to the interval’s characteristics, including its contents. In the JOR task, which is quite different, the subject experiences a long, monotonous series of events and is arbitrarily cued to retrieve one of them and judge the number of events separating it from the present moment. The beginning of the list is too far in the past to be a useful landmark, and the target interval, *per se*, does not have a representation in memory. The definitions of the target interval’s beginning, contents, and end constantly change as the experiment progresses. Under these conditions, subjects can do little but base their judgments on some retrieved property (or properties) of the target item’s memory.

Controlled JOR experiments such as these appear to tap just the leading edge of a subjective recency scale that spans several orders of magnitude, from seconds and minutes at one end, to years and even decades at the other (Friedman, 1993; Gallistel, 1990, chap. 15). To a first approximation, at least, this “distance-based” scale appears to follow a negatively accelerated (roughly exponential) function over its entire range. (For data on JOR over intervals of minutes to hours, see Wickelgren, 1972; for weeks to months, see Friedman, 1991; for months to years, see Linton, 1975, Squire, Chase, & Slater, 1975, and Underwood, 1977.)³ Such a scale is useful for making crude, order-of-magnitude judgments such as how many months ago one has read a certain article or whether one has seen a particular colleague today, but it lacks fine-grained temporal information. To make more exact judgments of time and order, people ordinarily supplement the crude recency scale with various reasoning and retrieval strategies, which make it difficult to study the distance-based scale in isolation. Nevertheless, in the course of mental development, the more sophisticated temporal concepts may all derive from the primitive sense of a memory’s age (Friedman, 1993).

Various lines of evidence suggest that the recency scale plays a powerful role in forgetting. As two memories drift back along the negatively accelerated curve, it becomes increasingly difficult to discriminate their relative ages, and they become more likely to mutually interfere in recall—particularly if they are similar on other relevant dimensions. Operating over seconds to minutes, this principle is exemplified by joint effects of presentation rate and serial position on the recency segment of the free recall function (see, e.g., Bjork & Whitten, 1974; Glenberg, 1987; Glenberg, Bradley, Kraus, & Renzaglia, 1983) and by effects of ITI and retention interval on proactive interference in the Brown–Peterson task (e.g., Bennett, 1975; Greene, 1992, chap. 4). Studies of proactive interference in paired-associate recall suggest that the same principle also operates over intervals of days (Alin, 1968; Underwood & Freund, 1968). That is, the crucial temporal variable appears to be not the length of the retention interval *per se*, but the ratio of the retention interval to the interlist interval.

Although the present experiments establish the importance of time in JOR, other studies using the same running judgment task show that time is not the only contributing factor. Hintzman (2003a) found JOR to be shorter for rare words than for common words, and JOR shorter for concrete words than for abstract words. In both cases, the difference was small, and the word category that was judged more recent was the one that supports better performance in recognition memory (e.g., Glanzer & Adams, 1985).⁴ This does not appear to be an artifact of recognition failure, because in this paradigm, JOR is contingent on reporting that the test item is “old.”

In addition, Hintzman (2002) compared JORs for names that had the same physical appearance at study and test with those for names whose appearance had changed. As predicted by a context-change explanation of JOR, the judgments were slightly shorter when the study and test trials physically matched than when they were different. It is important that in each of these cases, the direction of the difference was the same across the entire range of studied lags. Because the general tendency is for JOR to overestimate short lags and underestimate long ones, this means that the obtained effects cannot be characterized as overall differences in JOR accuracy; the condition that was more accurate at short lags was less accurate at long lags.

A theorist attempting to explain these effects seems to have two basic options: (1) to assume that all the effective manipulations affect a single underlying process or pathway, which exclusively determines JOR, or (2) to assume that JOR is mediated in parallel by more than one pathway. Let us first consider single-pathway hypotheses.

One idea is that in formulating a JOR, the subject retrieves contextual information from the memory trace and compares it with the cognitive context prevailing at the time of test (Hintzman, 2002). The more similar these contexts are, the shorter is the JOR. This hypothesis correctly predicts shorter JORs for items studied and tested in the same physical form (Hintzman, 2002). The effects of word class may be easily explained by assuming that rare and concrete words are better associated with context than are common and abstract words—perhaps because they attract more attention at study (e.g., Glanzer & Adams, 1990). However, to accommodate the present results, one would have to argue that cognitive context changes just as rapidly during a long ITI, when one is passively expecting the next stimulus to appear, as it does when one is actively processing and responding to items. There may be no way to definitively rule out this assumption, but it appears inconsistent with current notions of cognitive context (e.g., Block, 1990; Howard & Kahana, 2002).

Another possibility is that JOR is based on the degree to which the memory trace is activated by the test item, and that activation is a joint function of trace strength and similarity. In this view, the effect of time on JOR arises from time-dependent decay (Hinrichs, 1970). Physical attributes of the retrieval cue contribute to its

matching the trace, and traces of rare words and concrete words are stronger than those of the contrasting categories, perhaps because of heightened attention at encoding. Although this approach seems attractive, it faces several difficulties. In particular, recognition confidence ratings—which might also be assumed to reflect activation strength—change much more slowly than does JOR as a function of lag. Direct comparison of the two measures shows that a single unidimensional variable cannot simultaneously account for recognition confidence and JOR (Hintzman, 2003a). A dissociation between recognition and JOR is also seen in the data from the present Experiment 2, where presentation rate produced a cross-over interaction on recognition hit rate (Figure 2) but a main effect on JOR (Figure 3).

An additional problem with the strength hypothesis is that a decaying function must eventually approach zero, which would severely limit its effective range. It is hard to see how a “strength” that decays rapidly enough to mediate the discrimination of recencies measured in seconds could also discriminate recencies measured in months and years. It might be necessary to postulate several strengths, with drastically different decay rates, which support JOR over different scales of time. (McGaugh, 2000, summarizes neurobiological evidence that there are at least three separate traces, with respective durations of seconds to hours, hours to months, and months to a lifetime.)

The present data particularly favor hypotheses in which the basis of JOR is some factor directly related to time. One example is Wickelgren’s (1972, 1974) proposal that JOR is based on the degree of consolidation of the memory trace. In Wickelgren’s theory, consolidation (or resistance to forgetting) was assumed to occur strictly as a function of time. Another example is the proposal of G. D. A. Brown et al. (2000) that JOR is based on the difference between the present states of neural oscillators and those that are associated with the retrieved memory. Although these mechanisms explain why JOR is affected by time rather than by items, they do not suggest why it should also be influenced by a word’s frequency, concreteness, or physical appearance.

To explain such results with a theory incorporating a time-based process, it seems necessary to assume that another process separately affects JOR. This does not seem unreasonable, when one considers the low level of accuracy that the time-based scale presumably affords. If subjects know that the time-based code is very imprecise, they may also give some weight to another factor that they feel is correlated with recency. One such factor would be trace activation. If word frequency, concreteness, and context matching affect trace activation in the ways previously argued, that would explain why their effects are small but consistent over a range of recencies. It would also help explain why subjects give shorter JORs to items that they claim to recollect than to those that just seem familiar (Hintzman, 2001), and—on a much longer time scale—why past news events that people know more about are judged as being more recent

than are age-matched events that they know less about (N. R. Brown, Rips, & Shevell, 1985). Such results may be interpreted as showing that the amount of retrieved information is a factor contributing to JOR, but they do not imply that it is the only factor, or even that it is the most basic one.

Thus, it seems reasonable to propose that JOR is mediated primarily by strictly time-related information, and secondarily by trace activation. An apparent advantage of such a two-pathway hypothesis is that one of the processes (consolidation, oscillator states, or some other time-based process yet to be proposed) can be assumed to underlie recency judgments over time scales ranging from seconds to years. Of course, it is also possible that apparent recencies for different time scales are based on entirely different mechanisms.

CONCLUDING COMMENT

These findings show that subjects who are asked to judge an item’s recency in terms of intervening items do so with reference to a subjective sense of time. This contrasts with the literature on forgetting, where time has been generally credited with a minor role compared with that of items. To give two prominent examples, McGeoch (1932) famously argued that long-term forgetting was caused not by “disuse,” but by the interpolated learning of similar materials; and Waugh and Norman (1965) concluded that short-term forgetting was caused not by passive decay, but by displacement of the target item from primary memory (but see note 2). Readers familiar with that literature may consider the present findings surprising, as I did—but the quote that opens this article suggests that they would not have surprised Thomas Reid, the 18th-century Scottish philosopher.

REFERENCES

- ALIN, L. H. (1968). Proactive inhibition as a function of the time interval between the learning of the two tasks and the number of prior lists. *Journal of Verbal Learning & Verbal Behavior*, *7*, 1024-1029.
- ATKINSON, R. C., & SHIFFRIN, R. M. (1968). Human memory: A proposed system and its control processes. In K. W. Spence & J. T. Spence (Eds.), *The psychology of learning and motivation* (Vol. 2, pp. 89-195). New York: Academic Press.
- BENNETT, R. W. (1975). Proactive interference in short-term memory: Fundamental forgetting processes. *Journal of Verbal Learning & Verbal Behavior*, *14*, 123-144.
- BJORK, R. A., & WHITTEN, W. B. (1974). Recency-sensitive retrieval processes in long-term free recall. *Cognitive Psychology*, *6*, 173-189.
- BLOCK, R. A. (1990). Models of psychological time. In R. A. Block (Ed.), *Cognitive models of psychological time* (pp. 1-35). Hillsdale, NJ: Erlbaum.
- BROWN, G. D. A., PREECE, T., & HULME, C. (2000). Oscillator-based memory for serial order. *Psychological Review*, *107*, 127-181.
- BROWN, N. R., RIPS, L. J., & SHEVELL, S. K. (1985). The subjective dates of natural events in very-long-term memory. *Cognitive Psychology*, *17*, 139-177.
- FRIEDMAN, W. J. (1991). The development of children’s memory for the time of past events. *Child Development*, *62*, 139-155.
- FRIEDMAN, W. J. (1993). Memory for the time of past events. *Psychological Bulletin*, *113*, 44-66.

- GALLISTEL, C. R. (1990). *The organization of learning*. Cambridge, MA: MIT Press.
- GLANZER, M., & ADAMS, J. K. (1985). The mirror effect in recognition memory. *Memory & Cognition*, **13**, 8-20.
- GLANZER, M., & ADAMS, J. K. (1990). The mirror effect in recognition memory: Data and theory. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **16**, 5-16.
- GLENBERG, A. M. (1987). Temporal context and memory. In D. S. Gorfein & R. R. Hoffman (Eds.), *Memory and learning: The Ebbinghaus centennial conference* (pp. 173-190). Hillsdale, NJ: Erlbaum.
- GLENBERG, A. M., BRADLEY, M. M., KRAUS, T. A., & RENZAGLIA, G. J. (1983). Studies of the long-term recency effect: Support for a contextually guided retrieval hypothesis. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **9**, 231-255.
- GREENE, R. L. (1992). *Human memory: Paradigms and paradoxes*. Hillsdale, NJ: Erlbaum.
- HACKER, M. J. (1980). Speed and accuracy of recency judgments for events in short-term memory. *Journal of Experimental Psychology: Human Learning & Memory*, **6**, 651-675.
- HINRICH, J. V. (1970). A two-process memory-strength theory for judgment of recency. *Psychological Review*, **77**, 223-233.
- HINTZMAN, D. L. (2001). Judgments of frequency and recency: How they relate to reports of subjective awareness. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **27**, 1347-1358.
- HINTZMAN, D. L. (2002). Context matching and judgments of recency. *Psychonomic Bulletin & Review*, **9**, 368-374.
- HINTZMAN, D. L. (2003a). Judgments of recency and their relation to recognition memory. *Memory & Cognition*, **31**, 26-34.
- HINTZMAN, D. L. (2003b). Robert Hooke's model of memory. *Psychonomic Bulletin & Review*, **10**, 1-14.
- HOOKE, R. (1969). *The posthumous works of Robert Hooke: With a new introduction by Richard S. Westfall*. New York: Johnson Reprint Corporation. (Original work published 1705)
- HOWARD, M. W., & KAHANA, M. J. (2002). A distributed representation of temporal context. *Journal of Mathematical Psychology*, **46**, 269-299.
- LINTON, M. (1975). Memory for real-world events. In D. A. Norman & D. E. Rumelhart (Eds.), *Explorations in cognition* (pp. 376-404). San Francisco: Freeman.
- LOFTUS, G. R., & MASSON, M. E. J. (1994). Using confidence intervals in within-subject designs. *Psychonomic Bulletin & Review*, **1**, 476-490.
- MAYLOR, E. A., CHATER, N., & BROWN, G. D. A. (2001). Scale invariance in the retrieval of retrospective and prospective memories. *Psychonomic Bulletin & Review*, **8**, 162-167.
- MCCAUGH, J. L. (2000). Memory—a century of consolidation. *Science*, **287**, 248-251.
- MCGEOCH, J. A. (1932). Forgetting and the law of disuse. *Psychological Review*, **39**, 352-370.
- MURDOCK, B. B., JR. (1972). Short-term memory. In G. H. Bower (Ed.), *The psychology of learning and motivation* (Vol. 5, pp. 67-127). New York: Academic Press.
- ORNSTEIN, R. E. (1969). *On the experience of time*. Harmondsworth, U.K.: Penguin.
- REID, T. (2002). *Essays on the intellectual powers of man* (D. R. Brookes, Ed.) University Park: Pennsylvania State University Press. (Original work published 1785)
- REITMAN, J. S. (1974). Without surreptitious rehearsal, information in short-term memory decays. *Journal of Verbal Learning & Verbal Behavior*, **13**, 365-377.
- ROSENTHAL, R., & ROSNOW, R. L. (1985). *Contrast analysis*. Cambridge: Cambridge University Press.
- SQUIRE, L. R., CHASE, P. M., & SLATER, P. C. (1975). Assessment of memory for remote events. *Psychological Reports*, **37**, 223-234.
- UNDERWOOD, B. J. (1977). *Temporal codes for memories*. Hillsdale, NJ: Erlbaum.
- UNDERWOOD, B. J., & FREUND, J. S. (1968). Effect of temporal separation of two tasks on proactive inhibition. *Journal of Experimental Psychology*, **78**, 50-54.
- WAUGH, N. C., & NORMAN, D. A. (1965). Primary memory. *Psychological Review*, **72**, 89-104.
- WICKELGREN, W. A. (1972). Trace resistance and the decay of long-term memory. *Journal of Mathematical Psychology*, **9**, 418-455.
- WICKELGREN, W. A. (1974). Single-trace fragility theory of memory dynamics. *Memory & Cognition*, **2**, 775-780.
- YNTEMA, D. B., & TRASK, F. P. (1963). Recall as a search process. *Journal of Verbal Learning & Verbal Behavior*, **2**, 65-74.

NOTES

1. Because of a programming error, lag times could not be obtained for half the subjects. The values on the abscissa were computed from the 10 subjects who were tested after the error had been corrected.

2. Waugh and Norman (1965) concluded that forgetting in the probe-digit task was determined primarily by items and not by time. This was determined on the basis of a nonsignificant rate \times lag interaction ($F < 1$)—apparently an omnibus F with 8 and 3 df . Their article presented individual data from 4 subjects. When the data are analyzed using planned contrasts, the interaction is statistically reliable [$t(3) = 4.279, p = .023$]. Thus, the results show a significant effect of time, contrary to their original conclusion.

3. For a different perspective on “time-scale invariance,” see Maylor, Chater, and Brown (2001).

4. This characterization is further confirmed by an unpublished experiment from my laboratory that compared JOR for color photographs and names. Across all lags, pictures were judged to be more recent than names.

(Manuscript received February 2, 2004;
revision accepted for publication March 5, 2004.)