THE QUARTERLY JOURNAL OF EXPERIMENTAL PSYCHOLOGY, 1994, *47A* (1) 161-191

Overlapping Mental Operations in Serial Performance with Preview

Harold Pashler

University of California San Diego, La Jolla, California, U.S.A.

Dual-task research has revealed a response-selection bottleneck: response selection cannot occur simultaneously in two different tasks, though perceptual processing may overlap response selection (Welford, 1952; Pashler & Johnston, 1989). In serial performance with preview, the same task is performed repeatedly, but stimuli are available well before response. Does the dual-task bottleneck limit the rate of responding in this situation, despite the fact that task set need not change? Four experiments examined manual responses to letters, varying stimulus preview. Preview increased RTs for the first response and reduced the subsequent inter-response intervals (IRIs). Preview also reduced the effect of visual intensity and response duration on IRIs, whereas effects of stimulus-response mapping variables were unchanged. These results indicate that the response-selection bottleneck limits serial performance, just as it limits concurrent performance of two unrelated tasks. This implies that the response-selection bottleneck is not caused by the need to switch task set.

Many studies have examined the nature of the interference that arises when people attempt to perform two different tasks concurrently. Dual-task performance is of practical significance and has important theoretical implications about basic cognitive architecture. A great deal of human activity, however, involves *serial,* rather than dual-task performance. That is, the same task is performed over and over again, rather than two different tasks being performed concurrently. Typing and reading aloud are two obvious examples of serial performance. Human behaviour outside the laboratory is replete with other, less obvious examples, such as walking in a cluttered landscape, picking up a number of objects from the floor, or

---

Requests for reprints should be sent to H. Pashler, Department of Psychology C-0109, University of California San Diego, La Jolla, CA 92093, U.S.A., or send electronic mail to hpashler@ucsd.edu.

The work was supported by a grant to H. Pashler from NIMH (1-R29-MH45584-01). Thanks are due to Alann Lopes for excellent technical assistance with this project, to Allison Titus and Mike Lostetter for help in collecting the data, and to Clark Fagot and J. Toby Mordkoff for helpful comments and discussion.

painting a wall. In each of these cases, there is *stimulus preview,* i.e. the individual can potentially see stimuli well before the response to those stimuli are produced (in the three examples given, the primary stimuli are obstacles, objects on the floor, or half-painted patches of a wall). This makes it possible that certain mental operations involved in dealing with the successive stimuli may overlap in time (a possibility that does not exist when there is no preview).

Despite the fact that serial performance with preview is so widespread, surprisingly little research has been directed at analyzing the sequence (and possible overlap) of mental operations in such situations. The few exceptions include brief observations by Cattell (1886) and a report by Leonard (1953), which appeared in the *Quarterly Journal* almost 40 years ago. These studies provide the starting point for the present work.

Serial performance with preview is important not only because it is ubiquitous outside the laboratory, but also because it may illuminate the nature of the basic limits on simultaneous cognitive processes. In dual-task performance, a separate response mapping is used in each task. Thus, it may be that shifts in mental set are necessary in this situation. By contrast, no shift in task set is involved in serial performance. For that reason, serial performance provides a critical point of contrast for analyzing interference in the usual dual-task studies, because it may illuminate the possible contribution of changes of tasks set to dual-task interference.

Sources of Dual-Task Interference

The theoretical perspective guiding the studies reported below comes from an analysis of the nature of dual-task interference with tasks involving two distinct tasks performed concurrently. When people perform two stimulus-response tasks involving two stimuli separated by a short delay (the stimulus onset asynchrony, or SOA), the response to the second stimulus is slowed down, and the slowing is greater, the shorter the SOA (Telford, 1931). This phenomenon is often termed the psychological refractory period (or PRP) effect.

Welford (1952) first attributed the PRP effect to a bottleneck in response selection, causing second-task response selection to be postponed while the first-task response is being selected (see Figure 1). Although some early work seemed to favour this theory, its validity was subject to much dispute during the 1960s and 1970s (Bertelson, 1966; Smith, 1967).

In more recent work on PRP effect, fairly detailed (and sometimes counterintuitive) predictions have been derived from Welford's theory, and these predictions have been confirmed in several laboratories. One approach (which was utilized in the studies reported here) involves manipulating the duration of component stages in the two tasks (cf. Schweickert, 1978; Sternberg, 1969). When stages in the *first* task prior to the end of

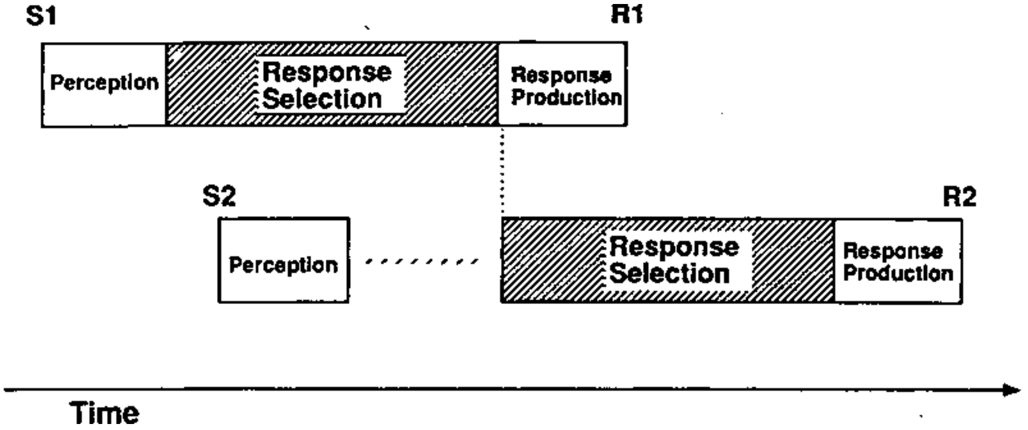

FIG. 1. The response-selection bottleneck model for dual-task performance: selection of the response in Task 2 cannot begin until the selection of the response in Task 1 is completed; other stages can overlap.

response selection are slowed, second task responses are correspondingly slowed at very brief SOAs (Pashler, 1984b; Smith, 1969). When response *production* of the first task (a stage *after* the bottleneck) is delayed, little of the slowing propagates onto the second task, regardless of SOA (Pashler & Christian, 1991). (The latter result has been observed with manual/vocal combinations, but not with two finger responses; in the manual/manual situation, an additional bottleneck in response production may exist.)

The most distinctive prediction of the model shown in Figure 1 is that when perceptual stages of the second task (*prior* to the response-selection bottleneck stage) are slowed, the presence of a bottleneck beyond these stages should reduce their effect at short SOAs. The reason for this is simply that under these circumstances the selection of the second response cannot commence when perception is complete, but must instead wait for the bottleneck mechanism to be freed up. This prediction has been confirmed in a number of experiments (Pashler, 1984b; Pashler & Johnston, 1989). Furthermore, when the bottleneck stage in the second task (i.e. selection of the second response) is slowed, this has its full effect on RT2, whatever the SOA (McCann & Johnston, 1992; Pashler, 1984b; Pashler & Johnston, 1989). This is simply because once the bottleneck operations commence, any slowing will delay the response.

These successful predictions confirm the existence of a bottleneck in the selection of responses. They rule out a bottleneck in perceiving stimuli in different modalities or in producing pairs of responses, at least when the responses are dissimilar (Fagot & Pashler, 1992; Pashler, 1993). Further confirmation of the response-selection bottleneck comes from experiments that combine a speeded first task with an unspeeded second task involving a brief display and show that the accuracy of perception of the brief display is little affected when the speeded task is performed concurrently (Pashler, 1989; 1990). Still further support for the model comes from studies examining how the speed of the second response depends on the speed of the first (e.g. Pashler, 1989; in press).

The conclusion from this work can be stated simply: if people have to perform two different choice tasks concurrently, they are unable to choose the response in one task while they are choosing a response in the other task. This generalization seems to apply to a great variety of different kinds of motor responses (Brebner, 1977; Pashler & Christian, 1991), although certain eye movements may prove to be an exception (Pashler, Carrier, & Hoffman, 1993). The effect is not easily eliminated by extended consistent practice (Gottsdanker & Stelmach, 1971).

Reasons for the "Bottleneck"

Why should any such processing limitation exist? The most obvious suggestion would be that some single mental mechanism is required every time a response is chosen; while this mechanism is choosing one response, the selection of another response must wait. There is another possible account, however; that it is simply not possible to maintain the response mappings for the two tasks active at the same time. The response-selection machinery might be perfectly capable of looking up two responses in parallel, but incapable of simultaneously maintaining the mappings for both tasks in a proper state of readiness. The fact that response selection proceeds serially might stem from the need to "re-program" the response selection rules for the second task once the first response has been selected.

There is an obvious critical test for the hypothesis that serial response selection arises because of the need to keep two different mappings readied simultaneously: simply to require the individual to produce multiple responses selected with the *same* task mapping. That is, the subject should perform the same task repeatedly, with stimuli presented early enough to reveal sequential operation of response selection, if this should occur. For this purpose, one could employ a PRP experiment with Task 1 and Task 2 being identical. There is no reason, however, to confine the examination to a sequence of only two such responses, especially given the fact that people more commonly do the same task more than twice in a row. Thus, in order to test whether the response-selection bottleneck occurs when there is no change in mapping, one is led back to the case of serial performance with preview, as first investigated by Cattell and by Leonard.

It should be noted that whether or not response selection queuing reflects an inability to prepare two mappings at once, preparation factors must account for *some* of the slowing observed in PRP experiments. Consider the common finding that responses to the second task are slower than single task control times, even when the second stimulus comes after the first *response* (Welford, 1980; Wilkinson, 1990). Presumably, this slowing cannot be attributed to the response-selection bottleneck. The most natural explanation (consistent with the response-selection bottleneck) is that when the subject must prepare for two tasks, neither can be performed as efficiently as it would be alone (even when the SOA is long enough that the bottleneck plays no role). So even if the response-selection bottleneck *itself* is not caused by the inability to keep two different tasks prepared at the same time, variation in degrees of preparation is likely to be a factor in producing dual-task slowing on top of that attributable to the bottleneck.

Serial Performance and Preview

Cattell (1886) had people read aloud a series of letters exposed by a revolving drum. He found that they were faster at doing this when they had the opportunity to preview several letters, and he concluded that the subject was "finding the name of the letter just gone by at the same time that he was seeing the letter then in view." Leonard (1953) required his subjects to perform a serial task involving moving a stylus to touch an illuminated target. The stylus moved along a track shaped like a cross, and four target lights were located at the four extreme ends of the cross. In the zero preview condition, the subject moved to touch an illuminated target light, then moved the stylus back to the centre. When the stylus reached the centre, the next target light was immediately illuminated. In the preview condition, as soon as the subject touched one target, the following target light was illuminated (but the subject had to return to the centre before proceeding to the next contact).

Leonard found that subjects were considerably faster in completing a sequence of 100 target contact when preview was provided. Subjects completed one response per 0.82 sec with preview, compared to one response per 1.04 sec without preview. Leonard observed that the speedup was mostly attributable to a reduction in the time for which the stylus dwelled at the centre, rather than a change in the speed of hand movements. He inferred from this that (in the preview condition) there was an overlap between the return movement of the hand and the planning of the next movement (which, in the no-preview case, would occur while the hand was at the centre). Leonard also noted that in the preview condition subjects reported an "uncanny lack of awareness of what they were doing" (p. 148).

Studies of Preview in Typing

With the exception of Jeeves (1961), who confirmed Leonard's observation using a rather similar choice reaction-time task, investigators looking at performance with preview have generally focussed on copytyping. When typists are deprived of the opportunity to see more than just the character they are typing, their typing is greatly slowed (e.g. Coover, 1923 [cited by

Salthouse, 1986]; Hershman & Hillix, 1965; Shaffer & Hardwick, 1970). The speedup due to preview has generally been interpreted as indicating that some mental operations overlap in time (see Salthouse, 1984b). However, another interpretation of the effect of preview is simply that it allows higher-level units (especially words) to be used for looking up responses (Shaffer & Hardwick, 1970; Thomas & Jones, 1970). According to this "chunking" account, preview speeds up performance because it enables the system to select responses corresponding to the typing of an entire word, and thereby reduces the number of responses to be selected.

There is a sizable preview effect for typing random letter strings, which suggests that some of the benefits of preview do not depend upon chunking (Salthouse, 1984b; Shaffer & Hardwick, 1970). It is nonetheless likely that chunking plays some additional role, because random letter strings are typed more slowly than are words (Fendrick, 1937; Salthouse, 1984a). Thus, the typing studies suggest that text preview facilities fast typing *both* by permitting chunking to occur and by allowing overlapping processing. It is not possible, however, to infer from the existing data which particular mental operations are overlapping in typing. For this purpose, the analysis of arbitrary choice reaction time tasks seems like the appropriate place to start.

Preview and the Response-selection Bottleneck

We begin with the hypothesis—derived from dual-task studies—that selection of responses constitutes a processing bottleneck (even when the stimulus-response mapping does not change), whereas recognition of a few familiar stimuli such as letters does not constitute a bottleneck. Consider the implications of this hypothesis for the sequence of stages in a serial-choice RT task performed with and without preview. Figure 2 (Lower Panel) shows the sequence of stages when there is zero preview (i.e. as soon as the subject responds to stimulus S_n, stimulus S_{n+1}. is then presented). The shaded stage represents the response selection stage (assumed to be a bottleneck on this view). Interresponse interval n (IRI_n) is defined as the time from response R_{n-1} to response R_n. RT_1 is defined as the time between S_1 and R_1. (RTs for subsequent stimuli could be defined correspondingly, but they will not figure prominently in the following discussion.)

It is obvious from the lower panel of Figure 2 that with zero preview, the bottleneck is irrelevant to performance. If the task is changed so that any of the stages are slowed by an amount t, this will produce an increase in all the IRIs of exactly t, whether the change affects the bottleneck stage or a stage before or after the bottleneck.

More interesting predictions arise with a preview of one (the top panel of Figure 2). In this condition, as soon as the subject has responded to S_n,

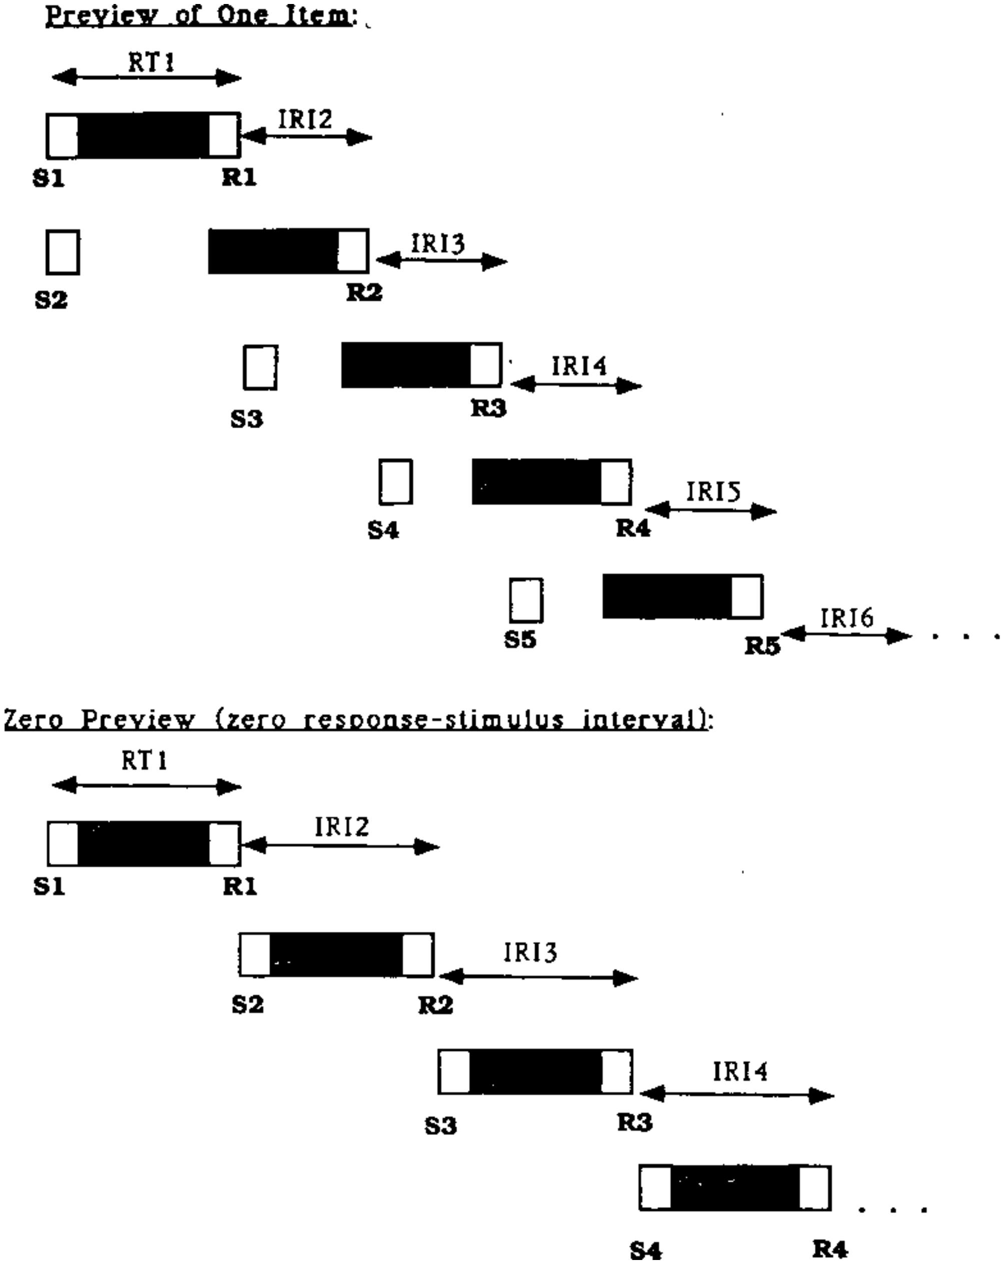

FIG. 2. The response-selection bottleneck model of serial performance. With preview of one (top panel), response selection is rate-limiting. With preview of zero (bottom panel), all stages are rate-limiting.

S_{n+2} is presented (except for the case of S_1 and S_2, which are presented simultaneously at the onset of the trial). As the figure makes plain, the model implies that certain mental operations will overlap, resulting in a reduction in IRIs compared to the no-preview condition. If the duration of the bottleneck stage is greater than or equal to the sum of the durations of the pre-bottleneck and post-bottleneck stages, then the effects of preview will saturate with a preview of one item. That is, previews of greater than one item will provide no further benefit in this case.[1]

If the response-selection bottleneck applies, then the sequence of stages should be as shown in the figure, with the selection of response R_n commencing as soon as the selection of response R_{n-1} is complete. Assuming that response selection takes no more time than the stages that precede and follow it, the duration required for the perception of stimulus *n* will not affect the latency for producing response *n*.

Now consider what happens when the duration of response selection is increased with preview (Figure 3). The slowing of response selection simply

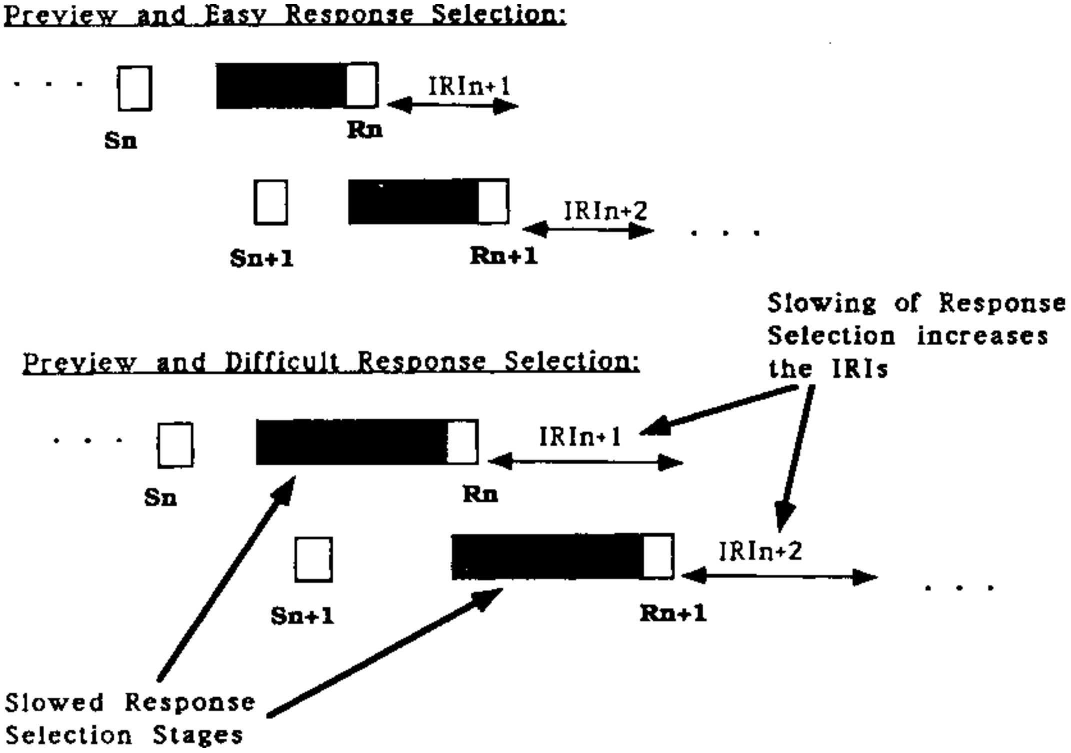

FIG. 3. Response-selection bottleneck model predicts that (with preview one) any increase in the duration of response selection will be reflected in each of the IRIs.

---

[1] The reader will find that the analysis of preview effects is straightforward when one assumes deterministic stage durations and a central bottleneck. When the duration of the bottleneck stage(s) is less than the duration of remaining stages, preview of greater than 2 is necessary for maximum benefit. In this case, when the preview is 2, oscillatory behaviour of successive IRIs is predicted. In general, when the duration of the bottleneck is greater than $1/n$ times the duration of the remaining stages, then IRIs will be the same with a preview of *n* and with a preview greater than *n*. The duration of the IRIs with this (asymptotic) preview is simply the duration of the bottleneck stage(s) (plus any possible time for the bottleneck mechanism to switch inputs).

increases the interresponse intervals accordingly. By contrast, consider what happens when the duration of perceptual processing is increased (Figure 4). As shown in the figure, this slowing should not affect the IRIs. However, note that this entire slowing should appear in the latency of the response to the *first* stimulus.

The predictions of the model can easily be summarized. Increases in the duration of the stage that constitutes a bottleneck (on the present hypothesis, response selection) should produce a constant effect on all the IRIs, whether preview is zero or one. In short, response selection should be *rate-limiting.* Assuming that response selection is always the most time-consuming stage, increases in the duration of stages that precede the bottleneck should produce corresponding increases in the duration of the first response, but subsequent IRIs should be unaffected.

Thus, the response-selection bottleneck model of serial performance with preview makes specific and testable predictions (closely related to tests previously employed with the PRP paradigm). But what of possible competing models? Do predictions described here discriminate between different types of models?

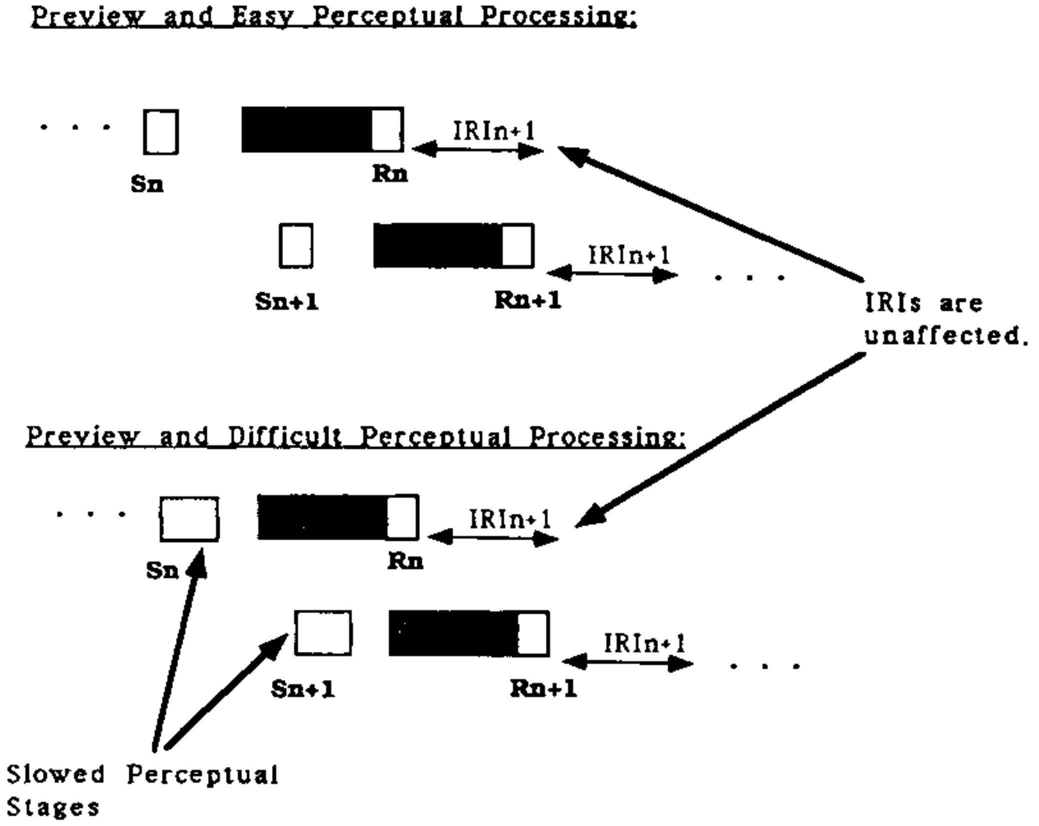

FIG. 4. Bottleneck model prediction for effect of increasing duration of stimulus recognition: the increase is absorbed by slack, leaving IRIs unaffected.

It is easiest to analyze competing models that postulate bottlenecks located in stages other than response selection. Although perceptual bottlenecks do not seem to arise in the PRP situation (at least when discriminations are easy; Pashler, 1989), they might nonetheless arise in serial performance tasks. Pashler (1984a) performed some single-task experiments requiring subjects to make a speeded choice response to a single item in a display of eight letters (cued by a bar-probe). In some trials, subjects had the opportunity to preview the entire display for 300 msec before the probe appeared. Despite this preview, responses were slowed when the letters were reduced in intensity or discriminability. This ruled out a "late-selection theory" analysis of this task, according to which all items are identified in parallel, and the cue simply directs a retrieval of the already-computed letter identities. That situation is related to the serial performance with preview presently under discussion. One interpretation of the results of Pashler (1984a) is that even if subjects can *identify* several characters in parallel (which it seems they can; cf. Duncan, 1980; Shiffrin & Gardner, 1972), they cannot maintain eight different letter identities available for retrieval by spatial location. Therefore, when the probe appears, the cued item must be identified "from scratch".

This limit is relevant in the present context because it *might* also prevent overlap of perception and other processes in the serial performance task, which also requires selection of identities based on location. This would result in a perceptual bottleneck. A perceptual bottleneck predicts that manipulations of the duration of stimulus recognition should have a constant (full) effect on IRIs, regardless of preview. Thus, this model should be empirically distinguishable from the response-selection bottleneck model.

A response *production* bottleneck is also a real possibility for serial performance. For example, response selection might drop out as a bottleneck due to the fact that the same mapping remains in place for all responses (a possibility raised above). At that point, limits on production of successive responses with the same effector system would presumably have to become rate-limiting. After all, the task inherently requires serial output of responses, and there must be *some* limits on the rate at which responses can be produced in the appropriate sequence. These rates might be so fast as to play no role whatever in the dual-task situation, but if the earlier bottleneck is eliminated in the serial task, then the production limit will emerge as rate-limiting. If the rate of response production limits the rate of serial performance, experiments of the sort discussed above should plainly be capable of revealing this, because with preview, manipulations of response selection should have much reduced effects on IRIs.

In the area of dual-task performance, models proposing graded sharing of mental capacity or resources have been widely discussed (e.g.

Kahneman, 1973). It is not immediately obvious how such models might apply in the case of serial performance. If processing stages are not queued up, but all operate in parallel, drawing on one or more common pools of resources, the effects of factors manipulating different stages should depend on many things, including the capacity allocation "policies" in force, and the functions relating speed of processing to capacity. Given these multiple factors, it is difficult to think of any principled reason why the pattern of results predicted by the model in Figure 2 should be observed. That is, even if some capacity allocation scheme could be discovered on a *post hoc* basis to mimic the predictions of the response-selection bottleneck, it would seem that such a result would be a remarkable coincidence if the particulars of capacity allocation determined performance.

In Experiment 1, we begin by examining the benefits of preview in a choice RT task with letter stimuli and manual keypress responses, comparing previews of zero, one, and the entire ensemble of stimuli. The results show that a preview of one provides virtually as much benefit as a preview of the entire stimulus ensemble, and the subsequent experiments therefore focus mostly on previews of zero and one. In Experiments 2 and 3, factors affecting perception and response selection are manipulated, and in Experiment 4, the duration of response production is varied by manipulating sequence length (one keypress versus two).

EXPERIMENT 1

The first experiment examined the effect of preview on choice RT performance. Subjects made a speeded response to each of 5 letters, pressing one of three keys, depending upon the identity of each letter. In the first three preview conditions, the letters were presented in a line from left to right. With preview zero, the subjects had no look-ahead: once they responded to a letter, the following letter was presented with as little delay as possible. With preview one, subjects had a look-ahead of one, so that when they responded to character *n,* character *n + 2* was presented. With complete preview, all five characters appeared at the outset.

A fourth preview condition was included, termed the preview one-successive condition. Here the subject had a look-ahead of one, but the letters appeared in succession in the same position on the display. The purpose was to determine whether any benefits of preview depend upon availability of the external display. [For half of the subjects (Version B) the second letter in the preview one-successive condition was actually delayed by 150 msec, which—it was thought—might improve recognition of the first letter.]

Method

Subjects. Sixteen undergraduates at the University of California, San Diego, participated as subjects in the experiment, in partial fulfillment of a course requirement. Eight participated in Version A and another eight participated in Version B.

Apparatus and Stimuli. The experiments were controlled by IBM PC microcomputers. Five letters were presented on NEC Multisync II monitors in each trial (controlled by a Paradise VGA card), and five responses were made on the IBM keyboard. Each letter was a T, a D, or an R, measuring 0.6 cm wide by 1.3 cm high (0.57° by 1.24° visual angle, based on a typical viewing distance of 60 cm). The letters were presented in white against a black background.

On each trial, letters were chosen at random, subject to the constraint that there could be no immediate repetitions of the same character (thereby avoiding any requirement for the subject to press the same key twice in succession).

In the first three preview conditions, the five letters were presented in a horizontal line. The width of this entire display was 6.5 cm (6.2°). In the fourth preview condition (preview one-successive), the fixation point and all five letters were presented in the centre of the screen.

Design. The experiment was divided into 24 blocks of 10 trials per block. The four different preview conditions (preview zero, preview one, complete preview, and preview one-successive) were presented in separate blocks. The 24 blocks cycled through six sets of the four different types of blocks, with subjects divided into four groups depending upon the preview condition with which they had begun.

Procedure. The subjects were given written instructions describing the tasks. The instructions stated that the subject should respond as quickly and accurately as possible to each character. They were discouraged from responding to the stimuli "in pairs or triplets", because pilot work showed that without this instruction some subjects fell into a strategy of producing responses in bunches (akin to response grouping; see Pashler & Johnston, 1989). Subjects responded to the letters *T, D*, and *R* by pressing the *V, B,* and *N* keys on the IBM keyboard, respectively, using the first three fingers of the dominant hand.

Each trial began with the presentation of a plus sign (fixation point) for 1 sec, located in the position where S_1 would occur. After a delay of 500 msec, the presentation of letters began.

In the preview zero condition, the first letter appeared in the leftmost position. As soon as the subject had responded to the first letter, the second character was presented in Position 2, and so on. (Thus, the response-stimulus interval was as close to zero as possible, given the limitations of generating the display, which took about 32 msec.) None of the letters disappeared until the previous response had been made.

In the preview one condition, the first letter appeared at the outset; after a delay of 100 msec, the second letter appeared in Position 2. Once the subject had responded to the first letter, the third letter appeared in Position 3, and so on.

In the complete preview condition, all five letters appeared at the start and disappeared only when the subject had responded to all the characters.

In the preview one-successive condition, the first letter was displayed for 100 msec. In Version A, the second letter was presented immediately at the offset of the first letter and remained present until the subject had responded to the first letter. Version B differed from Version A only in that there was a 150-msec delay between the offset of the first letter and the onset of the second letter. In either condition, S_2 disappeared as soon as R_1 had been made, and it was replaced by S_3 in the same position, and so on.

After the completion of a trial, there was a 3-sec intertrial interval before the beginning of the following trial. At the end of each block, feedback was provided consisting of mean IRIs and error rates for the block, and all preceding blocks.

Results

Figure 5 presents the subjects' mean response latencies as a function of preview condition and response (first through fifth response on a given trial). Strictly speaking, the first latency is RT_2, reflecting the time between the stimulus onset and the response, while the remaining times for response R_n $(n > 1)$ are the interresponse intervals (IRIs) measured from response R_{n-1} to response R_n. IRIs below 200 msec or greater than 1800 msec were discarded as deviant. The very small procedural difference between Versions A and B produced differences that were most apparent in the preview one-successive condition latency for RT_1: 1196 in Version A and 954 in Version B. This difference was to be expected, given that the two versions differed only in that S_1 was replaced by S_2 after only 100 msec in Version A, undoubtedly producing some backward masking of S_1. In the analyses that follow, the data from Versions A and B were combined.

Averaging across Responses 1-5, subjects were faster with preview one (494 msec) and complete preview (475 msec) than with preview zero (555 msec), but preview one-successive was the slowest of all (577 msec). These differences were significant, $F(3, 45) = 18.5$, $p < 0.001$, $MS_e = 10{,}115$. Subjects were also drastically slower on the first response ($RT_1 = 924$ msec) than on the remaining responses ($IRI_n = 391$ msec,

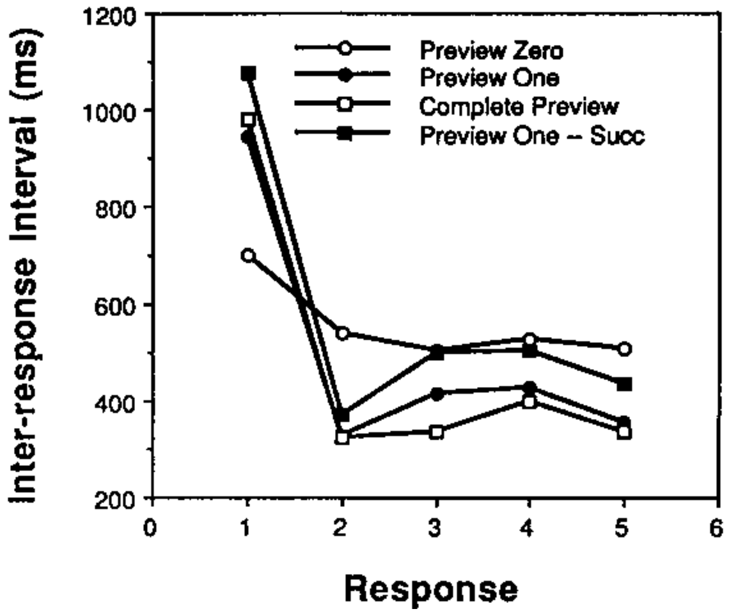

FIG. 5. Response latencies in Experiment 1 as a function of response and preview conditions. (Succ = successive).

439 msec, 464 msec, and 408 msec, for $n = 2, 3, 4$, and 5, respectively), $F(4, 60) = 240$, $p < 0.001$, $MS_e = 13{,}475$. As shown in the figure, preview slowed RT_1, while speeding up the remaining IRIs.

Error rates showed the same trends, although more weakly. There were fewer errors with preview one (4.1%) or complete preview (3.4%) than with preview zero (4.8%), and substantially more errors with preview one-successive (8.9%); the effect of preview was significant, $F(3, 45) = 14.3$, $p < 0.001$.

Interestingly, in all conditions there was a small speed-up from the fourth to the fifth (and final) IRI in a trial, with a corresponding increase in error rates, suggesting a speed-accuracy tradeoff.

Discussion

The experiment indicates that subjects' IRIs (after the first response) are substantially faster when they are able to preview stimuli. It seems to make little difference whether they preview just one stimulus, or all of the stimuli. However, preview slows the first response in the sequence very substantially.

One way to look at the data is in terms of the total time taken by subjects to respond to all five stimuli presented in a trial. The average total times are 2.77 sec for preview zero, 2.47 sec for preview one, 2.37 sec for complete preview. Total savings are modest (10-15%), but there are only five responses, and the proportion of savings would be expected to grow with number of responses (because no saving is expected on the first response). Of course, one should also note that although preview reduces the inter-response intervals and reduces total time, actual *response times* as conventionally defined from a stimulus to the corresponding response are lengthened by preview. (With preview one, response time *n* is simply the sum of IRI_{n-1} and IRI_n.)

The results for IRI_2 to IRI_n in the preview one condition and the complete preview conditions are essentially consistent with the model illustrated in Figure 2. This pattern of results is by no means critical evidence in favour of the model shown in Figure 2, however. It simply indicates that some processing overlap can be achieved with preview, and any scheme postulating overlap could fit these data.

The fact that the benefits of preview virtually saturate with a preview of one fits the expectations of a bottleneck model in which the duration of the bottleneck stage exceeds the sum of the durations of all stages *not* comprising a bottleneck (see footnote 1). As those predictions were generated by considering a hypothetical deterministic system, it is not surprising that a very small benefit for complete preview compared to a preview of one should appear in reality, where stage durations are all undoubtedly variable.

The response times on the *first* response clearly indicate effects not expected from the model shown in Figure 2, however. The model predicts that preview will have no effect on the latency of the first response. In fact, preview *slows* these responses. Why should this be? First, note that the first response is slower than any of the subsequent responses, even in the preview zero condition. One might suspect a "warm-up effect", whereby performing the task once optimizes preparation of the response mapping. Furthermore, the constraint against repetition of characters means that after the first, the stimuli become somewhat more predictable. With preview zero, however, there is additional slowing beyond that found with preview one. It would seem reasonable to suppose that there may be various sorts of processing "overhead" associated with taking full advantage of preview; this issue will be considered further in the General Discussion.

The effects of the preview one-successive condition are rather intriguing. Following the lead of Leonard (1953), it may be worth noting the subjective reports of individuals performing the tasks. Several individuals found the preview one and preview complete conditions easy and satisfying to perform, and after performing them, the preview zero condition (conventional serial RT) was experienced as irritating—as though stimulus delays were somehow creating an obstruction. By contrast, the preview one-successive condition was difficult to perform. Subjects reported little subjective confidence in their responses, although the actual increases in the IRIs and error rates in this condition were rather less than one might expect from such impressions. There are several features of this condition that could account for its relative difficulty, including the need to use internal buffer memory and the absence of spatial position as a cue to position in the sequence. These issues will not be investigated further here, but it is of some interest to find that stimuli must remain present in order for subjects to make effective use of preview in a serial RT task such as this one.

EXPERIMENT 2

The previous experiment confirmed the benefit of preview first noted by Cattell and Leonard, using a rather different sort of task than they used. The results also showed that a preview of one stimulus provided just about the full benefit available from preview of *all* the stimuli. For this reason, it was decided to compare previews of zero and one in the following studies. The remaining studies focused on testing predictions from the bottleneck model of serial performance with preview (shown in Figure 2). Experiment 2 required subjects to respond to characters with preview zero or one. There were three task-difficulty conditions. In the bright/easy-mapping condition, subjects saw four high-intensity white digits (*1, 2, 3, 4*) and made a compatible button push response, pressing one of four horizontally arrayed response keys. In the dim/easy-mapping condition, the same characters were presented at greatly reduced intensity, to increase the duration of perceptual processing. In the bright/hard-mapping condition, subjects saw the letters *Q, T, V,* and *L* (at high intensity) mapped onto the same row of keys. This relatively arbitrary mapping was chosen to slow the response selection stage (it seemed very unlikely that there would be much difference in the time required for recognition of the digits and letters; see Discussion of Experiment 2).

Method

Subjects. Twelve undergraduates at the University of California, San Diego, participated as subjects in the experiment, in partial fulfillment of a course requirement.

Apparatus and Stimuli. These were the same as in Experiment 1, except as noted. The same no-repetition constraint was employed as in Experiment 1. The digits were chosen from the set 1-4, and letters were chosen from the set *Q, T, V,* and *L.* Intensity was controlled by choosing white or dark grey (#63 or #4 from the VGA palette). This corresponded to about a 1.5 log unit difference (with characters in CRT displays, such measurements are necessarily very rough). Each trial consisted of 10, rather than 5 stimuli and responses per trial.[2] The display was 14 cm wide (13.1°).

Design. The experiment was divided into 24 blocks of 5 trials per block. There were six types of blocks (preview zero vs. preview one crossed with the three difficulty conditions: bright/easy-mapping, dim/easy-mapping, and bright/hard-mapping). The order in which six block types occurred was counterbalanced across subjects, as in Experiment 1.

Procedure. The procedure followed Experiment 1, except as noted. At the beginning of each block, the response mapping was presented (for 4 sec). In the preview one condition, S_2 was not delayed as in the first experiment. Subjects responded by pressing the *V*, *B, N,* or *M* keys, using the fingers of their dominant hand. In order to discourage errors and to avoid scoring ambiguities due to subjects' correcting responses or omitting responses from these long sequences, an entire trial was aborted and repeated whenever a subject made an error (new stimuli were selected for the repeat trial). Similarly, if any IRI fell outside the 200-1500-msec range, the trial was repeated. When this happened, an error message was displayed to the subject.

Results

Figure 6 presents the subjects' mean response latencies as a function of preview condition and response (first through tenth response on a given trial).

The top panel of Figure 6 compares the bright/easy-mapping and the dim/easy-mapping conditions. On the first trial, the intensity effect is slightly increased with preview zero (98 msec) compared to preview one (80 msec). However, on the second through ninth trials, the intensity effect is reduced to only 16 msec with preview one, compared to 49 msec with preview zero. This interaction of intensity with preview in IRI_n $(n > 1)$ is significant, $F(1, 11) = 38.8$, $p < 0.001$, $MS_\text{e} = 758$. All 12 subjects showed this interaction.

The bottom panel of Figure 6 compares the bright/easy-mapping and the bright/hard-mapping conditions. On the first trial, the mapping difficulty effect is smaller with preview zero (190 msec) than with preview one (251 msec). On the second through ninth trials, the mapping difficulty effect averages 155 msec with preview one and 130 msec with preview zero.

---

[2] In the previous experiment, the preview one-successive condition seemed overly difficult with ten stimuli, which was not true in the present tasks; therefore, the number of stimuli was increased.

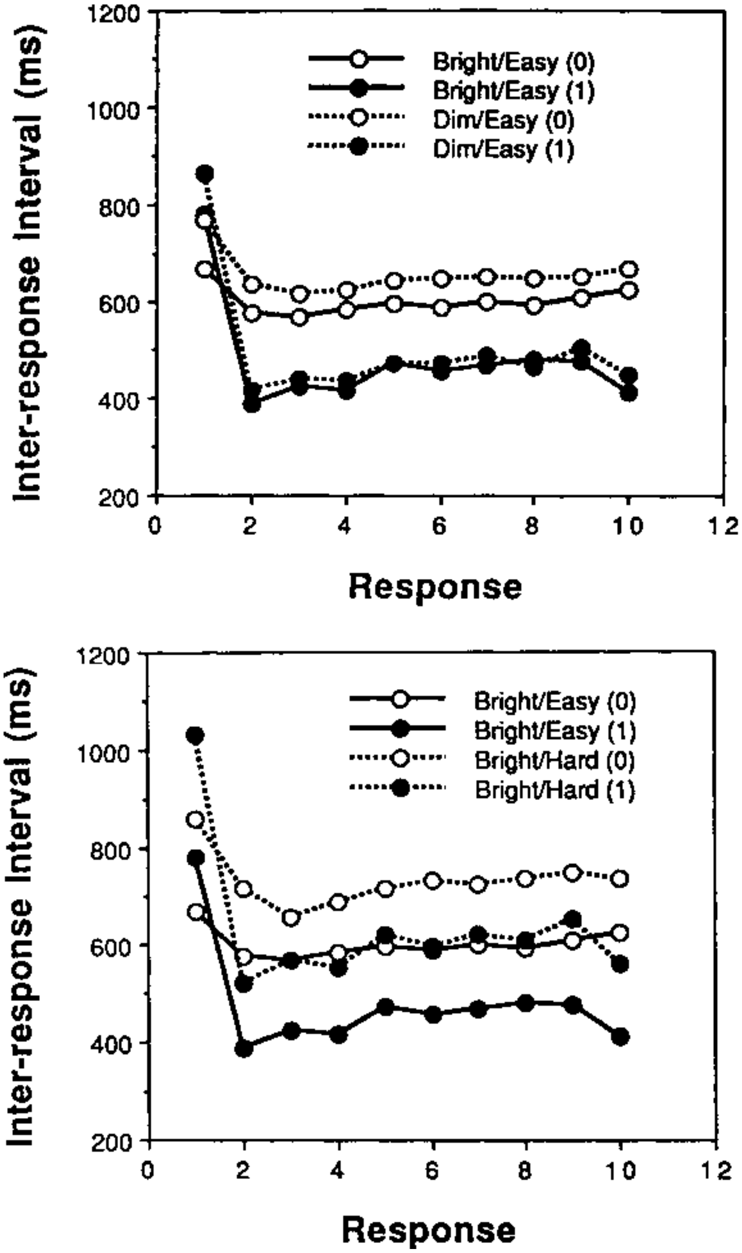

FIG. 6. Interresponse intervals in Experiment 2 as a function of response, preview, and difficulty. Top panel: effects of stimulus-response mapping difficulty; bottom panel: effects of stimulus intensity. (0) and (1) refer to preview.

This interaction of mapping difficulty with preview in IRI_n ($n > 1$) was not significant [$F(1, 11) = 3.6$, $0.10 > p > 0.05$, $MS_e = 3441$].

The total number of aborted trials per trial (trials on which an error or deviant RT occurred) was 0.11, 0.17, and 0.57 for the bright/easy, dim/easy, and bright/hard conditions, which was a significant effect, $F(2, 22) = 13.2$, $p < 0.001$, $MS_e = 0.111$. However, the effect of preview on the number of aborted trials was not significant.

Discussion

The results again show a beneficial effect of preview on the IRIs (with a slowing of the first response). More importantly, the results confirm several predictions of the response-selection bottleneck model of serial performance with preview (Figures 3 and 4). Preview reduced the effects on IRIs of a manipulation of the duration of perceptual processing (intensity). However, the first RT did not show such a reduction in intensity effects. This pattern is to be expected if response selection is more time-consuming than the other stages and constitutes a processing bottleneck (which together entail that response selection will be rate-limiting when there is a preview of one item). The fact that preview did *not* reduce the effects of a manipulation of the duration of response selection itself (mapping difficulty) further supports this hypothesis. (To be sure that the difference between the stimuli was not responsible for the mapping difficulty effect, an additional experiment was run, in which the digits 1-4 were mapped onto response keys in either a compatible or incompatible way, manipulated between blocks; the results were unchanged.)

EXPERIMENTS

In the previous experiment, preview was manipulated between blocks, so that subjects could anticipate whether preview would be provided. According to the model (Figure 2), essentially the same results should be obtained when preview is manipulated in mixed trials, so that subjects cannot anticipate whether preview will be provided. It is not uncommon, however, for effects of variables to differ as a function of whether subjects can and cannot anticipate what will happen on the next trial, and any such effects would require explanation.

Method

Twelve new subjects ran in an experiment that was identical to Experiment 2, except as noted. In the easy-mapping condition, the digits *1*, *2*, *3,* and *4* were mapped onto the four response keys used previously; in the difficult-mapping condition, the letters *Q, T, V,* and *L* were mapped onto the corresponding keys. Within each block, half the trials were preview zero, and half were preview one.

Results and Discussion

Figure 7 presents the subjects' mean time between responses as a function of preview condition and response (first through tenth response on a given trial). The results were virtually identical to the easy and hard mapping conditions in Experiments 2 and 3, and, as before, the mapping difficulty,

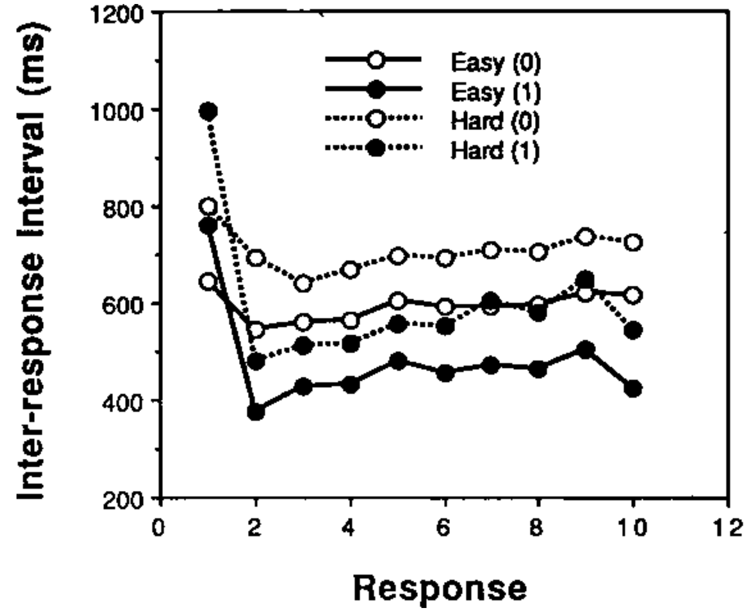

FIG. 7. Interresponse intervals in Experiment 3 as a function of response, preview (0 vs. 1), and stimulus-response mapping (easy vs. hard).

preview, and response (1-10) variables were significant. In responses 2-10 the effect of mapping difficulty averaged 107 msec with preview zero and 105 msec with preview one, $F(1, 11) = 0.06, p > 0.80, MS_e = 2241$.

The number of aborted trials was greater for the more difficult mapping (0.39 aborted trials per trial) than for the easy mapping (0.16), which was significant, $F(1, 11) = 29.1, p < 0.001$. There was no significant effect of preview on this measure, however.

The results again indicate that slowing response selection increases the times between responses, and the degree of slowing does not appear to differ depending upon whether or not preview is provided. This supports the model shown in Figure 2.

EXPERIMENT 4

Experiment 4 examined the effect of increasing the duration of response production, by varying the number of keypress responses required (sequence length). Both preview (zero, one, or two) and sequence length (one keypress or two) were varied between blocks, The basic goal was to test the assumption (implied by Figure 2) that response production processes at the "back end" of each task are exempt from bottleneck interference, just as perceptual processes near the "front end" are.

The effect of the length of response sequences in a single task has been extensively examined (e.g. Henry & Rogers, 1960; Klapp et al., 1979).

These studies have generally focused on how sequence length affects the time at which a response sequence *begins,* rather than ends. It is generally believed that greater sequence lengths are associated with longer RTs because a motor programming stage takes more time for longer sequences. Such effects are generally much greater in choice RT than in simple RT tasks. Klapp et al. argued that this is because the motor programming can be accomplished in advance in simple RT. In their view, motor programming occurs before the stimulus in simple RT, and between the stimulus and the initiation of the response in choice RT.

On the other hand, Sternberg, Monsell, and their collaborators (Sternberg, Monsell, Knoll, & Wright, 1978) have observed small but very systematic effects of sequence length on both the time to begin the sequence *and* the duration of the sequence using simple RT tasks of many different types (both vocal and manual). Monsell (1986) developed a model of the programming of motor sequences, which postulates that the production of response sequences includes a cyclic motor programming process whereby the sub-program controlling each element in the sequence is retrieved and then the appropriate motor commands are issued. This cyclic process continues until production of all the elements in the sequence has been commanded. The retrieval is assumed to take longer when there are more elements, whereas the command process may take longer depending on the composition of the subprogram—for example, when the sequence involves alternating hands.

Thus, on Monsell's view, one component of motor programming is a retrieval process that is time-locked to each of the individual elements in the sequence (through the final one). If this retrieval operation occupies the central bottleneck continually until the last commands have been issued, then increases in sequence length should obviously be rate-limiting for serial performance with preview.

Method

Twelve new subjects ran in an experiment that was identical to Experiment 2, except as noted. The letters *Q, T,* and *V* were mapped onto the *V, B,* and *N* keys, respectively. We examined previews of zero, one, and two (manipulated between blocks). Half of the blocks were in the one tap condition; here subjects were required to press the appropriate key once. The other half of the blocks were in the two tap condition, and here subjects pressed the same key twice. The order of presentation of the six types of blocks was counterbalanced as in the earlier experiments. In the two tap condition, stimuli and responses were all timed from the *second* tap. A trial consisted of six stimuli and responses and was repeated when an error occurred, or the IRI for the final keypress in a given response lay outside the range 80-1800 msec (the briefer lower cutoff was employed because we did not wish to exclude any extremely brief IRIs induced by the dual-tap requirement; in fact, such IRIs were rare).

Results

Figure 8 presents the subjects' mean response latencies as a function of preview condition, number of taps (one tap vs. two tap), and response (first through sixth response on a given trial). The top panel shows the latencies for the first keypresses, and the bottom panel shows latencies for the second keypresses. Note that in both panels IRI_2 to IRI_6 are timed from the *second* keypress in the two tap condition, which also triggered the stimulus presentation.

The number of aborted trials was not significantly affected by any of the experimental variables.

Consider now the latencies for IRI_2 to IRI_6 (until further notice, all times discussed are means of these IRIs). With no preview, the first of two keypresses in each response is produced after 558 msec in the one tap condition, and after 523 msec in the two tap condition. This is a reverse of the effect of sequence length commonly (but not invariably) observed in single-task reaction times (Klapp et al., 1979).

In the two tap condition without preview, the second tap occurs at 707 msec, 184 msec after the first tap. With a preview of one, IRI for the first keypress is 417 msec in the one tap condition (preview benefit of 141 msec). But in the two tap condition with a preview of one, the IRI for the first keypress is 282 msec. If the bottleneck lasted up to a point that was synchronized with the final keypress, there would be no reason for the IRIs for the first keypress to be any different in the one tap condition (417 msec) and the two tap condition (282 msec).

The interval between the first and second taps averages 184 msec with preview zero, 191 msec with preview one, and 195 msec with preview two. Thus, the relative timing of the two keypresses is little affected by preview.

The overall rate at which the task is being performed *is* slowed by the requirement to make two taps. With zero preview, the IRIs to complete the entire response are 558 msec in the one tap condition and 707 msec in the two tap condition (a difference of 149 msec). With preview one and preview two, the differences are 56 and 70 msec, respectively.

In summary, with or without preview, subjects perform the two tap task more slowly than the one tap task. With preview, the time between completions of responses in the two tap condition is in the order of 55-70 msec longer than in the one tap condition. However, this is substantially less than the 149-msec increment that the two tap condition adds to the total response time with zero preview, and it is much less than the time between keypress responses in sequences of two taps.

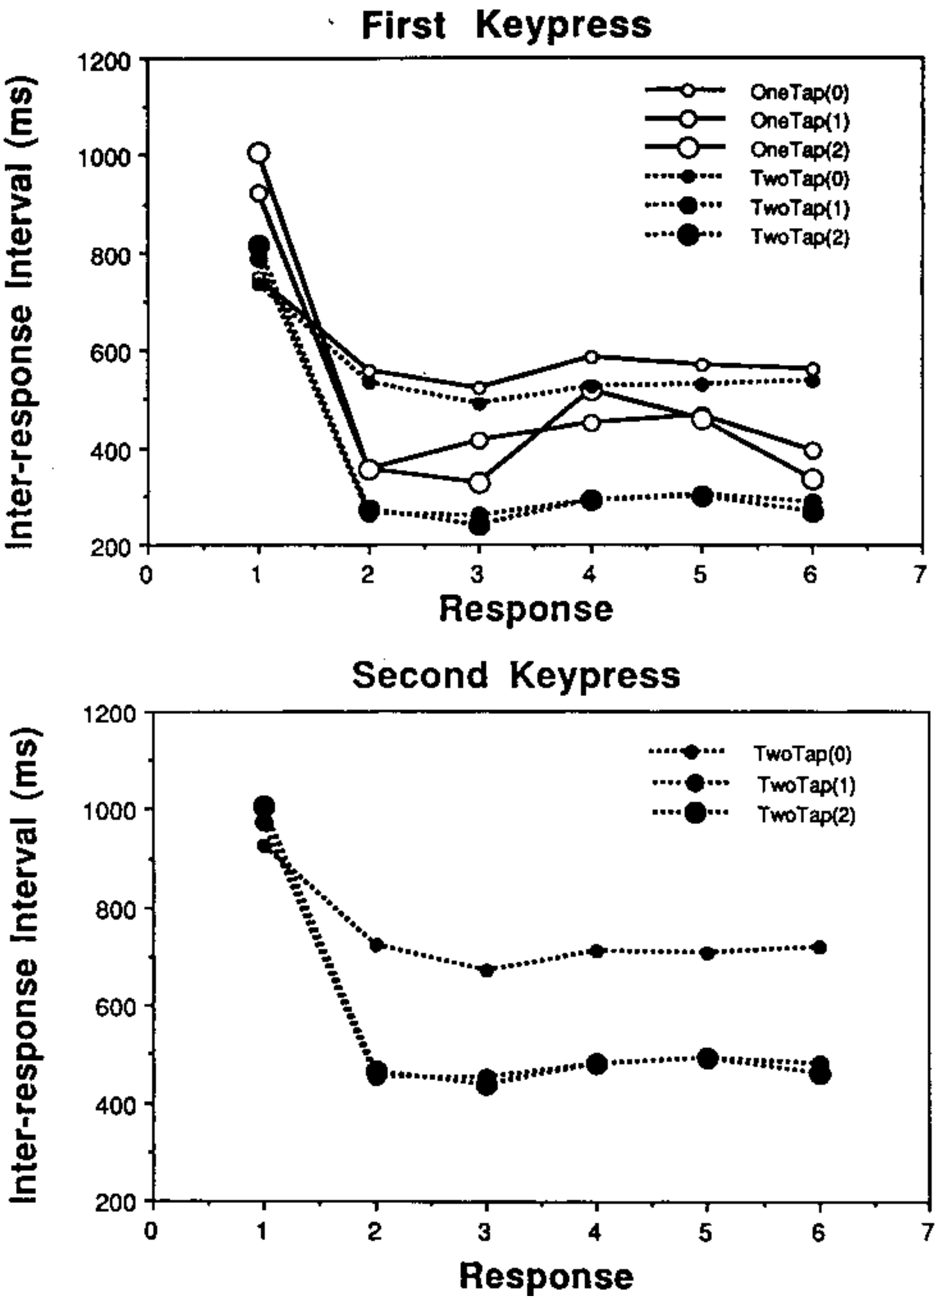

FIG. 8. Intel-response intervals in Experiment 4 as a function of response, preview (0, 1, and 2), and number of keypress responses required to each stimulus (one tap vs. two taps). Top panel: first keypress; bottom panel: second keypress.

Discussion

First of all, differences between preview one and preview two are quite minimal. As noted earlier, this supports a model in which the duration of the stages constituting a bottleneck generally exceeds the duration of stages not constituting a bottleneck (see footnote 1).

The second main result is that although requiring two taps increases the duration of the response production (defined to include triggering of both taps when two were required), this increase is by no means fully reflected in the rate at which the task can be performed with preview. In short, preview reduces the effects of response duration on rate of performance, as it reduced the effect of stimulus intensity (although the degree of reduction appears proportionately somewhat smaller).

We can conclude, therefore, that the bottleneck processes do not extend up to the moment the final keypress is actually produced. In fact, the implications are somewhat broader than that, but they must be rather carefully stated. The data rule out any model according to which the rate-limiting bottleneck extends through any stage that is time-locked to individual movements up to and including the final keypress movement. The concept of time-locking is here defined to apply to any two events that occur on every trial with a fixed separation in time. To consider a concrete example, take the suggestion that the bottleneck includes certain motor programming processes that trigger each individual keypress movement, so that this triggering precedes the movement by a fixed amount *t*. The subprogram retrieval operation model postulated by Monsell (1986) is an example of such a time-locked motor programming process. Any such hypothesis is ruled out by the data, even though they do not postulate that the bottleneck actually lasts up to the moment at which the final keypress actually takes place. The latter possibility is obviously too narrow, as it seems so likely that there must be *some* peripheral motor processes for which central processes do not wait. Obviously, the conclusion that motor programming processes time-locked to production of elements in the sequence cannot occupy the bottleneck does not argue against Monsell's model of sequence production; it only argues that if the retrieval substages of the type he postulates do occur, they do not occupy the central bottleneck.

The third result is that requiring two keypresses rather than one *does* decrease the rate at which the task is completed. It would appear that the bottleneck must include *some* processes whose duration is increased when two responses are to be produced rather than one. One possibility would be that (a) the bottleneck includes both response selection and a motor programming stage; (b) this motor programming stage takes longer when two keypresses are to be executed rather than one; and (c) increasing the number of keypresses lengthens the duration of motor programming more than it affects the latency of the first keypress. If this is the case, then the critical motor programming stages might (although it does not follow that they would have to) be operating *after* overt motor behaviour has begun. (They could not, however, be time-locked to the production of the response elements, like the subprogram retrievals postulated in Monsell's model.) Given the speed emphasis in the present task, this does not seem so implausible. (At a more molar level, people can obviously begin an activity before they have finished planning all of what they will do.) Both Rosenbaum (1980) and Goodman and Kelso (1980) have made quite similar suggestions based on very different kinds of evidence.

GENERAL DISCUSSION

Subjects performed serial choice reaction time tasks in which the difficulty of stimulus recognition, response selection, and response production were manipulated. As noted by Cattell (1886) and Leonard (1953), preview sped up the overall rate of responding. The model shown in Figure 2—which postulates a response-selection bottleneck as the rate-limiting factor governing the rate of performance with preview—is generally supported by the data. Some of the data suggest that the bottleneck may also encompass a motor programming stage that takes somewhat longer when a sequence involves more movements. (Recent PRP studies also seem to suggest that the bottleneck may encompass motor programming; Pashler & Christian, 1991.)

A preview of one item increases the rate of responding, whereas additional preview appears to have little effect. Slowing of response selection slows the rate of responding by essentially the same amount, regardless whether preview is provided. Slowing of stimulus recognition delays the first response, but its effect on the rate of subsequent responding is much reduced by stimulus preview. Finally, requiring two identical keypress responses to each stimulus also has a modest effect upon the rate of responding when preview is provided, and the effect is less than the extra time required actually to produce two keypresses rather than one.

The model shown in Figure 2, which asserts that selection of one response cannot overlap with selection of another response but allows stimulus recognition and response production to overlap with response selection, predicts this general pattern of effects. Nonetheless, the data included some unexpected findings suggesting that Figure 2 does not quite tell the whole story.

Unexpected Findings

The main unexpected results pertained to the latency of the *first* response in a trial. Responses on the first trial were always slower than subsequent responses. This effect was greatly magnified with preview. Presumably, when there is no preview, the first response is slower than the later responses because subjects become better prepared for the task as they do it. It is difficult to see why there would be more temporal uncertainty about these later stimuli (there may be less), so one must assume that what improves is not preparation for the time of arrival but preparation of the "mapping rules" for the task.

The next question is why providing preview should slow down performance on the first trial still further? Some slowing of RT_1 is commonly observed in dual-task experiments, but the magnitude of that slowing is usually less than the effect under discussion. It would seem plausible that scheduling the processing stages shown in Figure 2 may involve executing a plan that takes a certain amount of time to set in motion when S_1 appears. The plan must involve various activities that have not been discussed. For example, the subject must switch the focus of visual attention from one stimulus to the next at the appropriate time. The only condition in which such shifts are probably *not* necessary is in the preview one-successive condition in Experiment 1, and performance there was quite poor. Recent dual-task studies suggest that visual attentional shifts can overlap with response selection (Pashler, 1991). Almost certainly, the plan must also involve eye movement control. At least in Experiments 2-4, displays occupied enough space that it would not be possible to do the task without making eye movements. Certain eye movements can be generated independently of the central response-selection bottleneck (Pashler et al., 1993), but it is not clear whether the sort required in these tasks would qualify. It might be interesting to record eye movements while subjects performed serial tasks with and without preview, looking at the relative timing of eye movements and the serial task responses. For present purposes, however, the point is just that there is a substantial amount of "overhead" associated with coordinating the sequence of mental operations shown in Figure 2. It seems plausible that the unanticipated slowing of RT_1 may be attributed to the process of setting in motion the processes required to coordinate these various activities.

Task Preparation

It was suggested in the introduction that an analysis of overlapping mental operations in the serial performance task was of interest not only because serial performance is ubiquitous and little studied, but also because it might illuminate the role played by task set in dual-task interference. In these experiments, there was only a single mapping, and yet response selection stubbornly resisted temporal overlap. These results imply that the response-selection bottleneck in the PRP situation cannot be caused by the need to switch task sets between Task 1 and Task 2. This by no means implies that switching of task sets is *never* required in dual-task experiments, not does it imply that preparatory factors do not play an important role in dual-task interference.

In fact, it seems likely that task set switching *is* necessary only when two successive tasks map the *same* stimuli onto *different* responses. Allport and Styles (1991) called attention to (and replicated) some fascinating early observations of Jersild (1927). Jersild required his subjects either to perform the same operation over and over on a list of stimuli, or to alternate between two different operations. He found a very large slowing (often well over 0.5 sec per response) in the alternating condition. This effect is virtually eliminated when the type of stimulus cues the appropriate task—e.g. if the subject is to name each word in a list and add three to each digit (Spector & Biederman, 1976).

It was noted in the introduction that graded failure of preparation undoubtedly contributes to dual-task slowing. This is evident in the fact that responses to both tasks in the PRP situation are generally slower than single-task controls, even when the SOA is very long. As Welford (1980) noted, responses to R_2 are often slowed when S_2 *follows* R_1.

A straightforward way of accounting for these various effects, together with the current observations, is as follows. Preparation for a task involves "activating" a set of response mapping rules in advance of a trial. These rules govern the response selection stage. The more rules are active, the longer the response selection stage will take. Activation of rules is assumed to decay over a period of seconds, rather than milliseconds. Suppressing or refreshing such rules is assumed to be a relatively slow process as well. (Much of what passes for dual task interference in various studies—such as those involving the effects of concurrent memory loads—may actually reflect disruption of this preparatory activation process, resulting in a slowing of response selection, but no delay in its onset, as suggested by Logan, 1978.) As Gottsdanker (1980) suggested, the effect of the number of alternatives described by the Hick/Hyman Law may also be attributable to such an effect.

However, the response-selection bottleneck (and its consequences—the PRP effect and the maximal rate of serial performance) are *not* attributable to preparatory limitations, in this sense. Rather, they seem to stem from the fact that when response selection (and perhaps motor programming) actually commences, other response selection operations are forced to wait for its completion. As demonstrated in the serial performance studies reported above, this waiting occurs even when just a single stimulus-response mapping rule is activated, so that preparatory shifts are unnecessary.

According to the view advanced here, in the PRP situation the mappings constituting *both* tasks are generally prepared in advance.[3] As the mapping rules for both tasks are activated, either task will be performed more slowly than in a single-task control condition, even when S_2 follows R_1. If, as seems sensible, we reserve the term "task switching" for the mental opera-

---

[3] There is an unfortunate complication that must be noted. With two tasks both involving manual responses on different hands, there is evidence that one aspect of response selection in the tasks (namely the mapping from spatial response codes to finger movements; *may* be switched between tasks; Pashler, 1990). What follows seems to apply most straightforwardly to the case of responses in different modalities.

tion of suppressing one set of mapping rules and activating another, then task switching does not generally occur either in the PRP situation or in serial performance. Task switching is only necessary when the tasks are constructed in such a way that the mapping rule for one task will actually "accept" the wrong stimuli as inputs and produce an error (as in Jersild's alternating conditions).

Roughly speaking, the cost of having to prepare more response rules is generally in the range of from tens to a few hundred milliseconds, and are reflected in the Hick/Hyman Law effect and the difference between single task performance and long-SOA PRP performance. The response-selection bottleneck introduces delays of 0.25-0.5 sec in choice reaction time tasks, and these delays occur whenever two independent responses must be selected (regardless of whether there are one or many mapping rules active). Finally, task switching often produces costs exceeding 0.5 sec or even a full second, but this occurs only when the two tasks are structured so that having all the mapping rules active would lead to outright errors (as in Jersild's work).

This proposal shares some common ground with the suggestions put forward by Logan (1978) and Gottsdanker (1980), although it differs somewhat from either of their formulations, particularly with respect to the role of bottlenecks. More work will be necessary to test further and to elaborate these ideas. Another challenge will be to explore their application to the vast range of human behaviour beyond the realm of the rather constrained and artificial tasks examined here. The chronometric approach utilized here could be readily applied to the analysis of the timing of mental events in serial tasks that are more cognitively interesting than choice reaction time tasks, and also to many activities that are commonly performed outside the laboratory.

Generalizing Beyond Strict Stage Models

One final issue concerns another question about the generality of the analyses presented here. Starting with the first four figures, the concept of a bottleneck has been developed in the context of strictly successive stages (as in the work of Sternberg, 1969). The assumption that there are fairly discrete transitions between successive stages (particularly stimulus analysis and response selection) in fact enjoys empirical support (Miller, 1988). However, the tests of competing hypotheses carried out here do not rely on this assumption. Suppose, for example, there was some "cascading" of processing—i.e. that each stage ordinarily fed information into the next stage continuously, with different stages usually operating at the same time. If stage n was postponed by a processing bottleneck, this could allow the output of stage n - 1 to reach an asymptotic strength level that would be the same regardless of the level of difficulty of factors affecting stage $n - 1$. In that way the cascade model could yield the same types of predictions as the discrete version depicted in Figures 1-4. Another assumption that seems unrealistic is that the output of the pre-bottleneck processing simply waits passively for the bottleneck mechanism and then accesses it instantaneously. If the output state were subject to a (comparatively slow) decay, and if switching were taking some time, the basic character of the predictions explored in this article would not change.

REFERENCES

Allport, A., & Styles, E.A. (1991). Multiple executive functions, multiple resources? Experiments in shifting attentional control of tasks. Unpublished manuscript, Oxford University.

Bertelson, P. (1966). Central intermittency twenty years later. *Quarterly Journal of Experimental Psychology, 18,* 153-163.

Brebner, J. (1977). The search for exceptions to the psychological refractory period. In S. Dornic (Ed.), *Attention and performance, VI* (pp. 63-78). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

Cattell, J.M. (1886). The time it takes to see and name objects. *Mind, 11*, 63-65.

Coover, J.E. (1923). A method of teaching typewriting based upon a psychological analysis of expert typewriting. *National Education Association: Addresses and Proceedings, 61*, 561-567.

Duncan, J. (1980). The locus of interference in the perception of simultaneous stimuli. *Psychological Review, 87,* 272-300.

Fagot, C., & Pashler, H. (1992). Making two responses to the same object: Exploring the central bottleneck. *Journal of Experimental Psychology: Human Perception and Performance, 18,* 1058-1079.

Fendrick, P. (1937). Hierarchical skills in typewriting. *Journal of Educational Psychology, 28,* 609-620.

Goodman, D., & Kelso, J.A.S. (1980). Are movements prepared in parts? Not under compatible (naturalized) conditions. *Journal of Experimental Psychology: Human Perception and Performance, 109,* 475-495.

Gottsdanker, R. (1980). The ubiquitous role of preparation. In G.E. Stelmach & J. Requin (Eds.), *Tutorials in motor behavior* (pp. 355-371). Amsterdam: North Holland Press.

Gottsdanker, R., & Stelmach, G.E. (1971). The persistence of psychological refractoriness. *Journal of Motor Behaviour, 3,* 301-312.

Henry, F.M., & Rogers, D.E. (1960). Increased response latency for complicated movements and a "memory drum" theory of neuromotor reaction. *Research Quarterly, 31,* 448-458.

Hershman, R.L., & Hillix, W. A. (1965). Data processing in typing: Typing rate as a function of kind of material and amount exposed. *Human Factors, 7,* 483-492.

Jeeves, M.A. (1961). Changes in performance at a serial-reaction task under conditions of advance and delay of information. *Ergonomics, 4,* 329-338.

Jersild, A.T. (1927). Mental set and shift. *Archives of Psychology.* Whole No. 89.

Kahneman, D. (1973). *Attention and effort.* New York: Prentice Hall.

Klapp, S., Abbott, J., Coffman, K., Greim, D., Snider, R., & Young, F. (1979). Simple and choice reaction time methods in the study of motor programming. *Journal of Motor Behavior, 11,* 91-101.

Leonard, J.A. (1953). Advance information in sensori-motor skills. *Quarterly Journal of Experimental Psychology, 5,* 141-149.

Logan, G.D. (1978). Attention in character-classification tasks: Evidence for the automaticity of component stages. *Journal of Experimental Psychology: General, 107,* 32-63.

McCann, R., & Johnston, J.C. (1992). The locus of the single-channel bottleneck. *Journal of Experimental Psychology: Human Perception and Performance, 18,* 471-485.

Miller, J.O. (1988). Discrete and continuous models of human information processing: Theoretical distinctions and empirical results. *Acta Psychologica, 67,* 191-257.

Monsell, S. (1986). Programming of complex sequences: Evidence from the timing of rapid speech and other productions. In H. Heuer & C. Fromm (Eds.), *Generation and modulation of action patterns.* Berlin: Springer-Verlag.

Pashler, H. (1984a). Evidence against late selection: Stimulus quality effects in previewed displays. *Journal of Experimental Psychology, 10,* 429-448.

Pashler, H. (1984b). Processing stages in overlapping tasks: Evidence for a central bottleneck. *Journal of Experimental Psychology: Human Perception and Performance, 10,* 358-377.

Pashler, H. (1989). Dissociations and dependencies between speed and accuracy: Evidence for a two component theory of divided attention in simple tasks. *Cognitive Psychology, 21,* 469-514.

Pashler, H. (1990). Do response modality effects support multiprocessor models of divided attention? *Journal of Experimental Psychology: Human Perception and Performance, 16,* 826-842.

Pashler, H. (1991). Shifting visual attention and selecting motor responses: Distinct attentional mechanisms. *Journal of Experimental Psychology: Human Perception and Performance, 17,* 1023-1040.

Pashler, H. (1993). Dual-task interference and elementary mental mechanisms. In D. Meyer & S. Kornblum (Eds.), *Attention and performance, XIV (pp.* 245-264). Cambridge, MA: MIT Press.

Pashler, H., Carrier, M., & Hoffman, J.E. (1993). Saccadic eye movements and dual-task interference. *Quarterly Journal of Experimental Psychology, 46A,* 51-82.

Pashler, H., & Christian, C. (1991). Dual-task interference and the production of motor responses. Manuscript submitted for publication.

Pashler, H., & Johnston, J.C. (1989). Chronometric evidence for central postponement in temporally overlapping tasks. *Quarterly Journal of Experimental Psychology, 41A,* 19-45.

Rosenbaum, D. (1980). Human movement initiation: Specification of arm, direction, and extent. *Journal of Experimental Psychology: Human Perception and Performance, 109,* 444-474.

Salthouse, T. A. (1984a). Effects of age and skill in typing. *Journal of Experimental Psychology: General, 113,* 345-371.

Salthouse, T.A. (1984b). The skill of typing. *Scientific American, 250,* 128-135.

Salthouse, T.A, (1986). Perceptual, cognitive and motoric aspects of transcription typing. *Psychological Bulletin, 99,* 301-319.

Schweickert, R. (1978). A critical path generalization of the additive factor method: Analysis of a Stroop task. *Journal of Mathematical Psychology, 18,* 105-139.

Shaffer, L.H., & Hardwick, J. (1970). The basis of transcription skill. *Journal of Experimental Psychology, 84,* 424-440.

Shiffrin, R.M., & Gardner, G.T. (1972). Visual processing capacity and attentional control. *Journal of Experimental Psychology, 93,* 72-83.

Smith, M.C. (1967). Theories of the psychological refractory period. *Psychological Bulletin, 67,* 202-213.

Smith, M.C. (1969). The effect of varying information on the psychological refractory period. In W.G. Koster (Ed.), *Acta Psychologica, 30,* 220-231.

Spector, A., & Biederman, I. (1976). Mental set and mental shift revisited. *American Journal of Psychology, 89,* 669-679.

Sternberg, S. (1969). The discovery of processing stages: Extensions of Donders' method. In W.G. Koster (Ed.), *Attention and performance, II* (pp. 276-315). Amsterdam: North Holland.

Sternberg, S., Monsell, S. Knoll, R., & Wright, C.E. (1978). The latency and duration of rapid movement sequences: Comparisons of speech and typewriting. In G. Stelmach (Ed.), *Information processing in motor control and learning* (pp. 117-152). New York: Academic Press.

Telford, C.W. (1931). The refractory phase of voluntary and associative responses. *Journal of Experimental Psychology, 14,* 1-36.

Thomas, E.A.C., & Jones, R.G. (1970). A model for subjective grouping in typewriting. *Quarterly Journal of Experimental Psychology,* 22, 353-367.

Welford, A.T. (1952). The 'Psychological Refractory Period' and the timing of high-speed performance—A review and a theory. *British Journal of Psychology, 43,* 2-19.

Welford, A.T. (1967). Single-channel operation in the brain. *Acta Psychologica, 27,* 5-22.

Welford, A.T. (1980). The single-channel hypothesis. In A.T. Welford (Ed.), *Reaction time* (pp. 215-252). New York: Academic Press.

Wilkinson, R.T. (1990). Response stimulus interval in choice serial reaction time—Interactions with sleep deprivation, choice and practice. *Quarterly Journal of Experimental Psychology, 42,* 401-423.

Manuscript received 12 July 1992