



ELSEVIER

Acta Psychologica 90 (1995) 163–190

**acta
psychologica**

Adaptive executive control: Flexible multiple-task performance without pervasive immutable response-selection bottlenecks

David E. Meyer ^{a,*}, David E. Kieras ^b, Erick Lauber ^a,
Eric H. Schumacher ^a, Jennifer Glass ^a, Eileen Zurbriggen ^a,
Leon Gmeindl ^a, Dana Apfelblat ^a

^a *Department of Psychology, University of Michigan, 525 East University, Ann Arbor, MI 48109-1109, USA*

^b *Department of Electrical Engineering and Computer Science, University of Michigan, 1101 Beal Avenue, Ann Arbor, MI, 48109-2110, USA*

Abstract

A new theoretical framework, the EPIC (Executive-Process/Interactive-Control) architecture, provides the basis for accurate detailed computational models of human multiple-task performance. Contrary to the traditional response-selection bottleneck hypothesis, EPIC's cognitive processor can select responses and do other procedural operations simultaneously for multiple concurrent tasks. Using this capacity together with flexible executive control of peripheral perceptual-motor components, EPIC computational models account well for various patterns of mean reaction times, systematic individual differences in multiple-task performance, and influences of special training on people's task-coordination strategies. These diverse phenomena, and EPIC's success at modeling them, raise strong doubts about the existence of a pervasive immutable response-selection bottleneck in the human information-processing system. The present research therefore helps further characterize the nature of discrete versus continuous information processing.

1. Introduction

Traditionally, it has been hypothesized that a pervasive immutable cognitive response-selection bottleneck (RSB) exists in the human information-processing system (for comprehensive reviews, see Meyer and Kieras, 1995; Pashler, 1994). According to the RSB hypothesis, there is a stage of processing that selects

* Corresponding author. demeyer@umich.edu.

responses to stimuli for various tasks, and that only has enough capacity to accommodate one stimulus at a time. Thus, if two tasks (e.g., saying words in response to auditory tones, and pressing keys in response to visual letters) must be performed simultaneously, then even when different stimulus and response modalities are involved, the selection of a response for one task will supposedly have to wait until response selection for the other is completed. Theorists have typically attributed such delays to permanent "hardware" characteristics of central limited-capacity decision mechanisms (Craik, 1948; Davis, 1957; De Jong, 1993; McCann and Johnston, 1992; Pashler, 1984,1989,1990,1993; Pashler and Johnston, 1989; Ruthruff et al., 1995; Schweickert et al., 1992; Smith, 1967; Van Selst and Jolicoeur, 1993; Vince, 1949; Welford, 1952,1959,1967).

Assuming that it is valid, the RSB hypothesis has major implications for conceptualizing discrete and continuous information processing. These implications concern not only multiple-task but also single-task performance. Given a pervasive response-selection bottleneck, many individual tasks could appear to be performed through discrete processing stages that accomplish perceptual encoding, response selection, and movement production in a strict step-by-step fashion (cf. Luce, 1986; Meyer et al., 1988; Miller, 1988; Sanders, 1980; Sternberg, 1969; Townsend and Ashby, 1983). For example, suppose a task involves stimuli that have several orthogonal perceptual dimensions (e.g., shapes and sizes of visual objects). Also, suppose stimulus values on each perceptual dimension (e.g., round and square for shape; small and large for size) must be converted to values on other response dimensions (e.g., left and right hands; index and middle fingers). Then making responses to these stimuli may necessarily entail discrete substages, if there is a response-selection bottleneck. In particular, the start of selecting a correct response finger on the basis of an output from a relatively slow perceptual size-identification process might have to wait until the selection of a correct response hand on the basis of a faster shape-identification process has been completed (cf. Miller, 1982; Osman et al., 1992). Essentially, a response-selection bottleneck could limit the extent to which selections of response values along different dimensions can occur simultaneously.

Analyses of discrete and continuous information processing should therefore take the RSB hypothesis very seriously. Depending on whether or not there is a pervasive response-selection bottleneck, more or less constraint would be placed on what mental operations may temporally overlap each other and exploit their partial outputs. With these stipulations in mind, the remainder of this article has five parts. First, we survey some attractive features of the traditional RSB hypothesis. Second, strong arguments are made against taking the validity of this hypothesis for granted. Third, an alternative theoretical framework (Kieras and Meyer, 1994; Meyer and Kieras, 1992,1994,1995) is described, in which some components of the human information-processing system have substantially more capacity and flexibility than the RSB hypothesis allows. Fourth, using our framework, precise computational models are applied to account for results from representative studies of multiple-task performance. Fifth, we consider what the present theory implies about future research on discrete versus continuous information processing

in a Brave New World without pervasive immutable response-selection bottlenecks.

2. Attractive features of the RSB hypothesis

Several features of the traditional RSB hypothesis have made it especially attractive. Among them are conceptual simplicity, predictive precision, intuitive appeal, and authoritative endorsement. Compared to alternatives such as general capacity theory (Moray, 1967; Kahneman, 1973; McLeod, 1978; Gottsdanker, 1980) and multiple-resource theory (Navon and Gopher, 1979; Wickens, 1984), the RSB hypothesis involves relatively few assumptions and yields more precise quantitative predictions about certain aspects of multiple-task performance. These assumptions and predictions are consistent with some expert intuitions about the nature of human attention and the mental processes that mediate it. As William James (1890, p. 409) rhetorically remarked: “how many ideas or things can we attend to at once ... the answer is, not easily more than one ... there must be a rapid oscillation of the mind from one idea to the next, with no consequent gain of time.” When translated to modern information-processing concepts, James’ answer is what might be expected in terms of a pervasive response-selection bottleneck (Norman, 1969).

2.1. *Psychological refractory-period procedure*

Moreover, extensive behavioral evidence has been obtained to support the RSB hypothesis. A major source of this evidence is the psychological refractory-period procedure (Bertelson, 1966; Kantowitz, 1974; Pashler, 1994; Smith, 1967; Welford, 1967). On each trial of the PRP procedure, a warning signal is followed by a stimulus for the first of two tasks. In response to it, a subject must react quickly and accurately. Soon after the Task 1 stimulus, there is a stimulus for the second task. The time between the two stimuli is the stimulus-onset asynchrony (SOA). In response to the Task 2 stimulus, the subject must again react quickly and accurately. However, instructions for the PRP procedure typically state that regardless of the SOA, Task 1 should have higher priority than Task 2; subjects may also be required to make the Task 1 response first (e.g., Pashler, 1984; Pashler and Johnston, 1989). RTs are then measured to determine how much Task 1 interferes with Task 2.

2.2. *Representative PRP results*

Some idealized representative results from the PRP procedure appear in Fig. 1. Here mean RTs for both Tasks 1 and 2 are plotted versus the SOA as a function of Task 2 response-selection difficulty (easy or hard), which can be varied by manipulating factors like S–R compatibility and S–R numerosity in Task 2 (Becker, 1976; Hawkins et al., 1979; Karlin and Kestenbaum, 1968; McCann and Johnston, 1992).

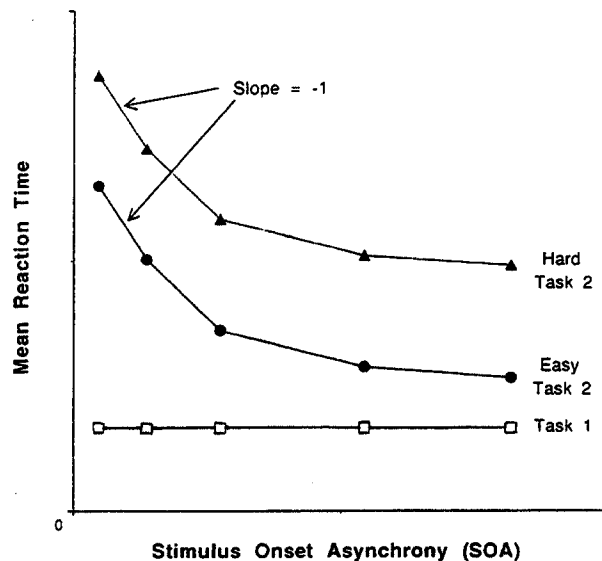


Fig. 1. Idealized representative results from the PRP procedure.

As Fig. 1 indicates, such manipulations may produce patterns of data that are consistent with a response-selection bottleneck.

Constant mean Task 1 RTs

A first salient aspect of Fig. 1 is that mean Task 1 RTs are affected by neither the SOA nor Task 2 difficulty. This conforms nicely to the RSB hypothesis. If Task 1 stimuli come first at each SOA, then presumably they always enter the response-selection bottleneck before Task 2 stimuli, and the bottleneck's all-or-none "admissions policy" precludes Task 2 from competing with Task 1 for the limited processing capacity therein. Satisfying this implication, many PRP studies have yielded essentially constant mean Task 1 RTs (Davis, 1957; Hawkins et al., 1979; Karlin and Kestenbaum, 1968; McCann and Johnston, 1992; Pashler, 1990; Pashler and Johnston, 1989; Welford, 1959).

Elevated Task 2 RTs at short SOAs

Second, the mean Task 2 RTs are higher at short SOAs in Fig. 1. This elevation, called the *PRP effect*, follows directly from the RSB hypothesis. When the SOA is short, response selection for Task 2 supposedly waits while the Task 1 stimulus passes through the bottleneck and is converted to a Task 1 response. As a result, the concomitant extra delay should raise mean Task 2 RTs above their single-task baseline. Virtually all past studies with the PRP procedure have yielded such increases (Pashler, 1994).

PRP curves with -1 slopes

In Fig. 1, however, mean Task 2 RTs gradually decrease as the SOA increases, forming so-called *PRP curves*. Specifically, under the RSB hypothesis, the slopes of these curves should equal -1 at short SOAs. This is because each increment of

the SOA, starting from zero, produces an equal opposite decrement in how long a Task 2 stimulus must wait to enter the response-selection bottleneck. Subsequently, once the SOA becomes long enough to yield some trials with no delay in Task 2, the PRP curves should become correspondingly shallower, bottoming out with a slope of zero at very long SOAs. Satisfying these additional expectations, PRP studies have typically manifested curves whose slopes range from -1 to 0 (e.g., Davis, 1957; Hawkins et al., 1979; Karlin and Kestenbaum, 1968; McCann and Johnston, 1992; Pashler, 1990; Pashler and Johnston, 1989; Welford, 1959).

Parallel PRP curves

Another quantitative property implied by the RSB hypothesis is that the SOA and response-selection difficulty for Task 2 should affect mean Task 2 RTs additively. Thus, as in Fig. 1, PRP curves ought to be “parallel” (i.e., vertically equidistant) when the difficulty of Task 2 response selection varies across conditions. This is because regardless of the SOA, each Task 2 RT depends directly on how long response selection takes for Task 2 after the Task 1 stimulus has passed through the bottleneck (Karlin and Kestenbaum, 1968; Keele, 1973; McCann and Johnston, 1992; Pashler, 1984; Pashler and Johnston, 1989; Schvaneveldt, 1969). Such additivity and parallelism have been obtained in several PRP studies, even when Tasks 1 and 2 involve neither the same stimulus nor response modalities (Becker, 1976; McCann and Johnston, 1992; Pashler, 1984; Pashler and Johnston, 1989; Ruthruff et al., 1995; Van Selst and Jolicoeur, 1993). Of course, this would be expected if between-task interference occurs at a central cognitive level rather than merely at peripheral perceptual-motor levels.

Other relevant findings

PRP effects are also robust in other respects. Typically, they persist with extended practice. Even after thousands of practice trials on the PRP procedure, subjects still produce elevated Task 2 RTs at short SOAs (Gottsdanker and Stelmach, 1971; Karlin and Kestenbaum, 1968). Again, this is expected from the RSB hypothesis, which assumes that the response-selection bottleneck is an immutable “hardware” component of the human information-processing system.

3. Arguments against the RSB hypothesis

Nevertheless, various arguments can be made against the traditional RSB hypothesis. For example, its assumptions seem neurophysiologically implausible. Contrary to them, information processing in the brain is “massively parallel” and “distributed” throughout components of many interconnected neural networks (Anderson and Hinton, 1981; Rumelhart and McClelland, 1986). There are no obvious brain sites that constitute immutable response-selection bottlenecks of the sort to which PRP effects and other multiple-task performance decrements have been attributed (Allport, 1980, 1987; Neumann, 1987).

A second related concern is that the RSB hypothesis lacks computational flexibility. It provides little accommodation for executive control processes that allocate available system resources efficiently and adaptively to different on-going tasks. Task scheduling through an immutable response-selection bottleneck has been assumed to happen simply on a first-come first-serve basis, whereby secondary tasks are completely blocked from access to essential resources during extended periods of time. Any performer who suffers from such rigidity would have great difficulty adapting successfully to major changes in task priorities and increased or decreased knowledge about impending environmental stimuli (Allport, 1980,1987; Neumann, 1987).

Because of its unrealistic limitations, the RSB hypothesis seems inconsistent with results from many multiple-task situations (Wickens, 1984). An immutable response-selection bottleneck does not even account fully for data from the standard PRP procedure. Instead, it has become gradually but compellingly evident that subjects can and do produce systematic patterns of RTs different than those in Fig. 1, extending well beyond the scope of the RSB hypothesis (Meyer and Kieras, 1992,1994,1995).

3.1. *Divergent PRP curves*

One observed RT pattern for which the RSB hypothesis cannot account very well involves divergent PRP curves. Such divergence occurs when the difficulty of selecting Task 2 responses affects mean Task 2 RTs less at short SOAs than at long SOAs, yielding a positive SOA-by-difficulty interaction. For example, consider the left panel of Fig. 2, which shows mean Task 2 RTs (solid curves) from a PRP study by Hawkins et al. (1979), who manipulated response-selection difficulty by varying the number of S–R pairs in Task 2. Here the Task 2 difficulty effect is only about 25 msec at the shortest SOA, whereas it is nearly 200 msec at the longest SOA.¹ Reliable positive interactions like this have also been reported by several other investigators (e.g., Ivry et al., 1994,1995; Karlin and Kestenbaum, 1968; Lauber et al., 1994).

A plausible interpretation of these results is that: (1) at short SOAs, response selection for Task 2 occurs independently and simultaneously with response selection for Task 1; (2) progress on Task 2 pauses temporarily before initiation of its response movement, letting Task 1 finish first, as required by instructions for the PRP procedure; and (3) the slack introduced by this pause absorbs the effects of response-selection difficulty on Task 2 when the SOA is short. Keele (1973) and others have discussed how the latter sequence of events could yield the type of positive interaction in Fig. 2 (left panel), whereas temporally separate response-selection stages for Tasks 1 and 2 would not. Consequently, the absence of strict additivity between the effects of SOA and Task 2 response-selection difficulty

¹ Mean Task 1 RTs equalled slightly more than 600 ms and were not affected very much by either the SOA or Task 2 response-selection difficulty (Hawkins et al., 1979).




Fig. 2. Left panel: divergent PRP curves obtained by Hawkins et al. (1979) with an auditory-manual Task 1 and visual-manual Task 2. The solid functions represent mean Task 2 RTs observed when response selection for Task 2 was “easy” (large circles) or “hard” (large triangles). Each observed mean RT has a standard error of approximately 10 msec. Dashed functions (small circles and triangles) represent simulated mean Task 2 RTs from the SRD model to be discussed later (Fig. 7). Right panel: convergent PRP curves obtained by Ivry et al. (1994, 1995, Exp. 2) with a visual-manual Task 1 and visual-manual Task 2. Mean Task 2 RTs plus-or-minus one standard error (based on the SOA-by-difficulty-by-subject interaction) are shown.

raises compelling doubts about the existence of a pervasive immutable response-selection bottleneck.

3.2. Convergent PRP curves

Another complementary RT pattern for which the RSB hypothesis cannot account very well involves convergent PRP curves. Such convergence occurs when the difficulty of selecting Task 2 responses affects mean Task 2 RTs more at short SOAs than at long SOAs, yielding a negative SOA-by-difficulty interaction. Several cases like this have been reported recently (e.g., Ivry et al., 1994, 1995; Lauber et al., 1994).² For example, consider the right panel of Fig. 2. Here mean Task 2 RTs are plotted from a PRP study by Ivry et al. (1994, 1995, Exp. 2), who manipulated response-selection difficulty by varying the spatial S–R compatibility in Task 2. At the shortest SOA, the difficulty effect on mean Task 2 RT is nearly 300 msec, whereas at the longest SOA, it is less than 200 msec, forming a substantial negative SOA-by-difficulty interaction. Given this pattern, Ivry et al. (1995) attributed their results to “resource sharing strategies.” In contrast, the RSB hypothesis offers no simple satisfying answers for why, across various experiments, PRP curves sometimes converge or diverge as a function of SOA and Task 2 response-selection difficulty.

² Converging PRP curves may occur even when there are very little effects of SOA and/or Task 2 response-selection difficulty on mean Task 1 RTs (Lauber et al., 1994).

3.3. Slopes steeper than -1

The plausibility of a pervasive immutable response-selection bottleneck is likewise reduced by carefully examining the slopes that PRP curves sometimes have. For example, consider the left panel of Fig. 3, which shows mean Task 2 RTs from a study by Lauber et al. (1994, Exp. 2). These data were obtained under conditions similar to those of Hawkins et al. (1979), except that the present Task 1 had more S–R pairs (viz., four instead of two). This change yielded “parallel” (i.e., vertically equidistant) average PRP curves with approximately additive effects of SOA and Task 2 difficulty, as the RSB hypothesis predicts. However, over their two shortest SOAs, Lauber et al. (1994, Exp. 2) found that the PRP curves in Fig. 3 had extremely negative slopes (almost -1.4 on average) that were reliably steeper than -1 . Such extreme steepness was also found by Hawkins et al. (1979) in some of their conditions. Why and how might this happen? In reply, the RSB hypothesis again has no ready answer. As mentioned before, it implies that the slopes of PRP curves should be -1 or shallower (cf. Fig. 1).

3.4. Systematic individual differences

Other implications of the RSB hypothesis may be refuted by comparing PRP curves from different subjects. If everyone has an immutable response-selection

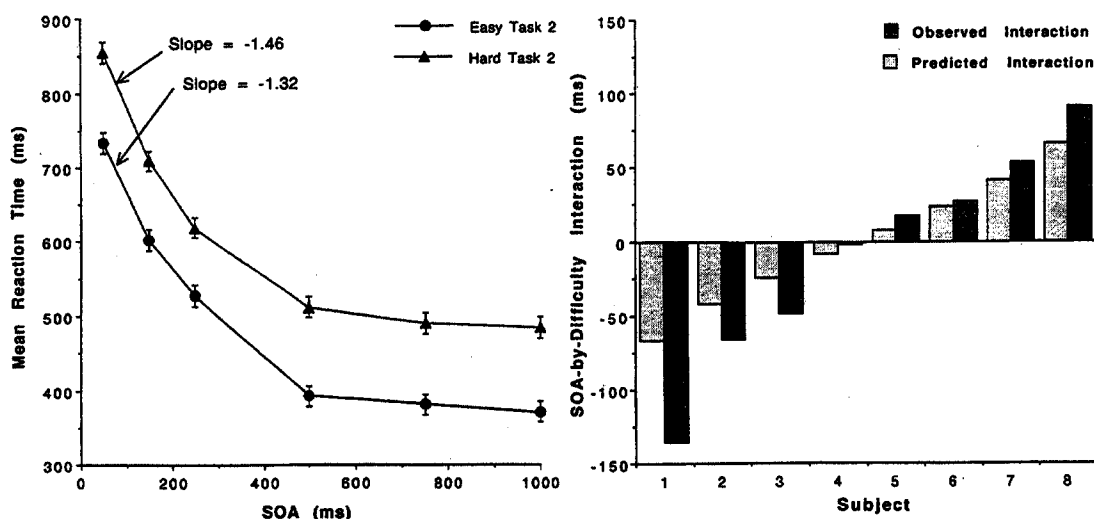


Fig. 3. Left panel: “parallel” (i.e., vertically equidistant) average PRP curves obtained by Lauber et al. (1994, Exp. 2) with an auditory-manual Task 1 and visual-manual Task 2. Mean Task 2 RTs plus-or-minus one standard error (based on the SOA-by-difficulty-by-subject interaction) are shown. Right panel: observed and predicted interactions between effects of SOA and response-selection difficulty on mean Task 2 RTs for eight subjects whose average PRP curves appear in the left panel. Nearly all of the dark vertical bars (observed interactions) are more extreme (RMS error = 29 ms; $p < 0.05$) than the light vertical bars (predicted interactions), which come from the RSB hypothesis.

bottleneck, then each subject in an experiment should produce the same qualitative pattern of mean Task 2 RTs. Nevertheless, occasional checks for such homogeneity have instead revealed striking systematic individual differences (e.g., Ivry et al., 1994,1995; Lauber et al., 1994).

For example, consider the right panel of Fig. 3. Here we have plotted observed interactions between the effects of SOA and response-selection difficulty on mean Task 2 RTs for each of eight subjects who contributed to the average PRP curves of Fig. 3's left panel. Across the horizontal axis of Fig. 3's right panel, these subjects are ordered according to the magnitudes and signs of their SOA-by-difficulty interactions. On the vertical axis, a zero interaction indicates that a subject produced equal Task 2 difficulty effects at the shortest and longest SOAs, consistent with "parallel" PRP curves. A positive interaction indicates that the subject's PRP curves diverged as the SOA increased, and a negative interaction indicates that they converged. The dark vertical bars show how positive or negative each subject's interaction was. Three subjects had marked negative interactions (dark vertical bars extending downward) and convergent PRP curves. One subject had a near-zero interaction and virtually "parallel" PRP curves. Four other subjects had various degrees of positive interaction (dark vertical bars extending upward) and divergent PRP curves.

In contrast, the light vertical bars of Fig. 3 (right panel) represent what the RSB hypothesis predicts for a sample of eight such subjects. These predictions were obtained by assuming that each subject belongs to a homogeneous population whose members all have theoretically additive effects of SOA and Task 2 response-selection difficulty. With this assumption, we estimated the distribution from which the subjects' observed interactions should come, given how much between-trial variance there was in each subject's data. Thus, if every subject had a response-selection bottleneck, then the light bars ought to match the dark bars closely.

However, this expected equivalence did not occur. A large majority (i.e., 7/8) of the dark vertical bars in Fig. 3 (right panel) are longer than the light bars paired with them, revealing interactions consistently more extreme than the RSB hypothesis predicts. Our results instead suggest that there are two (or more) distinct subgroups of subjects, including some (e.g., Subject 1) who produce significantly convergent PRP curves and others (e.g., Subject 8) who produce significantly divergent PRP curves. Such individual differences reinforce two conclusions: people do not have immutable response-selection bottlenecks; other mechanisms – whose parameters depend on personal predilections – are the source of observed PRP effects.

3.5. *Effects of special training*

Finally, consistent with the preceding conclusions, some studies have revealed that subjects' PRP curves can be modified predictably through special types of training (e.g., Koch, 1993,1994; Lauber et al., 1994; Sanders, 1964). Such results

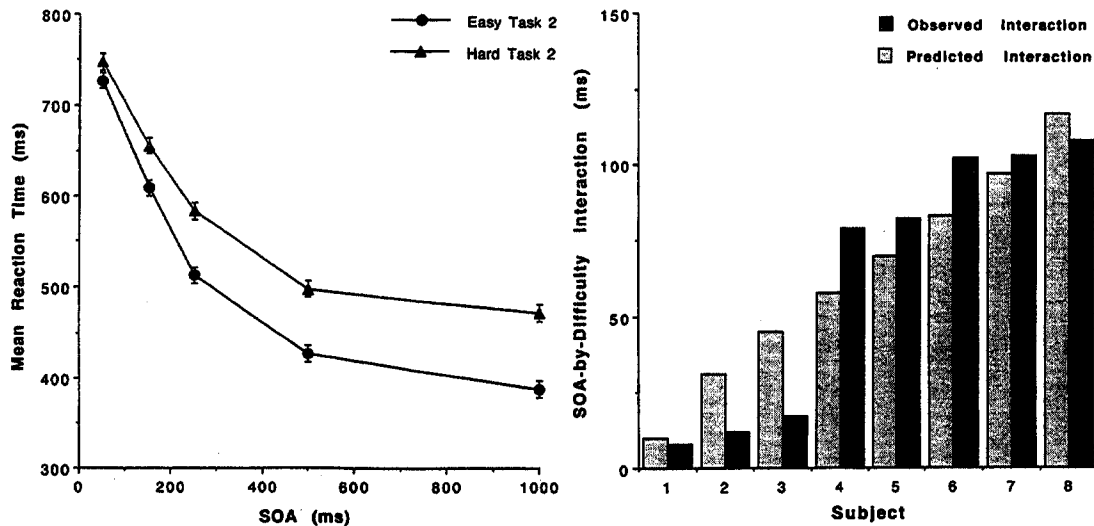


Fig. 4. Left panel: divergent average PRP curves obtained by Lauber et al. (1994, Exp. 3) after subjects received special preliminary training that encouraged concurrent response-selection for an auditory-manual Task 1 and visual-manual Task 2. Mean Task 2 RTs plus-or-minus one standard error (based on the SOA-by-difficulty-by-subject interaction) are shown. Right panel: observed and predicted positive interactions (RMS error = 17 ms) between effects of SOA and response-selection difficulty on mean Task 2 RTs for eight subjects who contributed to the average PRP curves in the left panel. The predicted interactions (light vertical bars) assume that these subjects belong to a single homogeneous population whose members produce different amounts of observed positive interaction (dark vertical bars) only because of inherent between-trial variability in their RTs.

affirm that whatever the source of the PRP effect, it is certainly not “immutable.” For example, during another study by Lauber et al. (1994, Exp. 3), additional subjects were tested under the same PRP procedure that yielded the data in Fig. 3. However, they received special preliminary training before the PRP procedure began. This training, which followed the suggestions by Gopher (1993) about how to optimize multiple-task performance, required concurrent auditory-manual and visual-manual tasks to be performed as quickly as possible with equally high priority and relaxed constraints on the order of the tasks’ stimuli and responses. As a result, subjects were strongly encouraged to overlap their response-selection processes for the two tasks. After this training finished, subjects then entered the standard PRP procedure.

Fig. 4 (left panel) shows average Task 2 PRP curves that Lauber et al. (1994, Exp. 3) thereby obtained. Unlike before (cf. Fig. 3, left panel), these new curves diverge substantially. At the shortest SOA, Task 2 response-selection difficulty has relatively little effect on mean Task 2 RTs, whereas at the longest SOA, there is still a substantial difficulty effect. Furthermore, during the PRP procedure, all of the subjects who received special training produced some positive interaction between the effects of SOA and Task 2 difficulty (Fig. 4, right panel). This latter outcome, combined with other prior ones (Figs. 2 and 3), seems rather problematic for the RSB hypothesis.

4. A new theoretical framework

If a pervasive immutable response-selection bottleneck does not mediate the PRP effect, then what is the effect's true source? As hinted already, an answer may be found in the instructions for the standard PRP procedure (Meyer and Kieras, 1992,1994,1995; Koch, 1993,1994). They typically request that Task 1 receive absolute priority. For example, in Pashler and Johnston's (1989) study, subjects were told that they "*should respond as rapidly as possible to the first stimulus,*" and "*the experimenter emphasized the importance of making the first response as promptly as possible.*" Similarly, in a study by Pashler (1984, Exp. 1, p. 365), subjects were instructed that "*the first stimulus must be responded to before the second.*" Because of such constraints, people may postpone completing Task 2 at short SOAs even though they have the potential capacity to perform concurrent tasks with no between-task interference. To satisfy the PRP procedure's instructions, perhaps optional temporary bottlenecks are programmed at one or more stages of processing for Task 2, deferring Task 2 responses until Task 1 has finished. If so, then the magnitudes of PRP effects and the forms of PRP curves may be under strategic control, and this could account for many of the phenomena (e.g., Fig. 2 through 4) that seem antithetical to the traditional RSB hypothesis. Given these possibilities, we have therefore begun to develop a new theoretical framework for describing and predicting human multiple-task performance through detailed executive computational models (Kieras and Meyer, 1994; Meyer and Kieras, 1992,1994,1995).

4.1. Basic assumptions

The first basic assumption of our framework is that in some respects, people have substantial cognitive capacity for performing multiple concurrent tasks. More precisely, we assume that various task procedures can be executed simultaneously with distinct sets of production rules (cf. Anderson, 1976,1983; Newell, 1973a). For example, while driving a car, a person may also be able to talk on a cellular telephone because the production rules used respectively for driving and conversing are distinct and applied in parallel. According to the present framework, there is no immutable decision or response-selection bottleneck for executing task procedures at a central cognitive level.

Instead, we attribute decrements in multiple-task performance to other sources such as limited peripheral sensory and motor mechanisms, which cause "structural interference" (cf. Kahneman, 1973). For example, while making phone calls in a car, most drivers cannot keep their eyes simultaneously on the phone dial and the road, nor can they keep both hands on the steering wheel and hold the telephone. Perhaps it is these sensory-motor constraints – not limited cognitive capacity – that restrict people's ability to drive safely and make phone calls at the same time.

We also assume that conflicts caused by sensory-motor constraints can be alleviated by properly scheduling the tasks at hand. In particular, concurrent tasks may be scheduled by efficient flexible executive processes that help people obey

instructions about relative task priorities. For example, when a driver sees a highway exit, his or her executive processes may end a phone call so that both hands can be put on the steering wheel to take the exit safely.

Of course, not all of our assumptions are entirely new. Some theorists have already begun to describe the functions of executive control in human multiple-task performance (Baddeley, 1986; Duncan, 1986; Logan, 1985; McLeod, 1977; Neisser, 1967; Newell, 1973b; Norman and Shallice, 1986; Shallice, 1972). By characterizing the nature of executive processes more precisely, and by implementing them in the framework of a detailed system architecture, we take further steps toward a comprehensive theory that supplants the traditional RSB hypothesis.

4.2. The EPIC architecture

To embody our basic assumptions, we have developed the EPIC (Executive-Process/Interactive-Control) architecture, which is intended to have many of the same basic properties as the human information-processing system (Kieras and Meyer, 1994; Meyer and Kieras, 1992, 1994, 1995). EPIC may be viewed as a conceptual neighbor of other previously proposed architectures such as the Model Human Processor (Card et al., 1983), ACT* (Anderson, 1983), and SOAR (Laird et al., 1987). Fig. 5 shows EPIC's major components. Among them are specific modules devoted to perceptual, cognitive, and motor processing. The perceptual processors include ones for vision, audition, and touch. The motor processors include ones for manual, articulatory, and ocular action. Each module has software routines, written in the LISP programming language, that send and receive symbolic information to and from other parts of the overall system. Inputs to EPIC's perceptual processors come from simulated sensors (eyes, ears, etc.) that monitor external display devices (CRT screen, headphones, etc.) in a virtual task environment (e.g., the PRP procedure); outputs by EPIC's motor processors go to simulated effectors (hands, mouth, etc.) that operate the environment's external recording devices (keyboard, joystick, voice key, etc.). Constructing models based on EPIC involves programming its cognitive processor to interact with the task environment through the architecture's perceptual and motor processors.

Perceptual processors

During task performance, EPIC's perceptual processors detect and identify stimuli (printed alphanumeric characters, geometric objects, auditory tones, speech, etc.) that occur in the virtual task environment, depositing their symbolic representations in working memory. Consistent with previous empirical research (e.g., Pashler, 1989), each perceptual processor is assumed to operate asynchronously, in parallel with other components of the architecture. The times taken for stimulus detection and identification are task-dependent parameters, whose values we estimate from current data or past literature.

Cognitive processor

EPIC's cognitive processor has no immutable decision or response-selection bottleneck per se. Instead, it relies on three major subcomponents that enable a

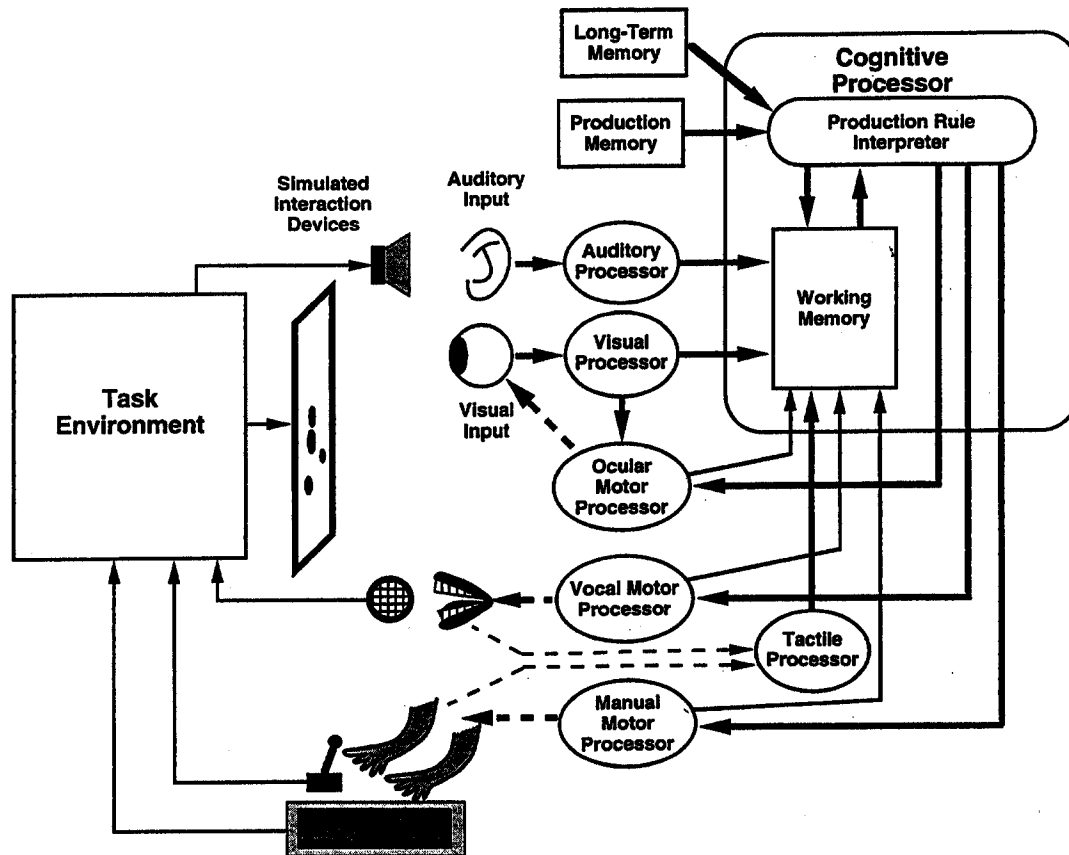


Fig. 5. A schematic diagram of the EPIC (Executive-Process/Interactive-Control) architecture and the virtual task environment with which its components interact during computational simulations of human multiple-task performance.

high degree of parallel processing. These subcomponents include an on-line declarative working memory, production memory, and production-rule interpreter that together implement sets of instructions whereby individual tasks are coordinated and performed simultaneously.

Working memory is assumed to contain various types of information, including (1) symbolic identities of external stimuli sent through EPIC's perceptual processors; (2) symbolic identities of selected responses waiting for transmission to EPIC's motor processors; (3) task goals; (4) sequential control flags; and (5) symbolic notes about the status of other system components (e.g., current motor-processor states). With this information, which evolves systematically over time, performance of each task may proceed efficiently from start to finish.

According to our assumptions, skilled performance is achieved by applying rules stored in EPIC's production memory. These rules, like those postulated by some other theorists (e.g., Anderson, 1976, 1983; Newell, 1973a), have the form "IF X THEN Y", where "X" refers to the current contents of working memory, and "Y" refers to actions that the cognitive processor executes. For example, during a primary auditory-manual choice-reaction task, the following rule might be used to

instruct EPIC's manual motor processor that it should prepare and produce a keypress by the left index finger in response to an 800 Hz tone:

```

IF
  ((GOAL DO TASK 1)
  (STEP DO CHECK FOR TONE 800)
  (AUDITORY TONE 800 ON)
  (STRATEGY TASK 1 RESPONSE MOVEMENT IS IMMEDIATE))
THEN
  ((SEND-TO-MOTOR (MANUAL PRESS LEFT INDEX))
  (ADD (TASK 1 RESPONSE UNDERWAY))
  (ADD (STEP WAIT FOR TASK 1 RESPONSE COMPLETION))
  (DEL (STEP DO CHECK FOR TONE 800))
  (DEL (AUDITORY TONE 800 ON))).

```

The actions of this rule, which not only instructs the manual motor processor but also adds and deletes specified symbolic items in working memory, would be executed whenever working memory contains all of the items in the rule's conditions. For each task that a person has learned to perform skillfully, there would be a set of such rules stored in EPIC's production memory. Also, complementing these task-rule sets, production memory is assumed to contain sets of executive-process rules that manage the contents of working memory, and that coordinate performance depending on task instructions and perceptual-motor constraints.

Task and executive rules in EPIC's production memory are tested and applied by the production-rule interpreter of EPIC's cognitive processor, using the Parsimonious Production System (PPS; Covrigan and Kieras, 1987). Under PPS, the interpreter operates through a series of processing cycles, whose durations are assumed to have a mean length of 50 msec.³ At the start of each cycle, the interpreter tests the conditions of all rules in production memory, determining which ones match the current contents of working memory. At the end of each cycle, for every rule whose conditions are completely matched by the current contents of working memory, all of the rule's actions are executed by the cognitive processor.

We assume that there is no limit on how many production rules can have their conditions tested and actions executed during any particular processing cycle. Also, the cycle durations do not depend on the number of rules involved. It is in this respect that EPIC's cognitive processor has no bottleneck per se. Through appropriate sets of task rules, the cognitive processor may simultaneously select responses and do other operations for multiple concurrent tasks, without between-task interference at this "central" level. Our computational models of multiple-task

³ During actual runs, the cognitive processor's cycle durations are sampled from a distribution whose standard deviation is typically 10 ms (i.e., 20% of the 50 ms mean), introducing realistic stochastic variation into simulated RT data (Kieras and Meyer, 1994; Meyer and Kieras, 1992, 1994, 1995).

performance avoid conflicts among the actions of task rules at peripheral perceptual and motor levels by including executive-process rules that help coordinate and schedule tasks harmoniously.

Motor processors

Upon receiving instructions from the cognitive processor, EPIC's motor processors convert symbolic identities of selected responses to specific features that desired overt movements should have. For example, a manual movement might have features that specify the style, hand, and finger (e.g., PRESS, LEFT, INDEX) to be used. We assume that the features for a response movement are prepared serially, with each feature adding a mean increment of 50 msec to the total movement-production time (cf. Rosenbaum, 1980). Under certain conditions, some features for anticipated response movements may be prepared in advance, if their identities are partially known beforehand. After all of the features for a response movement have been prepared, it is produced overtly through a final initiation step that likewise adds a mean increment of 50 msec. Because the motor preparation and initiation of overt movements are architecturally separate from the prior selection of symbolic response identities, EPIC enables precise control over the flow of information through its components. While response selection may occur simultaneously for multiple concurrent tasks, the production of distinct movements may be temporally staggered, depending on prevalent task instructions and available resources at the motor level.

As indicated already (Fig. 5), EPIC includes distinct motor processors for manual, vocal, and ocular action. Each of these motor processors is assumed to operate in parallel with the others. We also assume, however, that each motor processor only has the capacity to prepare and initiate one response movement at a time. Thus, at the motor level, there are explicit peripheral bottlenecks in EPIC (cf. Ivry et al., 1994,1995; Keele, 1973).

An especially relevant instance of this concerns manual movements. Based on results reported previously about manual movement production (e.g., McLeod, 1977), EPIC has only one motor processor devoted to preparing and initiating movements by the two (i.e., right and left) hands. Thus, for multiple manual tasks, substantial between-task interference is possible at the peripheral motor level. Such interference must be avoided through judicious executive control.

4.3. Adaptive executive-control models

Within the framework of the EPIC architecture, we have formulated a class of adaptive executive-control (AEC) models to characterize performance in the PRP procedure. Our AEC models incorporate executive processes that flexibly control the extent to which progress on a secondary task overlaps with a primary task. Fig. 6 illustrates how this control is achieved.

According to this view, performance of each task goes through several steps, including stimulus identification, response selection, and movement production, consistent with discrete stage models (Sternberg, 1969; Sanders, 1980). Further-

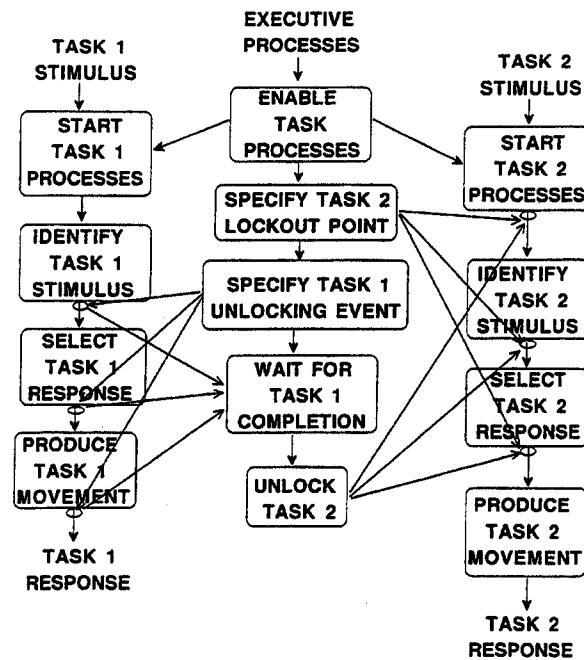


Fig. 6. Component processes for the class of adaptive executive-control (AEC) models whereby the tasks of the PRP procedure may be flexibly coordinated.

more, there is assumed to be an executive process that coordinates Tasks 1 and 2. Its supervisory functions include (1) enabling the primary and secondary tasks to proceed at the start of each trial; (2) specifying a Task 2 lockout point; (3) specifying a Task 1 unlocking event; (4) waiting for the Task 1 unlocking event to occur; and (5) unlocking Task 2 processes so that they may be completed.

Task 2 lockout points

By definition, the Task 2 lockout point is a point during the course of Task 2 such that when it has been reached, further processing for Task 2 stops until Task 1 enters a "done" state. There are at least three potential alternative Task 2 lockout points (Fig. 6, right-side ovals), which are located respectively before the start of stimulus identification, response selection, and movement production for Task 2. Depending on whether the executive process specifies a pre-movement, pre-selection, or pre-identification lockout point, progress on Task 2 would overlap more or less with Task 1.

Task 1 unlocking events

The extent of such overlap is also influenced by the specification of a Task 1 unlocking event. By definition, this is an event during the course of Task 1 such that when it occurs, Task 1 is deemed to be "done," and the executive process permits processing for Task 2 to continue beyond the Task 2 lockout point. There are several potential alternative Task 1 unlocking events (Fig. 6, left-side ovals); Task 1 may be deemed "done" immediately after either its stimulus-identification,

response-selection, or movement-production stage finishes. Again, depending on whether the executive process specifies a post-identification, post-selection, or post-movement unlocking event, progress on Task 2 would overlap more or less with Task 1.

Executive production rules

At the start of each trial, our AEC models' executive process specifies a particular Task 2 lockout point and Task 1 unlocking event by putting their designations in working memory. For example, the following executive production rule enables a post-response-selection lockout point for Task 2 and a post-movement-initiation unlocking event for Task 1:

```

IF
  ((GOAL DO PRP PROCEDURE)
   (STRATEGY AUDITORY-MANUAL TASK 1)
   (STRATEGY VISUAL-MANUAL TASK 2)
   (VISUAL FIXATION POINT DETECTED)
   (NOT (TRIAL UNDERWAY)))
THEN
  ((SEND-TO-MOTOR MANUAL RESET)
   (DEL (VISUAL FIXATION POINT DETECTED))
   (ADD (TRIAL UNDERWAY))
   (ADD (GOAL DO TASK 1))
   (ADD (GOAL DO TASK 2))
   (ADD (STRATEGY TASK 2 RESPONSE MOVEMENT IS DEFERRED))
   (ADD (STRATEGY UNLOCK ON MOTOR-SIGNAL MANUAL STARTED LEFT))
   (ADD (STEP WAIT FOR TASK 1 DONE))).

```

Subsequently, when EPIC's manual motor processor informs the cognitive processor that the Task 1 response movement (a left-hand key press) has been initiated, the following executive production rule unlocks Task 2 and lets it finish:

```

IF
  ((GOAL DO PRP PROCEDURE)
   (STRATEGY AUDITORY-MANUAL TASK 1)
   (STRATEGY VISUAL-MANUAL TASK 2)
   (STEP WAIT-FOR TASK 1 DONE))
   (TASK 1 DONE)
THEN
  ((DEL (STEP WAIT-FOR TASK 1 DONE))
   (DEL (TASK 1 DONE))
   (DEL (STRATEGY TASK 2 RESPONSE MOVEMENT IS DEFERRED))
   (ADD (STRATEGY TASK 2 RESPONSE MOVEMENT IS IMMEDIATE))
   (ADD (STEP WAIT-FOR TASK 2 DONE))).

```

As a result, response-selection but not movement-production stages for the two tasks could overlap. Other executive production rules may enable different lockout points and unlocking events instead of those just illustrated, regulating the amount of overlap that actually occurs.

Particular AEC models

Overall, the class of AEC models includes many particular cases. For each possible combination of Task 2 lockout point and Task 1 unlocking event, there is a different set of executive production rules that can implement this combination, achieving a certain discretionary amount of temporal overlap between the two tasks. Which executive process is used under what circumstances may vary with task instructions, strategic goals, perceptual-motor requirements, and past experience.

In particular, one of our AEC models can mimic a response-selection bottleneck. Its executive process does this by specifying a pre-selection lockout point for Task 2 and a post-selection unlocking event for Task 1, thereby precluding response selection during Task 2 until Task 1 response selection has finished. Within EPIC's framework, however, such specifications are neither obligatory nor immutable, contrary to the traditional RSB hypothesis. An optional response-selection bottleneck may, but need not, be imposed when the situation encourages making extremely sure that Task 2 responses never precede Task 1 responses.

Other particular AEC models can mimic additional types of bottleneck. For example, Keele (1973) has hypothesized that a movement-initiation bottleneck rather than a response-selection bottleneck exists in the human information-processing system. Consistent with this hypothesis, an executive process may defer Task 2 movement initiation by specifying a post-selection/pre-movement lockout point for Task 2 and a post motor-initiation unlocking event for Task 1. Again, however, such specifications are neither obligatory nor immutable in EPIC. An optional movement-initiation bottleneck may, but need not, be imposed when the situation encourages producing Task 2 responses as quickly as possible after Task 1 finishes.

5. Qualitative accounts of PRP phenomena

Unified qualitative accounts for a variety of PRP phenomena, including many beyond the scope of the traditional RSB hypothesis, are provided by the EPIC architecture and its AEC models.

PRP effect

In terms of our theoretical framework, elevated Task 2 RTs at short SOAs result mainly from having to satisfy task instructions for the PRP procedure. Due to these instructions, Task 2 cannot proceed freely from start to finish along with

Task 1, because doing so might yield premature Task 2 responses when Task 1 is relatively hard and the SOA is short. Thus, executive processes for the PRP procedure must, out of strategic necessity, specify some intermediate unlocking event and lockout point for Tasks 1 and 2 respectively, delaying overt Task 2 responses enough that they never precede Task 1 responses. Recently, Koch (1993,1994) has offered an independent account of the PRP effect that is similar to ours, thereby reinforcing some of the present article's main premises.

Diverse forms of PRP curves

Given the adjustability of their lockout points and unlocking events, our AEC models likewise imply that PRP curves may have diverse forms. If the executive process adopts a pre-selection lockout point for Task 2, then it can yield "parallel" (i.e., vertically equidistant) PRP curves of mean Task 2 RTs as in Fig. 1. This would seem especially plausible when Task 1 is relatively difficult and has a high probability of finishing after Task 2 at short SOAs unless the executive process strongly intervenes. In contrast, if Task 1 is relatively easy and encourages a more ambitious strategy that needs to guard less against premature Task 2 responses, then the executive process may adopt a post-selection lockout point for Task 2, thereby producing divergent PRP curves like those in the left panels of Figs. 2 and 4 (Meyer and Kieras, 1992,1994,1995).

Convergent PRP curves (e.g., Fig. 2, right panel) are also accommodated naturally by our AEC models (Meyer and Kieras, 1995). Suppose that at the start of each trial, the unlocking event and lockout point specified for Tasks 1 and 2, respectively, depend on the anticipated difficulty of response selection in Task 2. Also, suppose that the specified Task 2 lockout point is a relatively earlier one when Task 2 will be difficult than when it will be easy, whereas the Task 1 unlocking event is a relatively later one. Then less overlap may occur between Tasks 1 and 2 at short SOAs in the difficult Task 2 condition than in the easy Task 2 condition, causing the difficult Task 2 to manifest a larger PRP effect than does the easy Task 2. Combined with the main effect of Task 2 difficulty at long SOAs, this difference between PRP effects in the easy and difficult Task 2 conditions would yield a pair of converging PRP curves. A possible rationale for such difficulty-dependent task scheduling is that, although not necessary under EPIC, it may seem to help preclude a difficult Task 2 from interfering more with Task 1 than does an easy Task 2.

Slopes steeper than -1

Similarly, our AEC models account for PRP curves whose slopes are steeper than -1 . Suppose that at the start of each trial, the executive process specifies an initial cautious unlocking event and lockout point for Tasks 1 and 2, respectively. Also, suppose that after the Task 1 stimulus has arrived, no Task 2 stimulus is detected during a subsequent period of time (i.e., the SOA is somewhat greater than zero). Then the executive process may modify the Task 2 lockout point and/or Task 1 unlocking event by updating their designations in working memory,

because Task 1 now has a better chance of finishing first without much further delay in Task 2. Specifically, the executive process could make the modified Task 2 lockout point be later and/or Task 1 unlocking event be earlier than they were before, using what we call *progressive unlocking* (Meyer and Kieras, 1995). With progressive unlocking, mean Task 2 RTs at intermediate SOAs would be less than when the lockout point and unlocking event are static throughout each trial. The extra RT reduction, combined with the usual RT decrement caused by increasing the SOA, therefore yields PRP curves whose slopes are steeper than -1 , as in Fig. 3 (left panel). Indeed, such extreme steepness may be a hallmark of sophisticated executive processes that are sensitive to rapidly evolving contingencies in multiple-task performance.

Individual differences

Of course, if people have such executive control, then individual differences might occur in their patterns of PRP curves. Depending on personal factors, different subjects may be inclined to adopt different Task 2 lockout points and Task 1 unlocking events. If so, then this would yield mixtures of diverging, parallel, and converging PRP curves, as some investigators have reported (e.g., Ivry et al., 1994,1995; Lauber et al., 1994). Furthermore, the curves produced by any particular individual might change from one set of conditions to another, depending on how each condition meshes with the subject's predilections.

Training effects

Yet despite these individual differences, our AEC models also imply that executive processes can be shaped and homogenized through proper experience. If Task 2 lockout points and Task 1 unlocking events are adjustable, then certain types of training should induce more overlap between primary and secondary tasks. For example, subjects might come to adopt more optimal executive processes when responses must be produced rapidly in an unconstrained rather than constrained order (Koch, 1993,1994; Lauber et al., 1994; cf. Pashler, 1990). Consequently, upon later transfer to the standard PRP procedure, PRP curves may embody a pattern that is similar across subjects and indicative of concurrent response selection (e.g., Fig. 4). Moreover, if there are no constraints on the order in which subjects must make their responses, then the PRP effect may virtually disappear (Koch, 1993,1994), consistent with EPIC's capacity to select and execute multiple responses simultaneously.

6. Computational simulations of quantitative PRP data

Additional justification of present claims is provided by computational simulations that account quantitatively for data from the PRP procedure (Meyer and Kieras, 1992,1994,1995). Our simulations to date are based on one instructive member of the AEC class. We call it the strategic response-deferment (SRD) model.

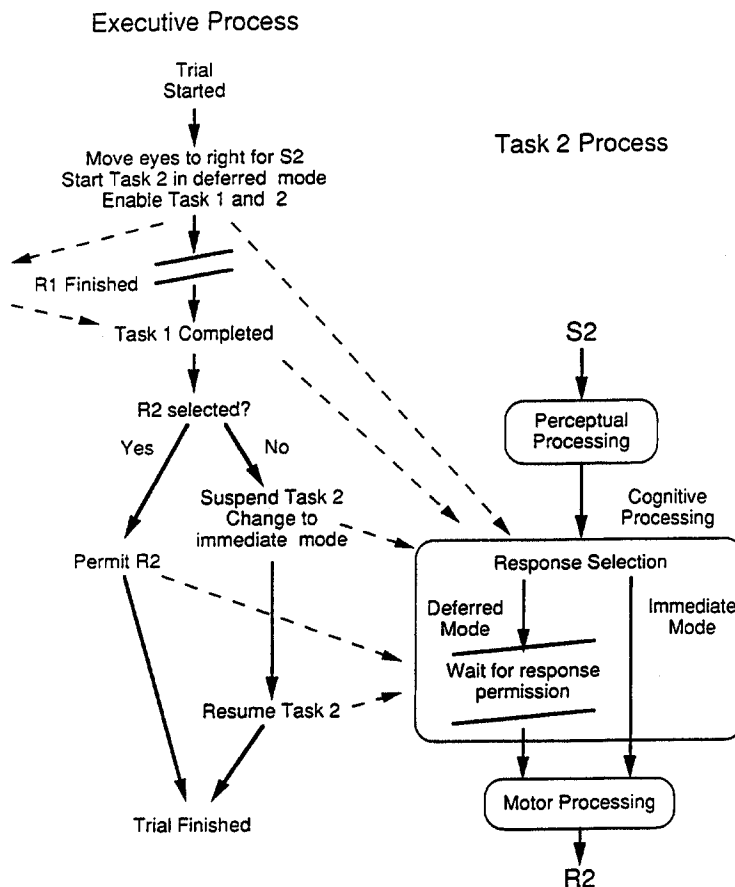


Fig. 7. Flowchart for executive and Task 2 processes under the strategic response-deferment model.

6.1. Strategic response-deferment model

Fig. 7 shows the SRD model's executive process, which starts each trial of the PRP procedure by putting Tasks 1 and 2 in "immediate" and "deferred" mode, respectively. While Task 2 is in deferred mode, the identities of Task 2 responses may be selected and sent to working memory, but Task 2 response movements are not produced by EPIC's motor processors. This constraint is imposed by assigning a post-selection/pre-motor lockout point to Task 2. Putting Task 1 in immediate mode lets its responses be selected and sent directly to their motor processor. When the Task 1 unlocking event occurs (e.g., the Task 1 response movement is initiated), the executive process temporarily suspends Task 2 (i.e., withdraws "GOAL DO TASK 2" from working memory) and shifts it to immediate mode, after which Task 2 is resumed again (i.e., "GOAL DO TASK 2" is reinserted in working memory). Following this transition, previously selected Task 2 responses are sent directly from working memory to their motor processor. If response selection has not yet finished for Task 2 before it enters the immediate mode, then Task 2 production rules may both select and send Task 2 responses to their motor processor. Because response selection for Task 2 is suspended briefly during the




Fig. 8. Left panel: goodness-of-fit between simulated mean RTs (small symbols on dashed curves) from the SRD model and empirical mean RTs (large symbols on solid curves) from Hawkins et al.'s (1979) PRP study with an auditory-vocal Task 1 and visual-manual Task 2. Filled circles and triangles represent Task 2 RTs when response-selection was respectively easy or hard; unfilled circles and triangles represent corresponding Task 1 RTs. Right panel: goodness-of-fit between simulated mean RTs from the SRD model and empirical mean RTs from McCann and Johnston's (1992, Exp. 2) PRP study.

transition from deferred to immediate mode, the SRD model has a flexible combination of temporary "soft" movement-initiation and response-selection bottlenecks (cf. De Jong, 1993; Kantowitz, 1974; Keele, 1973).⁴

6.2. Simulations with the SRD model

With the SRD model, we have successfully simulated quantitative data from many representative PRP studies (e.g., Hawkins et al., 1979; Karlin and Kestenbaum, 1968; McCann and Johnston, 1992; Pashler, 1990), confirming the utility of the EPIC architecture and the validity of our present theoretical claims (Meyer and Kieras, 1992, 1994, 1995).

PRP study by Hawkins et al. (1979)

One PRP study that has provided us with extensive relevant data is by Hawkins et al. (1979). In part of this study, subjects performed an auditory-vocal Task 1 (saying words in response to two alternative tones) and an easy or hard visual-manual Task 2 (pressing keys in response to either two or eight alternative printed

⁴ Unlike the movement-initiation bottleneck hypothesis of Keele (1973) and the multiple-bottleneck hypothesis of De Jong (1993), however, the SRD model assumes that these bottlenecks are optional – not immutable – ones programmed by the executive process to efficiently satisfy instructions of the PRP procedure.

digits). A comparison between Hawkins et al.'s empirical mean RTs from these conditions and simulated mean RTs from the SRD model appears in Fig. 8 (left panel, solid vs. dashed curves).

As this graph indicates, the SRD model accounts well ($R^2 = 0.99$) for the positive interaction that Hawkins et al. (1979) found between SOA and response-selection difficulty in Task 2, which yielded divergent PRP curves. This follows because the model's executive process lets response selection proceed simultaneously for Tasks 1 and 2 at short SOAs, so the difficulty effect on Task 2 is absorbed by waiting for the Task 1 unlocking event. With optional progressive unlocking (Meyer and Kieras, 1995), the executive process accurately reproduces a slope significantly steeper than -1 , which occurred in the PRP curve at short SOAs when Task 2 was relatively easy (Fig. 8, left panel). Our simulations of other data from Hawkins et al.'s (1979) study also produced good fits (e.g., see Fig. 2, left panel, solid vs. dashed curves).⁵

PRP study by McCann and Johnston (1992)

In addition, the SRD model accounts well for data from studies that have yielded "parallel" (i.e., vertically equidistant) rather than divergent PRP curves. For example, consider the right panel of Fig. 8, which shows some of McCann and Johnston's (1992, Exp. 2) data. Here the effects of SOA and response-selection difficulty on empirical mean Task 2 RTs (solid curves) are additive. Similarly, simulated mean Task 2 RTs (dashed curves) produced by the SRD model manifest such additivity ($R^2 = 0.99$).

The latter pattern can be understood more fully in terms of how McCann and Johnston's (1992, Exp. 2) subjects were tested. During the PRP procedure, they performed an auditory-vocal Task 1 (saying words in response to tones) together with an easy or hard visual-manual Task 2 (pressing keys in response to horizontal arrows or printed letters). On each trial, a relatively small Task 2 stimulus was displayed several degrees to the right or left of a central visual fixation point. Subjects could not predict beforehand exactly where the Task 2 stimulus would be located. Following the SOA, eye movements presumably had to be made for stimulus identification in Task 2. This requirement – which is mimicked by the SRD model – probably created a temporary peripheral perceptual bottleneck that precluded Task 2 response selection from overlapping with response selection for Task 1. Because Task 1 was relatively easy, subjects may have already finished it and entered the unlocking phase of the SRD model (Fig. 7) by when the Task 2 stimulus identity became available for response selection (Meyer and Kieras, 1994, 1995).

More generally, this interpretation raises an important meta-theoretical point. Results (e.g., "parallel" PRP curves) that are superficially consistent with the

⁵ Across these different conditions, the number of variable parameter values used by the SRD model was markedly less than the number of reliable one-degree-of-freedom contrasts between mean RTs in Hawkins et al.'s data (Meyer and Kieras, 1995).

traditional RSB hypothesis may actually have much subtler sources. Thus, future interpretation of data from the standard PRP procedure and other multiple-task paradigms should take such subtleties more fully into account.

7. Conclusion

In conclusion, our discourse on the RSB hypothesis, PRP procedure, EPIC architecture, and AEC/SRD models has potentially significant implications for characterizing discrete versus continuous human information processing. If the present theoretical claims are valid, then people's performance may entail a variety of concurrent discrete perceptual-motor and cognitive processes that provide symbolic partial outputs to each other. We therefore concur with at least some of the assumptions made by Miller's (1982, 1988) asynchronous discrete-coding model, under which stimulus identification and response selection overlap temporally, producing quantized intermediate outputs about relevant stimulus and response features, respectively. Furthermore, it now appears that when two or more tasks do not logically conflict, sets of production rules for them may be used simultaneously as if procedural cognitive processes have multiple channels rather than a single-channel response-selection bottleneck.

Another lesson from our research is that even in very elementary situations, sophisticated executive processes play a crucial role. For any task, there are many alternative paths from stimulus input to response output in the human information-processing system. The path that is actually taken, and the extent to which processing may seem "discrete" or "continuous," can depend on control strategies that subjects adopt. Future research on discrete versus continuous processing should take these strategies more fully into account. This may be facilitated by formulating a comprehensive system architecture and detailed computational models. An important role of such models will be to help specify how working memory is judiciously used so that procedural cognitive processes may interact effectively with limited-capacity peripheral perceptual-motor components.

Of course, now is not the first occasion on which human-performance theorists have needed to radically change their world view. More than fifty years ago, for example, a dominant model in sensory psychophysics was high-threshold theory (HTT). Analogous to the traditional RSB hypothesis, HTT claimed that people detect simple sensory stimuli (e.g., lights, tones, etc.) through a discrete all-or-none threshold mechanism. In order for a stimulus to be detected and reported, its subjective intensity supposedly had to exceed an absolute level within this mechanism. Because of the assumed threshold's rigidity, little or no accommodation was provided by HTT for subjects' possible judgment strategies. As a result, many problematic aspects of psychophysical data went unexplained. Then, however, statistical signal-detection theory (SDT) emerged on the scene, reconciling phenomena that had previously bedeviled HTT (Tanner and Swets, 1954).

Unlike HTT, this new framework assumed no discrete absolute high threshold; instead, SDT attributed subjects' detection performance to stochastic processes

that involve a continuum of sensory states and adjustable decision criteria. According to SDT, people set their decision criteria strategically to achieve various combinations of stimulus “hits” and noise “correct rejections,” depending on prevailing reward schemes. The adjustable decision criteria of SDT have much the same spirit as the flexible lockout points and unlocking events of our AEC models for the PRP procedure. As in our AEC models, a key insight of SDT has been that even the seemingly most elementary human performance – for example, detection of sensory stimuli – is governed by sophisticated programmable executive processes rather than just rigid peripheral mechanisms. Perhaps keeping this historical precedent in mind will help smooth the entry of human-performance theory into a Brave New World without pervasive immutable response-selection bottlenecks.

Acknowledgements

We gratefully acknowledge helpful comments, suggestions, and criticisms from Bert Mulder, Andries Sanders, and our many other colleagues. Financial support has been provided by the U.S. Office of Naval Research through grant N00014-92-J-1173 to the University of Michigan.

References

- Allport, D.A., 1980. ‘Attention and performance’. In: G.L. Claxton (ed.), *Cognitive psychology: New directions* (pp. 112–153). London: Routledge and Kegan Paul.
- Allport, D.A., 1987. ‘Selection-for-action: Some behavioral and neurophysiological considerations of attention and action’. In: H. Heuer and A.F. Sanders (eds.), *Perspectives on perception and action* (pp. 395–419). Hillsdale, NJ: Erlbaum.
- Anderson, J.R., 1976. *Language, memory, and thought*. Hillsdale, NJ: Erlbaum.
- Anderson, J.R., 1983. *The architecture of cognition*. Cambridge, MA: Harvard University Press.
- Anderson, J.A. and G.E. Hinton, 1981. ‘Models of information processing in the brain’. In: G.E. Hinton and J.A. Anderson (eds.), *Parallel models of associative memory*. Hillsdale, NJ: Erlbaum.
- Baddeley, A.D., 1986. *Working memory*. Oxford: Oxford University Press.
- Becker, C.A., 1976. Allocation of attention during visual word recognition. *Journal of Experimental Psychology: Human Perception and Performance* 2, 556–566.
- Bertelson, P., 1966. Central intermittency 20 years later. *Quarterly Journal of Experimental Psychology* 18, 153–164.
- Card, S.K., T.P. Moran and A. Newell, 1983. *The psychology of human-computer interaction*. Hillsdale, NJ: Erlbaum.
- Covrigaru, A. and D.E. Kieras, 1987. PPS: A parsimonious production system. Tech. Rep. No. 26, TR-87/ONR-26. Ann Arbor, MI: University of Michigan, Technical Communication Program.
- Craik, K.J.W., 1948. Theory of the human operator in control systems: II. Man as an element in a control system. *British Journal of Psychology* 38, 142–148.
- Davis, R., 1957. The human operator as a single-channel information system. *Quarterly Journal of Experimental Psychology* 9, 119–129.
- De Jong, R., 1993. Multiple bottlenecks in overlapping task performance. *Journal of Experimental Psychology: Human Perception and Performance* 19, 965–980.
- Duncan, J., 1986. Disorganization of behaviour after frontal lobe damage. *Cognitive Neuropsychology* 3, 271–290.

- Gopher, D., 1993. 'Attentional control: Acquisition and execution of attentional strategies'. In: D.E. Meyer and S. Kornblum (eds.), *Attention and performance XIV. Synergies in experimental psychology, artificial intelligence, and cognitive neuroscience*. Cambridge, MA: MIT Press.
- Gottsdanker, R., 1980. 'The ubiquitous role of preparation'. In: G.E. Stelmach and J. Requin (eds.), *Tutorials in motor behavior* (pp. 315–371). Amsterdam: North-Holland.
- Gottsdanker, R. and G.E. Stelmach, 1971. The persistence of psychological refractoriness. *Journal of Motor Behavior* 3, 301–312.
- Hawkins, H.L., E. Rodriguez and G.M. Reicher, 1979. Is time-sharing a general ability? ONR Technical Report No. 3, University of Oregon, Eugene, OR.
- Ivry, R.B., E.A. Franz, A. Kingstone and J.C. Johnston, 1994 (November). The PRP effect in a split-brain patient: Response uncoupling despite normal interference. Paper presented at the meeting of the Psychonomic Society, St. Louis, MO.
- Ivry, R.B., E.A. Franz, A. Kingstone and J.C. Johnston, 1995. The PRP effect following callosotomy: Residual interference despite uncoupling of lateralized response codes. Manuscript submitted for publication.
- James, W., 1890. *The principles of psychology*. New York: Holt.
- Kahneman, D., 1973. *Attention and effort*. Englewood Cliffs, NJ: Prentice-Hall.
- Kantowitz, B.H., 1974. 'Double stimulation'. In: B.H. Kantowitz (ed.), *Human information processing: Tutorials in performance and cognition* (pp. 83–131). Hillsdale, NJ: Erlbaum.
- Karlin, L. and R. Kestenbaum, 1968. Effects of number of alternatives on the psychological refractory period. *Quarterly Journal of Experimental Psychology* 20, 67–178.
- Keele, S.W., 1973. *Attention and human performance*. Pacific Palisades, CA: Goodyear.
- Kieras, D.E. and D.E. Meyer, 1994. The EPIC architecture for modeling human information-processing and performance: A brief introduction. Technical Report TR-94/ONR-EPIC-1, University of Michigan, Ann Arbor.
- Koch, R., 1993. *Die psychologische Refraktärperiode*. Doctoral dissertation, University of Bielefeld, Bielefeld, Germany.
- Koch, R., 1994 (December). Hick's Law and the psychological refractory period. Paper presented at the KNAW Symposium on Discrete versus Continuous Information Processing, Amsterdam, The Netherlands.
- Laird, J.E., A. Newell and P.S. Rosenbloom, 1987. SOAR: An architecture for general intelligence. *Artificial Intelligence* 33, 1–64.
- Lauber, E.J., E.H. Schumacher, J. Glass, E. Zurbruggen, D.E. Kieras and D.E. Meyer, 1994 (November). Adaptive PRP effects: Evidence of flexible attention to action. Paper presented at the meeting of the Psychonomic Society, St. Louis, MO.
- Logan, G., 1985. Executive control of thought and action. *Acta Psychologica* 60, 193–210.
- Luce, R.D., 1986. *Response times: Their role in inferring elementary mental organization*. New York: Oxford University Press.
- McCann, R.S. and J.C. Johnston, 1992. Locus of the single-channel bottleneck in dual-task performance. *Journal of Experimental Psychology: Human Perception and Performance* 18, 471–484.
- McLeod, P.D., 1977. A dual task response modality effect: Support for multiprocessor models of attention. *Quarterly Journal of Experimental Psychology* 29, 651–667.
- McLeod, P., 1978. Parallel processing and the psychological refractory period. *Acta Psychologica* 41, 381–396.
- Meyer, D.E. and D.E. Kieras, 1992 (November). The PRP effect: Central bottleneck, perceptual-motor limitations, or task strategies? Paper presented at the meeting of the Psychonomic Society, St. Louis, MO.
- Meyer, D.E. and D.E. Kieras, 1994. EPIC computational models of psychological refractory-period effects in human multiple-task performance. Technical Report TR-94/ONR-EPIC-2, University of Michigan, Ann Arbor.
- Meyer, D.E. and D.E. Kieras, 1995. A computational theory of human multiple-task performance: The EPIC architecture and strategic response-deferment model. *Psychological Review*, in press.
- Meyer, D.E., A.M. Osman, D.E. Irwin and S. Yantis, 1988. Modern mental chronometry. *Biological Psychology* 26, 3–67.

- Miller, J., 1982. Discrete versus continuous stage models of human information processing: In search of partial output. *Journal of Experimental Psychology: Human Perception and Performance* 8, 273–296.
- Miller, J., 1988. Discrete and continuous models of human information processing: Theoretical distinctions and empirical results. *Acta Psychologica* 67, 191–257.
- Moray, N., 1967. Where is capacity limited? A survey and a model. *Acta Psychologica* 27, 84–92.
- Navon, D. and D. Gopher, 1979. On the economy of the human-processing system. *Psychological Review* 86, 214–255.
- Neisser, U., 1967. *Cognitive psychology*. Englewood Cliffs, NJ: Prentice Hall.
- Neumann, O., 1987. 'Beyond capacity: A functional view of attention'. In: H. Heuer and A.F. Sanders (eds.), *Perspectives on perception and action* (pp. 361–394). Hillsdale, NJ: Erlbaum.
- Newell, A., 1973a. 'Production systems: Models of control structures'. In: W.G. Chase (ed.), *Visual information processing* (pp. 463–526). New York: Academic Press.
- Newell, A., 1973b. 'You can't play 20 questions with nature and win'. In: W.G. Chase (ed.), *Visual information processing* (pp. 283–308). New York: Academic Press.
- Norman, D.A., 1969. *Memory and attention: An introduction to human information processing*. New York: Wiley.
- Norman, D.A. and T. Shallice, 1986. 'Attention to action: Willed and automatic control of behavior'. In: R.J. Davidson, G.E. Schwartz and D. Shapiro (eds.), *Consciousness and self-regulation*, Vol. 4. New York: Plenum Press.
- Osman, A.M., T.R. Bashore, M.G.H. Coles, E. Donchin and D.E. Meyer, 1992. On the transmission of partial information: Inferences from movement-related brain potentials. *Journal of Experimental Psychology: Human Perception and Performance* 18, 217–232.
- Pashler, H., 1984. 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., 1993. 'Dual-task interference and elementary mental mechanisms'. In: D.E. Meyer and S. Kornblum (eds.), *Attention and performance XIV. Synergies in experimental psychology, artificial intelligence, and cognitive neuroscience* (pp. 245–264). Cambridge, MA: MIT Press.
- Pashler, H., 1994. Dual-task interference in simple tasks: Data and theory. *Psychological Bulletin* 116, 220–244.
- Pashler, H. and J.C. Johnston, 1989. Chronometric evidence of central postponement in temporally overlapping tasks. *Quarterly Journal of Experimental Psychology* 41A, 19–45.
- Rosenbaum, D.A., 1980. Human movement initiation: Specification of arm, direction, and extent. *Journal of Experimental Psychology: General* 109, 475–495.
- Rumelhart, D.E. and J.L. McClelland (eds.), 1986. *Parallel distributed processing*. Cambridge, MA: MIT Press.
- Ruthruff, E., J.O. Miller and T. Lachmann, 1995. Does mental rotation require central mechanisms? *Journal of Experimental Psychology: Human Perception and Performance*, in press.
- Sanders, A.F., 1964. Selective strategies in the assimilation of successively presented signals. *Quarterly Journal of Experimental Psychology* 16, 368–372.
- Sanders, A.F., 1980. 'Stage analysis of reaction processes'. In: G.E. Stelmach and J. Requin (eds.), *Tutorials in motor behavior* (pp. 331–354). Amsterdam: North-Holland.
- Schvaneveldt, R.W., 1969. Effects of complexity in simultaneous reaction time tasks. *Journal of Experimental Psychology* 81, 289–296.
- Schweickert, R., A. Dutta, C. Sangsup and R.W. Proctor, 1992 (November). Scheduling processes using working memory. Paper presented at the meeting of the Psychonomic Society, St. Louis, MO.
- Shallice, T., 1972. Dual functions of consciousness. *Psychological Review* 79, 383–393.
- Smith, M.C., 1967. Theories of the psychological refractory period. *Psychological Bulletin* 67, 202–213.
- Sternberg, S., 1969. On the discovery of processing stages: Some extensions of Donders' method. *Acta Psychologica* 30, 276–315.

- Tanner, W.P., Jr. and J.A. Swets, 1954. A decision-making theory of visual detection. *Psychological Review* 61, 401–409.
- Townsend, J.T. and F.G. Ashby, 1983. *Stochastic modeling of elementary psychological processes*. Cambridge: Cambridge University Press.
- Van Selst, M. and P. Jolicoeur, 1993 (November). A response-selection account of the effect of number of alternatives on dual-task processing. Paper presented at the meeting of the Psychonomic Society, Washington, DC.
- Vince, M.A., 1949. Rapid response sequences and the psychological refractory period. *British Journal of Psychology* 40, 23–40.
- 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., 1959. Evidence of a single-channel decision mechanism limiting performance in a serial reaction task. *Quarterly Journal of Experimental Psychology* 2, 193–210.
- Welford, A.T., 1967. Single channel operation in the brain. *Acta Psychologica* 27, 5–22.
- Wickens, C.D., 1984. ‘Processing resources in attention’. In: R. Parasuraman, J. Beatty and R. Davies (eds.), *Varieties of attention* (pp. 63–101). New York: Wiley.