Original →Copy Before Using

Psychological Bulletin
1984, Vol. 96, No. 1, 29–44

Copyright 1984 by the American Psychological Association, Inc.

On the Cost and Benefit of Cost and Benefit

John Jonides
University of Michigan

Robert Mack
IBM Thomas J. Watson Research Center

Cost–benefit analysis of reaction times has become a popular chronometric tool in the study of cognitive processes. We review the technique, assumptions underlying its application, and pitfalls that are encountered in actually implementing it in various experimental contexts. This review suggests that the unthoughtful application of the technique may cause one to draw improper conclusions about the underlying mechanisms that produce costs and benefits.

The keys to the mind are few in number. Because of this, the small number of paradigms available to experimental cognitive psychology are exploited ubiquitously to reveal characteristics of a variety of mental activities. When a new technique is invented, it is eagerly coopted by large numbers of investigators who hope to unravel yet more of the mysteries of cognitive life.

This is as it should be, given the limited repertoire of methodological resources. However, care must be taken to ensure that new empirical tools are examined critically and that their flaws are laid bare before they are put to use. Often this is not feasible until a technique has been used for some time, thereby permitting sufficient examples of its product to be scrutinized. Also, some period of fairly extensive application of an experimental procedure is required before investigators begin to amass occasionally discomforting feelings about its use. These feelings typically arise during day-to-day activities in the laboratory when otherwise mundane decisions about the details of an experimental design lead one to question some basic assumptions about the empirical technique that is being applied.

We have had such feelings of uncertainty. The paradigm that produces our symptoms is known as the cost–benefit analysis of reaction times, originally developed by Posner and Snyder (1975), who in turn modified procedures developed by others (e.g., Beller, 1971; Meyer & Schvaneveldt, 1971; Meyer, Schvaneveldt, & Ruddy, 1975). This paradigm has been used extensively in recent years to provide supposedly informative decompositions of reaction times (e.g., Fischler & Bloom, 1979; Neely, 1976, 1977; Stanovich & West, 1979, 1981). As such, the paradigm has followed on the heels of other recent developments in the use of reaction time as a measure of cognitive processes (e.g., Sternberg, 1969).

In this article, we review cost–benefit analysis, outline some issues that arise in its use, and provide examples of research that highlight these issues. In addition, we review certain suggestions that may offer a more principled basis for the application of cost–benefit methodology. First, let us review the basic methodology and motivation underlying the cost–benefit technique.

Throughout the history of psychology, there has been strong interest in the effects of past experience on cognitive processes. During the past 20 years or so, one manifestation of this interest has been the study of preparation effects on performance in various tasks. Preparation effects refer to various performance phenomena that all presumably depend upon the state of a subject's preparation upon engaging in a task. The study of preparation phenomena with techniques such as cost–benefit analysis promises to reveal some of the mechanisms that underlie the influence of past experience on ongoing processes.

Let us examine an early example of these experiments, a study by Leonard (1958).

Preparation of this manuscript was supported by National Science Foundation Grant BNS 80-24512 and Air Force Office of Scientific Research Grant 82-0297 to John Jonides. The authors gratefully acknowledge comments by Howard Egeth, David Meyer, and James Neely on an earlier version of this article.

Requests for reprints should be sent to John Jonides, Department of Psychology, University of Michigan, 330 Packard Road, Ann Arbor, Michigan 48104.

Leonard's subject responded with a key press in a one-to-one mapping task requiring a speeded response. The stimulus for a response on any trial was the offset of one of six lights. In one condition, the subject was given a selective warning about whether the critical light would be located in a left or right bank of three lights each. The subject was able to respond as quickly when the warning was provided 500 ms before presentation of the target as when the target was selected from just three lights. Leonard was able to demonstrate, as many others have since, that forewarning about the target set permitted the subject some advanced preparation that facilitated target processing.

What is the preparatory mechanism implicated by this and other similar experiments? This has been an often-asked question in many laboratories during the past several years, and it has motivated research on various candidate mechanisms. Broadly speaking, there are two classes of mechanisms that have been uncovered (Posner & Boies, 1971; Posner, Klein, Summers, & Buggie, 1973; Sudevan, 1984; Taylor, 1977). The first class consists of processes that cause a general warning effect. For example, imagine that one of the effects of the selective forewarning in Leonard's (1958) experiment was to heighten the subject's general alertness by whatever means, and thereby increase overall attentiveness to all the stimulus lights in the task. This heightened attentiveness might have caused the subject to respond more quickly.

The second class of processes consists of those that might selectively prepare a subject to process the cued signal more efficiently. Referring again to Leonard's (1958) experiment, consider the possibility that a selective warning signal altered the order in which the subject sampled from the two banks of stimulus lights in searching for an extinguished light, or that it caused the subject to foveate the critical bank of lights. Either of these mechanisms might have resulted in facilitated processing.

These classes of processes—general warning effects and specific preparatory effects—have been of interest to cognitive psychologists (Sudevan, 1984). Therefore, experimental techniques have been devised to isolate each class of effect and to study them further. The technique that concerns us here is cost–benefit analysis, which was developed to investigate selective preparatory effects. In the form devised by Posner and Snyder (1975), the paradigm requires two conditions. In one, subjects execute a task after having been warned by a neutral cue that presumably does not permit selective preparatory processes. For example, if one were to rerun Leonard's (1958) task to include a neutral cue, the cue might be a general warning tone presented prior to having one of the six lights extinguished. Performance in this neutral condition is then compared with performance in a condition in which a selective (informative) preparatory cue is presented. However, Posner and Snyder's (1975) paradigm, unlike Leonard's, typically has the selective cue be informative less than 100% of the time. Thus, for example, in our hypothetical replication of Leonard's task, a subject might have received a tone in one earphone that indicated whether the bank of lights on the left or right would contain the critical target. However, the tone might have been correct (valid) on only, for example, 80% of the trials and incorrect (invalid) on the remaining 20%.

This procedure results in two types of trials when a selective preparatory cue is presented: those on which the cue is valid, and those on which it is invalid. Consequently, two comparisons may be made against the neutral condition: one for the valid trials to assess a benefit in performance that may accrue with an informative cue, and one for the invalid trials that could show a cost in performance due to an informative but misleading cue. If the neutral and informative cues result in reliably different performance (in reaction time or error rate), one can argue for a selective preparatory effect of the informative cue that is different from, or at least added to, the general warning effect that might be produced by a neutral cue.

Cost–benefit analysis has become quite popular as a tool to diagnose preparatory processes in cognition, broadly construed. Applications in the last two decades have focused primarily on how prior context can provide advance information about a target stimulus. The technique has been applied to fairly simple perceptual tasks (e.g., Posner, Nissen, & Ogden, 1978; Jonides, 1981) to investigate how attention is allocated to stimuli in various parts of the visual field or to study different levels of stimulus encoding. It has also been applied to word identification tasks to examine the relation among the concepts underlying the words (e.g., Neely, 1976, 1977; Becker, 1976) and the influence of more complicated episodic and semantic relations between words (e.g., Blaxton & Neely, 1983; Durgunoglu & Neely, 1982; Neely, Schmidt, & Roediger, 1983). In addition, the technique has been applied to word identification in the context of sentence comprehension (e.g., Fischler & Bloom, 1979; Schuberth & Eimas, 1977; Stanovich & West, 1979, 1981).

Here the technique has been used to examine how sentence processing might influence the recognition of words that occur in sentences, both for adults (as in the studies just cited) and developmentally (e.g., Simpson & Lorsbach, 1983; West & Stanovich, 1978). Its broad application can be traced to at least three factors. First, as discussed previously, the paradigm has the potential to isolate specific preparatory processes from the melange of processes that might generally facilitate task performance.

Second, application of the technique has sometimes produced a pattern of costs and benefits that varies with the interval between cue and target stimulus. One such pattern, uncovered by Posner and Snyder (1975) and Neely (1977) among others, has been interpreted as evidence for the operation of both automatic and attentive preparatory mechanisms in target processing, which are supposedly separate from general alerting. These investigators found that at short intervals between cue and target there were only benefits and no costs associated with the cue. At longer intervals, however, there were both costs and benefits. This pattern of results has led to the proposal of a two-component model. One component is automatically activated and can cause very fast selective preparation for processing with little diversion of processing resources from other cognitive operations. This component is followed by one that requires more effort. Regarding this component, selective preparation for a target is accomplished at the expense of potential analysis of other stimuli. A two-component model of this sort has been suggested for a variety of preparatory effects (see, e.g., Fischler, & Bloom, 1979; Mack, 1981; Stanovich & West, 1979, 1981).

Third, the inclusion of valid, invalid, and neutral trials provides an opportunity to obtain quantitative estimates of the magnitudes of costs and benefits (including, simply, which is larger). These in turn can be used to test various specific models of preparatory mechanisms. For example, in the context of a visual search task, Jonides (1980, 1983) has obtained evidence that contradicts a particular model of array examination. Gathering this evidence required a selective examination of the cost of processing and of the relative magnitudes of costs and benefits. Becker (1980) also observed different quantitative patterns of cost and benefit in a study of semantic priming in word recognition. The observation that priming was dominated sometimes by inhibition and other times by facilitation led to hypotheses about qualitatively different mechanisms of semantic priming.

For these reasons and others, cost–benefit analysis has assumed a place of prominence among the diagnostic instruments of cognitive psychologists. However, in order to make the most of its diagnostic role, one must be cautious in applying the technique and interpreting results from its application. The major reason to exercise caution concerns the relation between neutral and informative cues. We now discuss the assumptions underlying the comparison of conditions with these two types of cues.

Cost–benefit analysis relies on the same rationale as Donders' (1969) subtraction method. Presumably, both informative and neutral cues serve general warning functions. For this reason the general facilitation of processing produced by both cues should be identical. They differ in that the informative cue, and not the neutral cue, provides the subject with some additional information about the target stimulus. Thus, one should be able to attribute any difference in performance between neutral and informative cues—that is, any costs or benefits—only to the specific preparatory processes that the subject engages.

It is clear that this rationale hinges on a critical assumption: Neutral and informative cues must be identical with respect to all their effects except that of information specific to the target. What if this were not the case?

The most straightforward implication of failing to satisfy this assumption is that the relative magnitudes of costs and benefits would be rendered meaningless. Suppose, for example, that for any of the reasons discussed below the neutral condition yielded response times that reflected not only a general warning effect of the cue, but also an effect of other factors that uniquely interfered with its processing. These other factors, by impacting on the neutral condition only, would not be subtracted out when evaluating costs and benefits. Thus, costs and benefits would be incorrectly attributed solely to the specific preparatory effects of the informative cue rather than to these specific preparatory effects combined with the factors that influenced the neutral condition alone. In this hypothetical example, performance with a neutral cue could not be meaningful compared with valid or invalid cues.

In principle, this line of reasoning is perfectly sensible. However, if one could not identify any factors such as those suggested, then the argument would not be worrisome. However, there are several such factors. We now discuss four factors that might play a role in influencing neutral and informative conditions differently.

Differential Attentiveness

An informative cue is a signal for subjects to engage in some additional processing beyond merely increasing their readiness to respond. Suppose that subjects interpret this, whether consciously or not, as a directive to be even more prepared to perform when an informative cue is given than when a neutral cue is given. That is, they not only might prepare specifically for the target, but they might also become more alert generally. The net effect of increased alertness due to an informative cue would be lowered costs and increased benefits when the informative cue's effects are compared with those of the neutral cue. We have identified four cases in which differential alertness should be suspected.

Different Cues

One circumstance is when the cues are quite different physically. This has occurred quite frequently in studies with sentences as primes. In this case, neutral cues have frequently been nonsentence stimuli like rows of Xs (e.g., Fischler & Bloom, 1979; Irwin, Bock, & Stanovich, 1982; Schuberth & Eimas, 1977; Mack, 1981) or single words like the determiner "the" (Stanovich & West, 1979, 1981). Fischler and Bloom (1979), for example, compared response latencies for target words preceded by anomolous sentence contexts in one experiment to latencies for words preceded by a neutral stimulus (consisting of Xs) obtained in an earlier experiment. Stanovich and West (1979) compared pronunciation latencies for target words that were preceded by a sentence context with those for words preceded by a nonsentence context (the determiner "the") at two different intervals of time between offset of reading the prior sentence and onset of the target word. These two investigations illustrate what is true of virtually all studies in this area: Neutral contexts consist of nonsentence stimuli. The potential problem with this lack of comparability between neutral and informative primes is that neutral contexts and meaningful contexts (consisting of meaningful sentences) may engage different processing demands.

These concerns have led some investigators to consider alternative neutral stimuli like lists of words (Forster, 1981) or sentence frames that lack a property associated with the process whose costs and benefits one wishes to assess (Fischler & Bloom, 1980; Stanovich & West, 1983; West & Stanovich, 1982). Forster (1981) predicted that theoretically facilitation effects resulting from sentence contexts should be relatively small or nonexistent (at least when specific lexical facilitation is factored out). However, he pointed out that supporting evidence for small facilitation effects (e.g., Fischler & Bloom, 1979) is ambiguous because facilitation effects could be underestimated by one's choice of neutral condition. Accordingly, Forster conducted sentence priming studies that used lists of words as a neutral condition. Forster drew on several considerations to argue that such stimuli are more appropriate than nonword (e.g., Xs) or nonsentence (e.g, the word "the") neutral primes. (Forster found evidence, incidentally, that sentence context effects are obtained only under more specific task conditions than generally assumed.)

In the case of word prime studies, neutral cues tend also to be nonword stimuli like Xs (e.g., Antos, 1979; Balota, 1983; Becker, 1976, 1980; Favreau & Segalowitz, 1983; Goodman, McClelland, & Gibbs, 1981; Koriat, 1981; Larochelle, McClelland, & Rodriquez, 1980; Neely, 1976; 1977). Even here investigators have questioned the choice of nonword neutral stimuli. For example, de Groot, Thomassen, and Hudson (1982) rejected the use of Xs in a word-priming experiment on the basis of evidence that response times to target stimuli were slowed as compared with target stimuli preceded by word primes. There is evidence that subjects expect primes to be words, so that targets preceded by nonword primes are actually treated as if they were the priming stimuli rather than the targets. The effect of this is to slow response to the target. Such an effect would tend to underestimate inhibition effects and overestimate facilitation effects. Consequently, de Groot et al. (1982) opted for a neutral word ("blank") as a prime.

This solution, however, raises other problems with respect to differential processing on neutral and non-neutral cues, as they point out. Neutral stimuli tend to be repeated from one neutral trial to another. In the case of words, the neutral prime would tend to become meaningless, making less memory demands than non-neutral word primes, which vary from trial to trial. Similarly, repeating the same nonword prime from one trial to another introduces redundancy in the neutral stimulus; the impact of redundancy on attentiveness is difficult to assess. There is evidence that people respond more slowly to a stimulus (in a simple or choice reaction time task) when it has been preceded by several presentations of itself (e.g., Kraut & Smothergill, 1978; Kraut, Smothergill, & Farkas, 1981). Although subjects do not respond directly to the priming stimuli in cost–benefit studies, repetition may engage a similar mechanism to diminish a subject's attentiveness to all stimuli on such trials.

This discussion makes clear that cost–benefit analysis hinges on the assumption that neutral and informative cuing conditions must be identical with respect to all processing consequences of the cue except the specific preparatory effect elicited by the informative cue. To the extent that this assumption is not satisfied, one is not justified in attributing benefits and costs in performance solely to the selective preparatory effect of the cue. Perhaps the most serious difficulty, though, is that it will be extraordinarily troublesome in any specific case to establish that one or another of the alternative accounts discussed above is not relevant.

Unfortunately, it is not easy to determine whether the data in any given study are tainted by a differential effect of attentiveness because no proper control is included that preserves the essential aspects of the experimental manipulation, but removes the conditions that may promote differential attentiveness. We suspect that in many cases investigators who worry about an attentiveness problem simply examine whether the relationships among neutral, valid, and invalid performance seem orderly or well behaved (i.e., valid is better than neutral, which is better than invalid). If so, they accept the appropriateness of the conditions. This is, of course, an extremely weak criterion. It may be that the ordering of absolute performance in the various conditions fortuitously results in the neutral condition falling between valid and invalid. As readers of the literature, we ought to have another worry as well. Because we do not know how many experiments with ill-behaved neutral conditions are not published, we cannot evaluate whether the published ones are simply a biased sample. Consequently, it is difficult to use the well behavedness of a neutral condition as an important criterion of the validity of the method.

Different Presentations

Another circumstance that may give rise to differential attentiveness for neutral and informative cues is when the two types of cues are presented during different major epochs in an experiment (e.g., sessions or blocks of trials). Consider, for example, data from a neutral condition that was not included in an experiment by Jonides (1980, Experiment 2). The purpose of this experiment was to examine the effect of validity of an informative cue on costs and benefits. Accordingly, for each level of validity, trials with a neutral cue and trials with an informative cue were intermixed. Before adopting this procedure, however, an alternative procedure had been considered, with a separate neutral condition given to a different group of subjects so that neutral trials would not have to be interspersed with each level of informative cue validity. Pilot experimentation with this technique led to its rejection because the data from the separate neutral condition were quite different from those of the neutral condition that was intermixed with the valid and invalid trials. For example, when the informative cue had a validity of 70%, the intermixed neutral trials yielded a mean response time (RT) of 610 ms, whereas the separate neutral condition resulted in a mean RT of 662 ms. (The valid trials in this condition had a mean RT of 512 ms, and the invalid trials had a mean RT of 746 ms.) It is easy to see that the choice of one or the other neutral control would have had quite different consequences for the magnitudes of costs and benefits. In this case, the experimenter chose the intermixed control, because there was a distinct possibility that subjects in the separate neutral condition were performing suboptimally. Choosing the intermixed condition at least permitted the argument that attentiveness was roughly equated between conditions because the order of conditions was unpredictable. We stress, however, that even this is not a completely convincing argument.

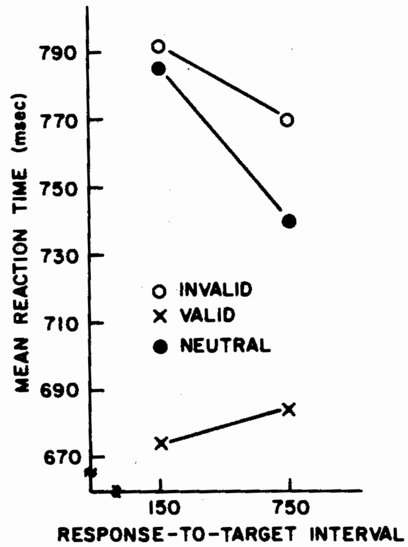

Figure 1. Mean reaction time as a function of response-to-stimulus interval from Stanovich and West's (1979) Experiment 2.

Different Frequency

A third circumstance that may induce differential attentiveness is when neutral cues occur less frequently than non-neutral cues. This tends to be a function of the size of the set of target stimuli. For example, differences in frequency are typically smaller for studies using a small set of target stimuli (e.g., words from semantic categories as in Neely's, 1976, 1977, studies). However, all studies using cost–benefit analysis suffer this problem, and it is not clear what impact this might have on assessing baseline performance for target stimuli preceded by neutral cues. There is, of course, evidence from studies of recognition memory and simple choice reaction time for words that repetition and semantic relations between a target word and items preceding it (priming context construed broadly) can influence task performance for the target word (e.g., for choice reaction time, see Kraut & Smothergill, 1978; Kraut, Smothergill, & Farkas, 1981; for recognition memory for words, see Carroll & Kirsner, 1982; Durgunoglu & Neely, 1982; Neely, Schmidt, & Roediger, 1983). In recognition memory tasks, semantic relations do not always produce facilitation. The methodological implications of such effects still need to be examined and require a better theoretical understanding of the phenomena. (The interested reader should see Durgunoglu & Neely, 1982; Neely, Schmidt, & Roediger, 1983; Roediger, Neely, & Blaxton, 1983).

Different Processing Time

Consider now a final case in which there may be differences in the alerting function of neutral and informative cues. We illustrate this point by referring to Stanovich and West (1979, Experiment 2), in which subjects pronounced target words in a timed test. The targets were preceded either by a sentence fragment context that was congruous (valid) or incongruous (invalid) or by a neutral context (the word "the"). Targets were presented either 150 ms or 750 ms after the subject read the context. Results are presented in Figure 1. At the response-to-target interval (RSI) of 150 ms, the authors reported a reliable benefit and no reliable cost. At RSI 750 ms, they reported both a reliable benefit and a reliable cost. This pattern of results, as well as other results derived from a variation in target quality, led

Stanovich and West to favor the two-component model of Posner and Snyder (1975).

However, Kleiman (1980) made critical comments about this conclusion. He recognized that much of the increase in cost between the two RSI values was due to the sharp drop in RT for the neutral cue. Further, he argued that if there had been an attentive mechanism operating at RSI of 750 ms, but not at RSI = 50 ms, then the absolute level of RT for invalid trials should have increased at the longer RSI. There should not have been just an increase in the difference between neutral and invalid RTs.

What could account for the pattern of results reported by Stanovich and West (1979)? Kleiman (1980) reasoned that an increase in RSI caused an increase in the general warning effect of the neutral cue. This increase overshadowed any comparable increase for the informative cue, because at a 150-ms RSI the informative cue was actually available for much longer, namely the time subjects took to read the sentence fragment plus the RSI. According to this reasoning, the increase in RSI with the neutral cue provided subjects with a better opportunity generally to prepare themselves for the upcoming target. Hence, their response times declined. This decline, in turn, led to the appearance of a difference in response time between neutral and invalid trials. According to this explanation, then, the presence of a cost in responding at one RSI but not at another is best attributable to a change in the general preparation induced by the neutral cue, not to a change in specific preparatory processes attributable to the informative cue.

This alternative interpretation of the Stanovich and West (1979) data relies on the fact that stimulus onset asynchrony (SOA) was not equated for neutral and informative cues, because informative cues took longer to read. However, the argument can be made more general than this as well. What is critical to the most general form of the argument is the possibility that a neutral cue may be noticeably less potent than an informative cue in generally alerting subjects about an upcoming target. In addition, as seems reasonable, its potency may well be enhanced by stretching out the interval between its presentation and presentation of the target. Increasing SOA, according to this line of reasoning, may have two effects. It may bring specific preparatory processes into play with informative cues, but it may also heighten the general warning effect of the neutral cue and alter its processing time.

Fortunately, it ought to be relatively easy to detect the contribution of changes in the effectiveness of the neutral cue in any specific case. As Kleiman (1980) makes clear, a proper analysis of one's data ought to include examination of both cost–benefit effects and absolute levels of performance in the various conditions. This does not permit an unambiguous evaluation of an interaction such as that in Figure 1, but it does allow one to test whether there is any contribution of general, as well as specific, preparatory effects.

Overall, it should be clear that a variety of conditions exist that may cause differences in attentiveness to neutral and informative cues. Guarding against some conditions, such as those that may be induced by a blocked presentation of cues, may merely require some thought in experimental design. Others may be more insidious.

Extraneous Processing Demands

By virtue of the fact that neutral and informative cues have different signaling functions, they may exert different processing demands, which may in turn influence reaction time to a target stimulus.

To illustrate the problem, consider an experiment by Schuberth and Eimas (1977). In the main condition, subjects were engaged in a lexical decision task. Prior to this, either a congruous or incongruous context was presented. The incongruous context yielded RTs longer than the congruous context by 59 ms. To assess costs and benefits, Schuberth and Eimas (1977) included an isolated target condition in which no context preceded the target. With this control, costs and benefits were 20 ms and 25 ms, respectively. (These values do not total 59 ms because separate isolated target conditions were used as controls for congruous and incongruous contexts.) What can be made of these values?

Schuberth and Eimas (1977) recognized an obvious difference between neutral and informative conditions in this experiment. If subjects attend to the sentence contexts in this experiment and keep them in mind during analysis of the target, the extra memory load that results may hinder target processing. To examine this possibility, Schuberth and Eimas included another neutral condition in which strings of unrelated digits preceded each target. Subjects were required to recall the sentence or digit contexts on each trial to ensure that they were held in memory. The result was clear: The digit contexts produced slower lexical decisions than the isolated target condition (by 67 ms). This indicates that memory load does affect lexical decision times; so much so that if one were to calculate the costs and benefits of sentence versus neutral digit contexts, they would be 95 ms and −49 ms, respectively (yes, that is *negative* 49 ms!). Consequently, there are two sets of costs and benefits depending on the control condition. Which is more believable?

There is no easy way to decide. An extra memory load caused by a sentence compared with no context seems to be implicated by these data, but to control for it one needs a good estimate of the memory load imposed by the sentences that could be mimicked in a neutral condition. It was only chance that Schuberth and Eimas (1977) guessed correctly about this memory load in their digit control. Consequently, we are left with no truly good estimate of actual costs and benefits, or even an idea of whether there are costs and benefits. The conclusion drawn by the authors that "sentence contexts facilitated the processing of congruous words and interfered with the processing of incongruous words" is premature.

Manipulating the Cue-to-Target Interval

All theories of processing in the cost–benefit task claim that it takes time to accrue the consequences of a cue. Furthermore, according to some theories, changes in costs and benefits with time might be diagnostic about the process underlying changes in performance with a cue. In this regard, the model proposed by Posner and Snyder (1975) has been of great interest to investigators. This model, as described above, asserts that there are two component processes that can cause a difference in performance for an informative versus a neutral cue. The first is an automatic component whose operation results in facilitation with no inhibition. This component can be activated shortly after cue presentation. The second component is one that relies on effortful, attentive allocation of processing resources to the specific stimulus indicated by an informative cue. Assuming there is a limit on the total good of resources available for processing at any time, this component will confer a cost of processing if a cue is invalid in addition to a benefit if it is valid. Furthermore, because the allocation of resources is relatively slow, consequences of activating this component will take measurably more time to detect than will the consequences of automatic activation.

This description should clearly explain why a substantial number of studies that have used cost–benefit analysis have also varied the interval between cue and target. Interpreting effects of this interval, however, is not a simple matter. We need to consider possible differences in encoding time due to our choice of neutral cue and due to the manipulation of the cue-to-target interval.

Encoding Time Difference for Neutral and Nonneutral Primes

An informative cue provides subjects with more information than does a neutral cue. It may well take more time to encode and extract this information from the informative cue than it does to encode the neutral cue. Furthermore, once extracted, data about informativeness must be used by engaging the appropriate processes to prepare specifically for the indicated target. The net effect of added encoding time for an informative cue would be to add some increment to all informative cue response times if the cue is not completely encoded by the time the target appears. That is, the subject's processing of the target would presumably be delayed by the subject's continued encoding of the cue (in whole or in part depending on whether processing subsequent to encoding overlaps encoding itself). Assuming a constant level of performance in the neutral condition, this would result in an underestimation of benefits and an overestimation of costs in processing.

Are there examples of this artifact in the literature? Once again, this is difficult to evaluate because the issue has not been addressed directly with controlled experimentation. However, if there are examples, they would be studies in which the interval between presentation of the cue and target (ISI) is short. For example, in the experiment by Schuberth and Eimas (1977), there was an interval of 0 ms (although cuing sentences were presented for 1.5 s). The problem is eliminated, of course, if subjects are given unlimited opportunity to encode the cue prior to presentation of the target, which is presumably the event that marks the beginning of reaction time measurement in most experiments. However, if processing the cue is permitted to intrude on the duration of target processing, one may expect an effect of the sort under discussion.

One way of ensuring that this artifact does not contaminate response times is to have subjects process cues *ad libitum*. However, this is not possible if one wishes to examine the effect of the interval between cue and target on response times. In this case, one strategy might be to estimate, using a separate task, the amount of extra time required to encode the informative cue compared with the neutral, and then to correct all reaction times on the neutral and informative trials by this difference. (See Fischler & Bloom, 1979, Experiment 2, for a similar attempt to determine differences in encoding time for informative versus neutral cues.) This use of Donders' (1969) logic, however, is somewhat fragile (see Fischler & Bloom, 1979, Experiment 3, for their attempt to overcome this fragility). It depends on one's confidence in the independent measurement of the encoding time for cues, and this is potentially troublesome because the measurement must be accomplished in a task that is closely comparable with the task of interest; this condition is difficult to satisfy.

An alternative experimental strategy to address this problem might rely on the use of catch trials inserted into a normal trial sequence.[1] Occasionally, one could replace the target that typically appears after a prime with a probe stimulus to which the subject has been trained to give a simple response. If probes inserted after neutral cues lead to simple reaction times different from those that appeared after informative cues, this would provide grounds to worry about differential encoding operations. If probe times were the same regardless of the prime that they followed, one might argue that neutral and informative probes had equivalent alerting functions, and that encoding times for the primes were equivalent. (Results of variation in the time of appearance of the probe might strengthen this argument.) An example of a similar technique is a study of letter priming by McClean and Shulman (1978). They were interested in more than the methodological question raised here, but they did indeed find longer probe latencies following non-neutral cues than neutral.

Encoding Time Differences Resulting From Manipulating Cue-to-Target Interval

The cost–benefit methodology involves at least three temporal relations: cue duration, interstimulus interval, and interval between the onset of the cue and the onset of the target (or SOA). Consider which interval one should vary. In principle, of course, a subject can begin to encode a cue as soon as it is presented. This suggests that the interval that most reasonably reflects the time available to activate processing induced by a cue is the SOA. Most investigators concerned with temporal issues have, in fact, manipulated this interval. In one case, however, this is simply not feasible. If it takes appreciable time to encode the cue, as it would with a sentence as cue, then manipulating SOA runs the danger of introducing values too short to permit full encoding of the cue prior to onset of the target. The result in this case is that at short SOAs, one may grossly underestimate the impact of the cue had it been fully analyzed.

Realizing this, some investigators, such as Fischler and Bloom (1979), have presented subjects with cues for a sufficiently long time to ensure their processing. Of course, this precluded variation in the cue-to-target interval (although one could manipulate the latter independently of cue duration in order to vary the overall SOA interval; see Mack, 1981). Alternatively, Stanovich and West (1979) had subjects read sentence cues aloud *ad libitum*, and they varied the interval between reading the end of each sentence and the onset of the target (or RSI). In a sense, they tried to achieve both complete processing of the cue and variation in cue to target interval. The result of

[1] The authors thank David Meyer for this suggestion.

this study was that a 150-ms RSI produced only a reliable benefit, whereas a 750-ms RSI produced reliable cost and benefit. The authors attributed this pattern of results to the two-component model discussed above even though at a short RSI subjects still had substantial time to process the cues (i.e., long SOA). Kleiman (1980) offered an alternative interpretation of these results.

Regardless of the results of previous research, however, the problem remains: What is the proper interval to vary? The solution is not simply a choice between one or another temporal parameter. Ideally, what one needs control over is the internal interval between when the cue is first recognized and when the target's processing is initiated. Setting SOA, ISI, or any other temporal parameter can only hope to approximate setting this internal interval. Furthermore, it cannot be guaranteed that if the nominal intervals set for neutral and informative cues are the same, then the internal intervals will also be the same.

Let us consider two examples that illustrate the potential effect of this uncertainty about internal SOA. Both examples draw on the critical point made originally by Taylor (1977) that when cue and target differ substantially in their processing demands, the relation between internal and external SOA is not easily determined. We extend this argument by suggesting that the internal interval can be different for neutral and informative cues, thus further confusing the calculation of costs and benefits as a function of nominal stimulus onset asynchrony.

Consider first a hypothetical experiment that schematizes the results typically used to implicate Posner and Snyder's (1975) two-component model of priming effects. In this experiment, there is a neutral cue, a valid cue with a validity higher than chance, and variation in stimulus onset asynchrony. The prototypical reaction time results of this experiment that supposedly support the two-component model are displayed in Figure 2 (for simplification, error data will not be discussed). Figure 2 presents these data for neutral, valid, and invalid trials as a function of SOA. These data might lead an investigator to conclude that at SOA value A, only an automatic facilitation effect is operating, whereas at SOA value $A + X$, an attentive preparatory strategy has taken over, resulting in costs and benefits.

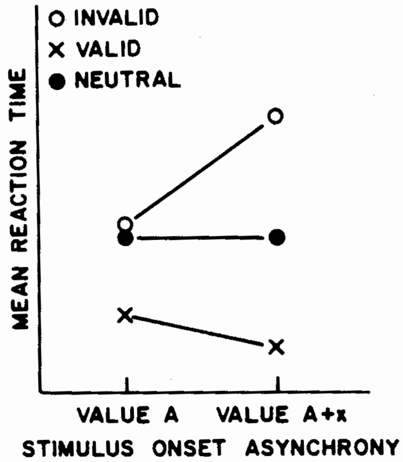

Figure 2. Hypothetical reaction time data as a function of stimulus onset asynchrony that are typical of the sort used in support of the two-component model of Posner and Snyder (1975).

Now let us consider an alternative view of these results. Suppose that the neutral and informative cues differ substantially in this experiment in terms of the amount of time required to encode them, as they might, for example, if the informative cue were a complex sentence fragment and the neutral cue were a simple row of crosses. Suppose, in fact, that the difference in encoding times for the cues were approximately equal to the difference between SOA values A and $A + X$ in Figure 2. If this were so, then it does not seem proper to compare the results of neutral versus informative cues with the same nominal values of SOA, because the internal asynchrony values of the cues are not at all equal. In fact, one might argue that the proper comparison in this case would be between value A for the neutral condition and value $A + X$ for the informative cues, in which case the asymmetry between cost and benefit disappears. This argument would be based on the fact, as hypothesized for this case, that when the informative cue precedes the target by value $A + X$, by the time it is properly encoded the actual asynchrony between it and the target is actually value A.

To make this point perfectly clear, consider Figure 3. Hypothetical values are presented for the encoding times of neutral and informative cues, and the values of asynchrony plotted in Figure 2 are displayed. According to the values in this figure, when the infor-

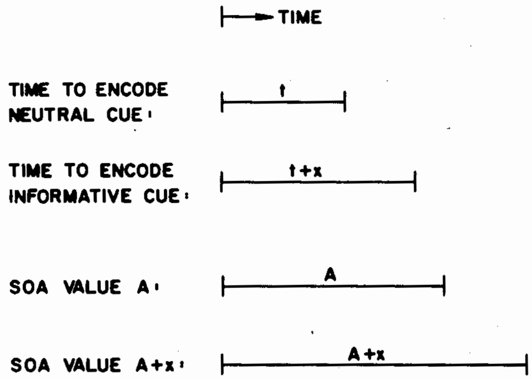

Figure 3. Schematic of various temporal values from the hypothetical experiment whose results are plotted in Figure 1.

mative cue is presented at SOA value $A + X$, its effective SOA (i.e., the time elapsing between the end of encoding of the cue and presentation of target) is $(A + X) - (t + X)$, or $A - t$. This is obviously not equal to the effective SOA of the neutral cue when it is presented at SOA value $A + X$, which is $(A + X) - t$. However, the effective SOA of the informative cue at the nominal SOA value $A + X$ is equal to the effective SOA of the neutral cue at the nominal SOA value A, that is $A - t$. Thus, by this argument, Figure 2 is actually improperly plotted. Given only the values in the figure, we can make only one proper comparison—between the RT of the neutral condition at value A and the informative condition at value $A + X$. This comparison indicates both costs and benefits in RT. It thus questions the relevance of the two-component model that was improperly implicated by analysis of nominal SOA values.

Of course, it will rarely be the case that, by coincidence, differences in encoding times for the cues will happen to be exactly equal to the interval over which an effect restricted to just benefits changes into an effect of costs and benefits. The point is that as long as there is an unknown difference in encoding times between the cues, there will be uncertainty about the proper comparison of cues as a function of SOA. This will have an impact on drawing conclusions about relative magnitudes of costs and benefits.

In order to dispel the impression that this argument holds only for the pattern of results presented in Figure 2, consider another experiment. This is an actual experiment whose results do not support the two-component model. In Figure 4 some of the results of Taylor's (1977) Experiment 1 are plotted. In this experiment, subjects were presented one of two targets on each trial: the letter K, which required a left-hand key response, or the letter T, which required a right-hand key response. At some stimulus onset asynchrony with the target (from −250 ms to 250 ms), two tokens of a context letter were presented flanking the target. The context could be the letter K, T, or O. Figure 4 plots the data from valid, invalid, and neutral trials that were equiprobable, as a function of the SOA values from −100 ms to 100 ms (negative values are defined as the target preceding the context). The figure suggests that costs increase with SOA sooner than benefits, a result that has implications for the model one proposes for these data. However, consider the possibility that it takes longer for subjects to encode the informative than the neutral cue. This might be so because the neutral cue in this experiment was a letter not assigned to a response, whereas the informative cues were each associated with unique responses. Consequently, if subjects prepared for the target by readying one response when given an informative cue, then it may have taken them an extra processing step to identify which response to prepare once they identified that

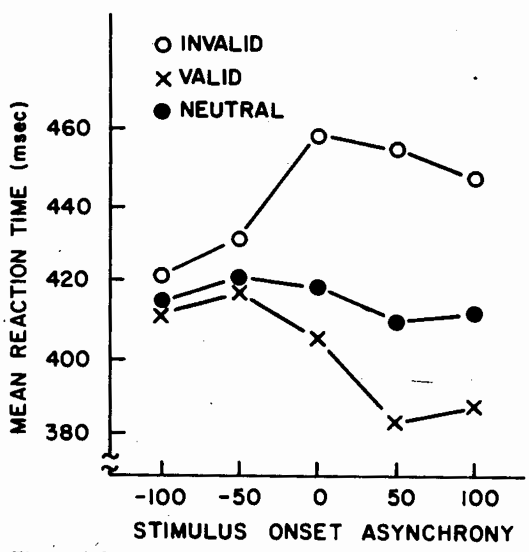

Figure 4. Mean reaction time as a function of stimulus onset asynchrony for selected data from Taylor's (1977) Experiment 1.

the cue was informative, assuming this order of processes. The implication of this analysis is that the functions for valid and invalid responses in Figure 4 should be shifted to the left by some unknown amount to achieve a proper comparison with the neutral trials. Although we cannot estimate the extent of this shift, even small amounts will obviously have a profound effect on the relative magnitudes of costs and benefits at any SOA.

Speed–Accuracy Trade-Off

The problems discussed above are inherent in the application of cost–benefit analysis. There is a methodological problem which, although not inherent in cost–benefit analysis, should also be mentioned. This pertains to the possibility of speed–accuracy trade-offs in response time measures.

Antos (1979) has shown that subjects can trade speed for accuracy between informative and neutral cue conditions. Subjects may, for example, exhibit a superiority in reaction time performance for a valid condition compared with a neutral condition, having achieved this via less accuracy in the valid condition. This fact, coupled with the serious difficulties that arise in interpreting reaction time data for accuracy levels near ceiling (Pachella, 1974), suggests that one must be cautious in deciding whether there are any costs and benefits in a particular task.

Now typically informative cues that result in benefits and costs in response times also result in benefits and costs in errors. Disregarding the arguments discussed previously, this pattern of results permits confidence in the existence of cost and benefit effects. However, it should not permit one to be confident about the relative magnitudes of the two effects.

We now consider an experiment presented by Posner and Snyder (1975) as an example. In this experiment, pairs of letters were presented for matching. The target pairs were preceded by either a neutral cue or a letter cue that matched none, one, or both of the letters in the target pair. When the letter cues were 80% informative about the target pair, the response times for neutral, valid, and invalid cues were 414 ms, 329 ms, and 450 ms, respectively, indicating cost and benefit in response times. The corresponding error rates were 8.7%, 2.4%, and 39.5%, respectively, also showing cost and benefit effects. If one were interested in assessing whether those priming effects were due to automatic or strategic processes, one might wish to determine which effect—the cost or the benefit—was larger. If we study only the response times, the benefit is larger; if we study only the errors, the cost is larger. In this case, then, one would need a precise calibration of the trade-off between accuracy and time to answer the question; this trade-off must be empirically determined. If we generalize from the results of Antos (1979), the cost is probably substantially larger in this study. (See Becker, 1980, Experiment 5, for another example of this same problem in a different task.)

The problem of a trade-off can have an even more serious consequence, however, than simply obscuring the relative magnitudes of cost and benefit effects. Frequently, investigators have been interested in whether there is any cost for a particular task. This is of interest because it has been proposed that a benefit with no cost is a symptom of automatic facilitation by a cue with no attendant inhibition (Posner & Snyder, 1975). However, one must keep in mind that both response time and error effects must be examined to evaluate whether there are costs. Antos (1979), for example, has noted possible evidence for a speed–accuracy trade-off in an aspect of Neely's data (1977), which suggests that lack of inhibition at short SOA intervals may result (in whole or part) from a trade-off with errors. However, there is converging evidence for Neely's conclusion in the full study, as he has pointed out.

Conclusions and Recommendations

We have reviewed multiple sources of problems that require care in interpreting the results of cost–benefit analysis. Students of the recent cognitive literature may have recognized the arguments presented as problems that occurred to them while reviewing the many articles that have adopted the cost–benefit technique. Indeed, ideas similar to ours found in the literature have been briefly discussed, although these discussions are often buried in empirical papers (de Groot et al., 1982, Fischler & Bloom, 1980; Forster, 1981; Kleiman, 1980; Schuberth & Eimas, 1977; Taylor, 1977). Our purpose here was not to raise a wholly novel set of issues about the cost–benefit technique. The ones that have lurked in the minds of experimenters will do well by themselves. Rather, we have collected the central problems all in one place. The purpose is threefold. First, it brings these concerns about the technique to the surface for all to study. Second, it may change the prior subjective odds of accepting conclusions by readers of the cognitive literature who encounter articles that use cost–benefit analysis. Third, it may promote better informed application of the technique and wiser evaluation of experiments in which it is applied.

We now comment on this last point because the analysis presented above leads to two recommendations about applying cost–benefit analysis:

If Possible, Do Not Include a Neutral Condition

Frequently, experiments are conducted with the goal of discovering whether there are specific preparatory effects elicited by informative cues. To meet this goal, it is sufficient to discover merely whether performance differs with valid and invalid cues and whether differences between these conditions grow with the interval between cue and target. It is not necessary to include a neutral condition and to calculate cost and benefits (e.g., see Goodman et al., 1981; Jonides, 1980; Koriat, 1981; Larochelle et al. 1980; Meyer et al., 1975; Schuberth, Spoehr, & Lane, 1981; Tweedy, Lapinski, & Schvaneveldt, 1977).

The inclusion of only valid and invalid trials allows one to go beyond the mere detection of a specific preparatory effect if one examines the absolute level of performance with these trials as well as their differences. To illustrate consider experiments by Tweedy et al. (1977), Koriat (1981), and Jonides (1980, Experiment 2). In all these papers, the authors were interested in determining whether the probability of an informative cue being correct (i.e., the validity) had an influence on the preparatory effect induced by the cue. To this end, Tweedy et al. (1977) used three groups of subjects given three different cue validities in a task requiring lexical decisions. Likewise, Jonides (1980) had three groups of subjects who differed in the cue validities assigned to them in an experiment that investigated visual search performance. Both studies discovered monotonic effects of validity on the influence of cue informativeness. The more informative the cue, the greater the difference (in RT or errors) in performance between valid and invalid cues. In this study, Jonides (1980) also included a neutral condition that was intermixed with the valid and invalid trials for each of the groups. However, the added information gained from having neutral trials was negligible. It merely provided an opportunity to assess whether overall performance levels among groups were comparable. Strictly speaking, this was not necessary to the main analysis of invalid–valid performance. Furthermore, by including a neutral condition there is the risk of obtaining varied patterns of costs versus benefits for the three groups. Such patterns might have been taken as serious evidence for one or another model of processing even though, as discussed, these patterns might have been caused merely by various factors that produce fluctuation in neutral condition performance.

To summarize, most of the arguments we have reviewed raise questions about the appropriateness of neutral conditions. If possible, the most straightforward strategy to finesse these questions is to avoid them by avoiding the conditions that prompt them.

If a Neutral Condition Is Included, Be Cautious

As suggested earlier, it is sometimes valuable to have estimates of the magnitudes of costs and benefits to evaluate certain models, such as those of Posner and Snyder (1975) or Becker (1980). In these cases, it seems prudent to observe two sorts of precautions.

First, the cues used for neutral and informative conditions should be matched as closely as possible. This includes matching them on physical appearance, on potential to alert subjects generally, and on ease of encoding. Comparability is especially pertinent when complex stimuli (e.g., sentences) are used as non-neutral cues. The problem is not necessarily intractable. Several studies suggest that sentences themselves can be used as relatively neutral stimuli.

Stanovich and West (1979, 1981), for example, used neutral stimuli consisting of the relatively contentless frame "it was the." Kleiman (1980) and Fischler and Bloom (1979, 1980) compared performance for target words that were predictable from a preceding sentence frame with performance for words that were not predictable but were semantically appropriate. (Kleiman, however, did not analyze his data in terms of costs and benefits for other reasons.) In effect, they could have used these sentence cues as neutral stimuli, at least relative to effects of specific expectations.

This strategy seems appropriate where the investigator can identify a specific attribute of sentences (e.g., degree to which they predict the target word), and can manipulate the presence or absence of this attribute in features of the sentence stimuli. Thus, instead of focusing on sentential context effects per se as in early studies, research might instead focus on the effect of having specific expectations, as in more recent studies. In this case, neutral stimuli could consist of sentences that are syntactically and semantically congruous with the sentence but not predictable (as in Kleiman, 1980). What this proposal requires is a more theoretically motivated analysis focused on the specific effects one wishes to study. This strategy assumes, of course, that one does not incur the other potential problems we discussed, such as the redundancy of the neutral cue (which would apply to a neutral cue that consisted of the same sentence frame on all neutral trials).

A similar strategy can be used for studies of single-word recognition. Investigators can try to use words as neutral cues as well as using them as non-neutral priming cues. Merrill, Sperber, and McCauley (1981), for example, compared Stroop naming performance for target words that shared compatible (congruous), incompatible (incongruous), or irrelevant features with a prior priming stimulus word. This last condition was treated as a neutral condition. As in the sentence cases, the strategy of trying to equate neutral and non-neutral priming cues in terms of physical properties assumes that one can identify specific processes (whose costs and benefits one is interested in assessing) and control their presence or absence in words.

A related but alternative approach would be to have subjects respond to the neutral cue in a way that made its processing demands more comparable with the non-neutral cue, at least with respect to the mechanisms in which one is interested. Durgunoglu and Neely (1982) suggested an example in the case of word recognition studies. Subjects might be asked to make a discriminative response to a stimulus other than the target, and to use the neutral cue to prepare for the stimulus. This assumes that the mechanisms for which one wishes to examine costs and benefits can be controlled by conscious strategies, and that the relative processing demands can indeed be equated.

Finally, the effectiveness of the matching process ought to be properly evaluated as well, as Becker (1980) tried to do, for example. In this way, there may be more confidence in the assertion that neutral and informative cues differ only in the extent of their elicitation of specific preparatory processes.

The second major caution investigators can exercise is to include experimental conditions that may provide some converging measures of the conclusions to be drawn from the comparison of neutral and informative cues. Of course, the choice of conditions to meet this goal hinges on the particular model being tested by the experimental study. In the case of the two-component model of Posner and Snyder (1975), various converging operations have been suggested and applied. These tend to focus on alternative measures of attentional involvement (or its lack thereof). They include variation in cue validity (Jonides, 1980; Taylor, 1977; Tweedy et al., 1977), variation in instructions to attend or ignore the cue (Jonides, 1981; Taylor, 1977), variation in strategy used in processing the cue (Smith, 1979), use of multiple neutral cues (McClelland & O'Regan, 1981; but see Rayner & Slowiaczek, 1981), and measurement of performance in a secondary dual task (Becker, 1976; Jonides, 1981; McLean & Shulman, 1978). Clearly, for this model there are various candidate measures that one could accumulate. Careful thought might reveal measures for other models as well. (See, e.g., Eisenberg & Becker, 1982, for converging evidence of specific semantic strategies for word recognition in reading and performance in a lexical decision task.)

Chronometric methods have obviously assumed a place of prominence in the repertoire of the cognitive psychologist. We recognize that the benefit of cost–benefit analysis is that it has expanded the scope of mental chronometry in the study of preparatory effects. However, we also recognize that this benefit has not accrued without cost, a cost that could be minimized with more careful application of the method.

References

Antos, S. (1979). Processing facilitation in a lexical decision task. *Journal of Experimental Psychology: Human Perception and Performance, 5,* 527–545.

Balota, D. A. (1983). Automatic semantic activation and episodic memory encoding. *Journal of Verbal Learning and Verbal Behavior, 22,* 88–104.

Becker, C. (1976). Allocation of attention during visual word recognition. *Journal of Experimental Psychology: Human Perception and Performance, 2,* 556–566.

Becker, C. (1980). Semantic context effects in visual word recognition: An analysis of semantic strategies. *Memory and Cognition, 8,* 493–512.

Beller, H. (1971). Priming: Effects of advance information on matching. *Journal of Experimental Psychology, 87,* 176–182.

Blaxton, T. A., & Neely, J. H. (1983). Inhibition from semantically related primes: Evidence of a category-specific inhibition. *Memory and Cognition, 11,* 500–510.

Carroll, M., & Kirsner, K. (1982). Context and repetition effects in lexical decision and recognition memory. *Journal of Verbal Learning and Verbal Behavior, 21,* 55–69.

de Groot, A. M. B., Thomassen, A. J. W. M., & Hudson, P. T. W. (1982). Associative facilitation of word recognition as measured from a neutral prime. *Memory and Cognition, 10,* 358–370.

Donders, F. C. (1969). On the speed of mental processes. *Acta Psychologica, 30,* 412–431.

Durgunoglu, A., & Neely, J. H. (1982, November). *Episodic and semantic priming in episodic and semantic memory.* Paper presented at the meeting of the Psychonomic Society, Minneapolis, MN.

Eisenberg, P., & Becker, C. A. (1982). Semantic context effects in visual word recognition sentence processing and reading: Evidence for semantic strategies. *Journal of Experimental Psychology: Human Perception and Performance, 8,* 739–756.

Favreau, M., & Segalowitz, N. (1983). Automatic and controlled processes in the first and second language reading of fluent bilinguals. *Memory and Cognition, 11,* 565–574.

Fischler, I. (1977). Associative facilitation without expectancy in a lexical decision task. *Journal of Experimental Psychology: Human Perception and Performance, 3,* 18–26.

Fischler, I., & Bloom, P. (1979). Automatic and attentional processes in the effects of sentence contexts on word recognition. *Journal of Verbal Learning and Verbal Behavior, 18,* 1–20.

Fischler, I., & Bloom, P. (1980). Rapid processing of the meaning of sentences. *Memory and Cognition, 8,* 216–225.

Forster, K. I. (1981). Priming and the effects of sentence and lexical contexts on naming time: Evidence for autonomous lexical processing. *Quarterly Journal of Experimental Psychology, 33,* 465–495.

Goodman, G. O., McClelland, J. L., & Gibbs, R. W., Jr. (1981). The role of syntactic context in word recognition. *Memory and Cognition, 9,* 580–586.

Irwin, D., Bock, K., & Stanovich, K. (1982). Effects of information structure cues on visual word processing. *Journal of Verbal Learning and Verbal Behavior, 21,* 307–325.

Jonides, J. (1980). Towards a model of the mind's eye's movement. *Canadian Journal of Psychology, 2,* 103–112.

Jonides, J. (1981). *Voluntary versus automatic control over the mind's eye's movement.* In J. B. Long & A. D. Baddeley (Eds.), *Attention and performance* (Vol. 9, pp. 187–204). Hillsdale, NJ: Erlbaum.

Jonides, J. (1983). Further toward a model of the mind's eye's movement. *Bulletin of the Psychonomic Society, 21,* 247–250.

Kleiman, G. (1980). Sentence frame contexts and lexical decisions: Sentence acceptability and word relatedness effects. *Memory and Cognition, 8,* 336–344.

Koriat, A. (1981). Semantic facilitation in lexical decision as a function of prime-target association. *Memory and Cognition, 9,* 587–598.

Kraut, A., & Smothergill, D. A. (1978). A two-factor theory of stimulus-repetition effects. *Journal of Experimental Psychology: Human Perception and Performance, 4,* 191–197.

Kraut, A., Smothergill, D., and Farkas, M. (1981). Stimulus repetition effects on attention to words and colors. *Journal of Experimental Psychology: Human Perception and Performance, 7,* 1303–1311.

Larochelle, S., McClelland, J., & Rodriquez, E. (1980). Context and the allocation of resources in word recognition. *Journal of Experimental Psychology: Human Perception and Performance, 6,* 686–694.

Leonard, J. A. (1958). Partial advance information in a choice reaction task. *British Journal of Psychology, 48,* 89–96.

Mack, R. (1981). *A comparison of spreading activation and attention mechanisms for sentence context effects* (Research Report RC 9056). Yorktown Heights, NY: IBM Thomas J. Watson Research Center.

McClelland, J., & O'Regan, J. (1981). Expectations increase the benefit derived from parafoveal visual information in reading words aloud. *Journal of Experimental Psychology: Human Perception and Performance, 7,* 634–644.

McLean, J., & Shulman, G. (1978). On the construction and maintenance of expectancies. *Quarterly Journal of Experimental Psychology, 30,* 441–454.

Merrill, E., Sperber, R., & McCauley, C. (1981). Differences in semantic encoding as a function of reading comprehension skill. *Memory and Cognition, 9,* 618–624.

Meyer, D., & Schvaneveldt, R. (1971). Facilitation in recognizing pairs of words: Evidence of a dependence between retrieval operations. *Journal of Experimental Psychology, 90,* 227–234.

Meyer, D., Schvaneveldt, R., & Ruddy, M. (1975). Loci of contextual effects on visual word recognition. In P. Rabbit & S. Dornic (Eds.), *Attention and performance* (Vol. 5, pp. 98–118). New York: Academic Press.

Neely, J. (1976). Semantic priming and retrieval from lexical memory: Evidence for facilitory and inhibitory processes. *Memory and Cognition, 4,* 648–654.

Neely, J. (1977). Semantic priming and retrieval from lexical memory: The roles of inhibitionless spreading activation and limited capacity attention. *Journal of Experimental Psychology: General, 106,* 226–254.

Neely, J., Schmidt, S., & Roediger, H., III. (1983). Inhibition from related primes in recognition memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 9,* 196–211.

Pachella, R. G. (1974). The interpretation of reaction time in information processing research. In B. Kantowitz (Ed.), *Human information processing: Tutorials in performance and cognition* (pp. 41–82). Potomac, MD: Erlbaum.

Posner, M., & Boies, S. (1971). Components of attention. *Psychological Review, 78,* 391–408.

Posner, M. I., Klein, R., Summers, J., & Buggie, S. (1973). On the selection of signals. *Memory and Cognition, 1,* 2–12.

Posner, M., Nissen, M., & Ogden, W. (1978). Attended and unattended processing modes: The role of set for spatial location. In A. Pick & E. Saltzman (Eds.), *Modes of perceiving and processing of information* (pp. 137–158). Hillsdale, NJ: Erlbaum.

Posner, M., & Snyder, C. (1975). Attention and cognitive control. In R. L. Solso (Ed.), *Information processing and cognition: The Loyola symposium* (pp. 55–86). Hillsdale, NJ: Erlbaum.

Raynor, K., & Slowiaczek, M. (1981). Expectations and parafoveal information in reading: Comments on McClelland and O'Regan. *Journal of Experimental Psychology: Human Perception and Performance, 7,* 645–651.

Roediger, H. L., III, Neely, J. H., & Blaxton, T. A. (1983). Inhibition from related primes in semantic memory retrieval: A reappraisal of Brown's (1979) paradigm. *Journal of Experimental Psychology: Human Learning, Memory, and Cognition, 9,* 478–485.

Schuberth, R., & Eimas, P. (1977). Effects of context on the classification of words and non-words. *Journal of Experimental Psychology: Human Perception and Performance, 3,* 27–36.

Schuberth, R. E., Spoehr, K. T., & Lane, D. M. (1981). Effects of stimulus and contextual information on the lexical decision process. *Memory and Cognition, 9,* 68–77.

Simpson, G. B., & Lorsbach, T. C. (1983). The development of automatic and conscious components of contextual facilitation. *Child Development, 54,* 760–772.

Smith, M. C. (1979). Contextual facilitation in a letter search task depends on how the prime is processed. *Journal of Experimental Psychology: Human Perception and Performance, 5,* 239–251.

Stanovich, K., & West, R. (1979). Mechanisms of sentence context effects in reading: Automatic activation and conscious attention. *Memory and Cognition, 7,* 77–85.

Stanovich, K., & West, R. (1981). The effect of sentence context on ongoing word recognition: Tests of a two-process theory. *Journal of Experimental Psychology: Human Perception and Performance, 7,* 658–672.

Stanovich, K. E., & West, R. F. (1983). The generalizability of context effects on word recognition: A reconsideration of the roles of parafoveal priming and sentence context. *Memory and Cognition, 11,* 49–58.

Sternberg, S. (1969). The discovery of processing stages: Extensions of Donder's method. In W. Koster (Ed.), *Attention and performance* (*Vol. 2,* pp. 276–315). North-Holland: Dordrecht.

Sudevan, P. (1984). *Chronometric analysis of preparation.* Unpublished manuscript.

Taylor, D. A. (1977). Time course of context effects. *Journal of Experimental Psychology: General, 106,* 404–426.

Tweedy, J., Lapinski, R., & Schvaneveldt, R. (1977). Semantic context effects on word recognition: Influence of varying the proportion of items presented in an appropriate context. *Memory and Cognition, 5,* 84–89.

West, R., & Stanovich, K. (1978). Automatic contextual facilitation in readers of three ages. *Child Development, 49,* 717–727.

West, R. F., & Stanovich, D. E. (1982). Source of inhibition in experiments on the effect of sentence context on word recognition. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 8,* 385–399.

Received December 13, 1982
Revision received February 10, 1984 ■