

Switching Between Simple Cognitive Tasks: The Interaction of Top-Down and Bottom-Up Factors

Eric Ruthruff, Roger W. Remington, and James C. Johnston
NASA - Ames Research Center

How do top-down factors (e.g. expectancy) and bottom-up factors (e.g. recency) interact to produce an overall level of task readiness? We addressed this question by factorially manipulating task expectancy and task repetition in a task-switching paradigm. Although the effects of expectancy and repetition on response time tended to interact underadditively, this interaction occurred only because the traditional binary task-repetition variable lumps together all switch trials, ignoring variation in task lag. When the task recency variable was scaled continuously, all four experiments showed additivity between expectancy and recency. These results argue against several a priori plausible hypotheses, including that participants suppress recently-used task sets when anticipating a task switch. Our findings suggest instead that expectancy and recency selectively influence distinct stages of mental processing. One specific possibility (the Configuration-Execution model) is that task expectancy affects the time required to program upcoming central mental operations, whereas task recency affects the time required to execute those central operations once they have been programmed.

Executive control processes play a critical role in our ability to flexibly adapt to changing environments. For instance, control processes are needed to ensure that the cognitive machinery are configured to perform the task appropriate to the current goals rather than some competing task (such as the one performed most recently or most frequently). In most cognitive psychology experiments, however, executive control processes are held constant rather than investigated. As a result, relatively little is known about them. In particular, little is known about how bottom-up, stimulus-driven processes interact with top-down, executive task-control processes. The present paper addresses this issue, building on some recent advances in task-control research.

Background: The “Switch Cost”

One of the earliest approaches to studying task preparation was to compare performance on mixed lists of two or more tasks to performance on pure lists of a single task. The main question was whether there is a time cost associated with switching between

cognitive tasks. Jersild (1927), for example, conducted several experiments in which he asked participants to perform arithmetic operations on digit pairs. In the pure-list condition, participants repeated the same operation (e.g. addition) on every digit pair. In the mixed-list condition, participants switched between two different operations (e.g. addition and subtraction). In several instances, Jersild found that participants took much longer to complete mixed lists than pure lists, an effect that has become known as the task-switch cost, or “switch cost” for short. Spector and Biederman (1976) later replicated these results with several technological and methodological improvements. For the sake of clear exposition, note that we use the term “task switch” to refer to experimental paradigms (such as Jersild’s) in which the task sometimes switches from trial to trial; when referring to a putative mental reconfiguration of task set (which might or might not occur in the task-switch

paradigm), we use theoretical terms such as “shift of set” or “task-set reconfiguration.”

The task-switch paradigm used by Jersild to study task preparation is simple and straightforward. However, it also suffers from the serious drawback that the task-repetition and task-switch conditions are run within separate lists of items. As a result, participants need to keep only one task available in pure lists, but must keep both tasks available in mixed lists. Thus, it is difficult to determine how much of the mixed-list slowing is due to the costs of task-shifting, per se, and how much is due to the costs of keeping both task sets available at the same time (cf. Gottsdanker, 1980). The use of separate lists of items also leaves open the possibility that participants adopt very different response criteria or task strategies in the repetition and switch conditions. In addition, participants might experience higher levels of arousal, or simply exert more effort, in one condition or the other (see Rogers and Monsell, 1995).

To avoid the problems inherent in Jersild’s design, Rogers and Monsell (1995) used an alternating-runs paradigm in which the task-repetition and switch conditions were mixed together within the same block of trials. Participants performed a run of trials (usually two) of one task before switching to the next task; for example, tasks A and B were typically presented in the sequence AABBAABB, etc. Note that half of these trials were task switches and half were task repetitions. In the experiments reported by Rogers and Monsell, one task required a vowel/consonant classification on a letter stimulus and the other task required an odd/even classification on a digit stimulus. Typically, each display contained both a letter and a digit, forcing participants to suppress processing of whichever character was irrelevant on that trial. The switch cost found by Rogers and Monsell was typically very large (roughly 300 ms), essentially replicating the basic results of Jersild (1927).

Single- versus Dual-Affordance

Perhaps the most important determinant of the size of the switch cost is whether the presented stimulus unambiguously indicates which task should be performed; that is, whether the stimulus offers affordances for just one of the relevant tasks, or more than one. When the two tasks share the same stimulus set (e.g. addition vs. subtraction of digit pairs), the stimulus clearly offers affordances for both tasks. We refer to this as a dual-affordance condition. Dual-affordance can occur even when the tasks have distinct stimulus sets, provided that each trial contains one stimulus for each task. Rogers and Monsell (1995), for example, typically presented both a letter and a digit at the same time; it is not immediately obvious from such a display whether one should perform the odd/even judgment on the digit or the vowel/consonant judgment on the letter. Rogers and Monsell also studied (Experiment 4) a single-affordance condition in which only one stimulus or the other (a letter or a digit) was presented on each trial (along with a neutral character).

The switch cost, which is usually very large in dual-affordance conditions, tends to be relatively small in single-affordance conditions (Allport, Styles, and Hsieh, 1994; Jersild, 1927; Spector and Biederman, 1976; Rogers and Monsell, 1995). In Rogers and Monsell (1995), for example, the switch cost was about 300 ms in the dual-affordance condition (Experiments 1-3), but only about 50 ms in the single-affordance condition (Experiment 4).

The modest switch costs observed in the single-affordance condition might reflect the time required for a shift-of-set operation (e.g. discarding the old task set and/or installing the new task set) on task-switch trials. A further possibility, not necessarily exclusive with the first, is that certain

critical processing stages simply proceed more slowly on task-switch trials because the task has not been performed recently, and is therefore not highly “activated” (or prepared). These potential sources of the switch cost would, of course, also occur in the dual-affordance condition as well. However, the dual-affordance condition introduces several new sources of interference not present in the single-affordance condition. For example, additional time might be required to suppress activation of the inappropriate task set. In addition, if this suppression is not completely successful then processing of the inappropriate task might interfere with processing of the appropriate task. It is even possible that on some trials participants begin to perform the inappropriate task, exclusively, then realize their mistake and start over with the appropriate task.

To summarize, there are many more potential sources of switch cost in the dual-affordance condition than in the single-affordance condition. The need to disentangle these multiple sources would be a daunting obligation. We opted to concentrate instead on the single-affordance condition. It seemed prudent to first try to understand the simpler scenarios before studying more complicated ones. In addition, we believe that the single-affordance condition is more representative of real-world task-switching scenarios than is the dual-affordance condition.

Top-Down Influences: Deliberate Task Preparation

Another important determinant of performance in task-switching paradigms is the degree of deliberate (i.e. top-down) task preparation. McCann, Remington, and Folk (2000) studied deliberate task preparation using a paradigm in which they cued one task prior to each trial (see also Sudevan & Taylor, 1987), then presented either that task (the valid condition) or another task (the invalid condition). The overall effect of task expectancy was about 90 ms. McCann et al. hypothesized that a task-set reconfiguration

stage needs to be inserted into the processing stream on unexpected trials, but not on expected trials. This inserted task-set configuration stage begins once perceptual processes have revealed that the current stimulus does not belong to the expected task, but perhaps before difficult central operations (e.g. response selection) can begin.

Meiran (1996) has provided further evidence for the benefits of deliberate task preparation. By presenting participants with a task cue (100% valid) either early or late during a fixed inter-trial interval, he was able to vary preparation time without varying the time elapsed since the previous trial. Meiran found that increases in preparation time decreased the switch cost. This result indicates that the switch cost is reduced by an active preparatory process (e.g. a task-set configuration) during the inter-trial interval, not just a passive decay of interfering representations from the previous trial.

Bottom-up Influences: Task Repetition

Although deliberate task preparation during the inter-trial interval can greatly improve performance, it is not necessarily sufficient to eliminate the switch cost. Rogers and Monsell (1995), for example, found that task switches were performed relatively slowly even when ample time was available to prepare for the upcoming task. This component of the switch cost not eliminated by deliberate preparation has become known as the “residual” switch cost. The existence of a residual switch cost suggests that there is something about actually performing a task that – in a bottom-up rather than top-down fashion – readies critical mental mechanisms to rapidly perform that task again in the near future.

What is the specific cause of the residual switch cost? Presumably, on task repetition trials the proper task set remains

in place from the previous trial, whereas on task switch trials the task set must be reconfigured. Rogers and Monsell (1995) proposed that this reconfiguration on task switch trials cannot be completed (despite ample preparation time) until after the actual stimulus has been presented. Thus some critical stage of mental processing (i.e. a stage that requires advance preparation) might be postponed until the online task-set reconfiguration can be completed. Alternatively, it is possible that the critical stages of mental processing are prepared well enough to begin immediately on task switch trials but, being less highly activated, they proceed relatively slowly.

As will be seen, the present data indirectly support the latter account. Note that on this view it is not clear whether the effect of task repetition is best thought of as a switch cost or a repetition benefit (e.g. recent task performance primes critical mental mechanisms to perform that task again). Partly for this reason and partly because the term switch cost might carry unwanted theoretical baggage, we chose to avoid this particular terminology altogether. Instead, we refer to any difference between the repetition and switch conditions simply as the “task repetition effect.”

Goals of the Present Study

Previous studies of task switching have revealed the importance of both top-down factors (e.g. task expectancy; Meiran, 1996; McCann et al., 2000) and bottom-up factors (e.g. task recency; Rogers and Monsell, 1995). The main goal of the present research was to determine how these top-down and bottom-up factors combine to produce an overall level of task-readiness. For instance, do both types of factors have essentially the same effect on cognitive processing (e.g. both influence the activation level of a task set)? Or does each factor influence a completely separate component of the cognitive architecture?

The basic approach employed in this paper was to factorially manipulate task expectancy (by comparing validly and invalidly

cued tasks) and task recency (by comparing task switches and task repetitions) and then examine their interaction on response time (RT)¹. Our experimental paradigm was a marriage of the alternating runs paradigm of Rogers and Monsell (1995) and the task-cuing validity paradigm of McCann et al. (2000). As in Rogers and Monsell, participants generally performed the tasks in runs of two (e.g. AABBAABB...). The inter-trial interval was relatively long (2.5 sec) so that participants could fully prepare for the upcoming trial; hence, any residual effects of task repetition should be due to the bottom-up effects of having recently performed the task. Following McCann et al. (2000), we occasionally violated the standard task sequence by presenting a task that participants were not expecting.

This design contains all four conditions resulting from a factorial manipulation of task expectancy and task repetition: expected-repetitions (e.g. ...BBAABB), expected-switches (e.g. ...BBAAB), unexpected-repetitions (e.g. ...AABBB), and unexpected-switches (e.g. ...AABBAB). The expected-switch and unexpected-repetition conditions are of special interest, because they pit the effects of expectancy against the effects of repetition.

Theory and Expected Outcomes

Although previous task-switching studies have not directly examined the interaction between task expectancy and task repetition², researchers in other domains have studied the interaction of expectancy and recency in an effort to better understand the interplay between top-down and bottom-up factors. Epstein and Rock (1960), for example, examined the roles of stimulus expectancy and stimulus recency in the perception of ambiguous figures. Their materials consisted of an ambiguous figure (A), along with two unambiguous figures (X and Y) corresponding to the two primary

interpretations of the ambiguous figure. Participants viewed an alternating sequence of the two unambiguous figures terminated by one presentation of the ambiguous figure (e.g. XYXYXYA). Epstein and Rock found that participants' interpretations of the ambiguous figure (A) corresponded to the figure they saw most recently (Y), not the figure they expected to see (X). This finding has been interpreted as evidence that perceptual organization is determined more strongly by bottom-up influences than by top-down influences.

In a set of semantic priming experiments, Neely (1977) presented a prime stimulus followed by a probe stimulus, to which subjects made a rapid lexical decision (word vs. nonword). Half of the probes were related to the prime category and half were not. To manipulate category expectancy, Neely told participants to expect a bird name following the prime "BIRD" but to expect a body-part name following the prime "BUILDING." Thus the probe stimulus could be related to the primed category only (e.g. BUILDING followed by HOUSE), to the expected category only (e.g. BUILDING followed by ARM), to both primed and expected categories (e.g. BIRD followed by ROBIN), or to neither (BIRD followed by ARM). The basic question was whether the activation of the lexical system would be determined more by top-down factors (i.e. category expectancy) or by bottom-up factors (category repetition/priming). When the delay between the prime and probe words was long, both category expectancy and category repetition had substantial facilitatory effects. When the delay was short, however, only category repetition had a substantial effect. This pattern of results led Neely to conclude that the prime produces a fast-developing, automatic, bottom-up spreading of activation to semantically related words, whereas expectancy produces a slow-developing, deliberate, top-down activation of words in the expected category (see also Balota, 1983; Balota, Black, and Cheney, 1992).

These previous studies of the relation between expectancy and recency in other domains reached quite different conclusions. Whereas Epstein and Rock (1960) concluded that recency dominates expectancy in the perception of ambiguous figures, Neely (1977) concluded that expectancy and recency have independent effects on semantic activation. Thus, these conclusions provide little guidance regarding what outcome to expect in the domain of task preparation. Within this domain, there is already ample evidence that both task expectancy (Meiran, 1996; McCann et al., 2000; Rogers & Monsell, 1995; Sudevan & Taylor, 1987) and task repetition (Rogers & Monsell, 1995) can have substantial effects on response time. The main question, therefore, is how the effects of these top-down and bottom-up factors will interact. Below we describe several a priori plausible models and derive predictions from them for the interaction between expectancy and repetition on RT. These candidate models can be conveniently classified into three groups based on whether they predict underadditivity, overadditivity, or additivity.

Underadditivity. One very plausible hypothesis is that participants deliberately suppress the heightened activations resulting from a recently performed task if that task is not expected to occur on the upcoming trial. We refer to this as the Activation Suppression model. As an example, after completing a run of Task A trials participants might suppress the activations associated with that task in order to enhance the relative accessibility of the stimulus-response mapping for Task B. According to this model, participants should benefit from task repetition when repetition is expected, but not when repetition is unexpected. Consequently, this model predicts an underadditive interaction between

expectancy and repetition on RT, such as the one shown in Figure 1a.

An underadditive interaction is also predicted by certain models that attribute task-repetition effects and expectancy effects to a process of task-set reconfiguration. Rogers and Monsell (1995), for instance, hypothesized that reconfiguration for a task switch can be completed only when prompted by recognition of a stimulus associated with that task. On this view, some critical stage of processing is delayed while the reconfiguration takes place, resulting in a task-repetition effect (i.e. a switch cost). Similarly, McCann et al. (2000) proposed that when participants receive a task they are not expecting, they must complete an online task-set reconfiguration. Combining these two hypotheses, we arrive at a model where task-set reconfiguration is required whenever the current task is either (a) not repeated or (b) not expected. According to the Dual-Purpose Configuration model, the very same reconfiguration process is used in both cases. On this view, a task that is both not repeated and not expected (an unexpected-switch) should be performed about as fast as a task that is only not repeated (an expected-switch) or only not expected (an unexpected-repetition). Consequently, this model predicts an underadditive interaction between expectancy and repetition.

Overadditivity. Other plausible models make essentially the opposite prediction, an overadditive interaction between task expectancy and task recency. Perhaps the most obvious examples are models in which expectancy and repetition influence the same stage of mental processing (say, response selection). Suppose, for instance, that expectancy and repetition independently modulate the rate constant of that stage. On unexpected trials the processing rate would be reduced by a proportionality factor K_e relative to expected trials, and on switch trials the processing rate would be reduced by a proportionality factor K_r relative to repetition

trials. Consequently, on unexpected-switch trials the processing rate would be reduced by $K_e * K_r$ relative to expected-repetition trials. This Multiplicative model predicts that the effects of expectancy and repetition will combine overadditively, as shown in Figure 1b.

Overadditivity could also occur if task expectancy and task repetition provide two different routes to achieving roughly the same state of task readiness. We refer to this as the “Redundant Routes” model (see Marcel and Forrin, 1974, for a related model of category repetition effects). It is quite plausible that the primary effect of both expectancy and repetition is to increase the activation of the relevant stimulus-response mappings relative to other stimulus-response mappings. Provided that either route can produce virtually the full amount of activation possible, it should not help much if the task is activated through both routes (e.g. an expected-repetition) rather than just one. Consequently, the effects of expectancy and repetition should interact overadditively (see Figure 1b).

Additivity. Additive interactions between task expectancy and task repetition (see Figure 1c) could occur if expectancy and repetition selectively influence distinct stages of mental processing arranged in a serial architecture³ (see Sternberg, 1969 and Roberts and Sternberg, 1993). We refer to any such model as a Separate-Stages model. One appealing, specific separate-stage model asserts that task expectancy affects the time required to prepare upcoming central mental operations (i.e. task-set reconfiguration), whereas task repetition affects the time required to execute those central operations once they have been programmed. We refer to this subcase, shown in Figure 2, as the Configuration-Execution model.

Summary. The preceding discussion, although certainly not

exhaustive, covers the models of executive task control that, a priori, we considered to be the most plausible. As it turns out, these models predict very different patterns of interaction between task expectancy and task repetition, ranging from underadditive to additive to overadditive. Therefore, no matter what type of interaction is observed, we should be able to make considerable progress in pruning the space of viable theories.

Experiment 1

To measure the interaction between task expectancy and task repetition, we used a modified version of Rogers and Monsell's (1995) alternating-runs paradigm. Stimuli were presented in a 2X2 grid of cells, like the one shown in Figure 3a. The first stimulus of each block was presented in the upper-left cell; thereafter, each stimulus appeared in the cell located immediately clockwise from the previous stimulus. For half the participants, a stimulus appearing in one of the upper two cells was usually a colored rectangle, while a stimulus appearing in one of the lower two cells was usually a letter. The assignment of tasks to cells was reversed for the other half of the participants. Note that there were two redundant cues as to which task (color or letter) was likely to be performed on the upcoming trial: (1) the repeating task sequence and (2) the location of the stimulus within the 2X2 grid.

Our method differed from that of Rogers and Monsell (1995) in that on a modest proportion of trials (13.3%) we presented a task that the participant was not expecting. In these cases, the stimulus still appeared in the expected location (i.e. immediately clockwise from the previous stimulus location), but the stimulus did not belong to the task that was expected to occur in that location. These unexpected trials could occur when participants were anticipating a task switch and when they were anticipating a task repetition. For example, we occasionally presented task sequences such as AABBAB where the final trial is an unexpected task-switch, and sequences such as AABBB where

the final trial is an unexpected task-repetition. This design thus produces not only expected-switches and expected-repetitions, but also unexpected-switches and unexpected-repetitions. In other words, the design contains all four conditions resulting from a 2X2 factorial combination of task expectancy and task repetition.

The experiment consisted of two practice blocks followed by eight experimental blocks. Each of the experimental blocks consisted of 8 unexpected-task trials and 52 expected-task trials, resulting in an overall task cueing validity of 86.7% in those blocks. Several efforts were made to encourage participants to establish a strong task expectancy. First, participants were explicitly instructed to prepare for the expected task. Second, the expected task was presented on every trial of the two practice blocks as well as the four warm-up trials of the subsequent experimental blocks. Third, every presentation of an unexpected task in the experimental blocks was followed by at least two consecutive additional presentations of the expected task. These "recovery" trials were not analyzed; they served only to reinstate the proper task expectancy.

As noted earlier, we chose our tasks so that each stimulus had affordances for only one of the two tasks. One task was to identify a letter (I, S, O, or X) and the other was to determine the color of a rectangular patch (red, green, blue, yellow). Not only do these tasks involve distinct stimulus sets, but they also require clearly distinct mental operations. Therefore, there was little potential for the mental operations of one task to be mistakenly applied to the other task. The two tasks were designed to be approximately equal in difficulty, so that participants would not find it advantageous to selectively prepare for the same task (e.g. the more difficult one) on every trial regardless of whether it was expected or not.

Method

Participants. Thirty-two students from local colleges and universities received partial course credit or money in exchange for their participation. All reported normal or corrected-to-normal visual acuity and normal color vision. None participated in more than one of the experiments reported in this paper.

Stimuli. The stimulus for the letter task was an uppercase letter drawn randomly from the set {I, S, O, X}. The letters, which extended approximately 0.5 by 0.4 degrees from a typical viewing distance of 60 cm, were presented in white on a black background. The stimulus for the color task was a filled rectangle drawn in red, green, blue, or yellow. The colored rectangle extended approximately 0.4 by 0.5 degrees. Each stimulus was presented within one of four squares arranged in a 2X2 grid (see Figure 3a). Each square subtended approximately 1.0 by 1.0 degrees.

Apparatus. Stimulus presentation and response timing were performed by IBM compatible computers connected to NEC Multisynch monitors.

Procedure. Participants were told about the assignments of tasks to cells of the 2X2 grid. In addition, during the practice blocks the words COLOR and LETTER were presented next to the cells that usually contained the color and letter stimuli, respectively. Participants were instructed to always prepare for task most likely to occur in the upcoming trial.

Responses to both tasks were made using the 'M' and '<' keys. The letters I and S were assigned to one of these response keys, and the letters O and X were assigned to the other response key. Similarly, the colors red and green were assigned to one of these response keys and the colors blue and yellow were assigned to the other response key. The assignments were counterbalanced across participants. Note that because the same set of response keys was used for both tasks, participants could not use task expectancy to selectively prepare one set of effectors.

The sequence of events within a trial was as follows. First, a fixation marker ('+') was presented for 1000 ms in the cell containing the upcoming stimulus. The stimulus was presented 500 ms later and remained present until the participant responded to it. If the response was incorrect, the word "Incorrect" was displayed for 200 ms just outside the cell containing the stimulus. The next trial began 1000 ms later. Thus, a total of about 2.5 seconds elapsed between the response to one stimuli and the onset of the next stimulus. This presumably gave participants sufficient time to prepare for the upcoming task.

The two practice blocks contained no unexpected tasks; the subsequent eight experimental blocks each contained eight unexpected trials each. The unexpected trials were distributed randomly within each block of 60 trials, with the restrictions that the first four trials (warm-ups) always consisted of expected tasks and that each unexpected task be followed by two recovery trials with the expected task. At the end of each block, the computer provided feedback on the speed and accuracy of each task and participants were allowed to take a short break.

Analysis. Separate ANOVAs (task X expectancy X repetition) were conducted on RT and percent correct. The two "recovery" trials following each unexpected task-trial were omitted from these analyses. Also eliminated were trials in which the stimulus was identical to the stimulus shown the last time that task was performed (regardless of how many trials of the other task intervened). By eliminating exact stimulus repetitions we were able to directly measure the effect of task repetition, unconfounded by the effects of stimulus repetition. Furthermore, incorrect responses were omitted from the RT analyses. In addition, any RTs less than 200 ms or greater than

2000 ms (< 1% of trials in all experiments) were excluded as outliers. All analyses used an alpha level of 0.05.

Results and Discussion

Error Rates. Participants made 3.4 % errors. Participants made more errors to task switches (4.1 %) than to task repetitions (2.7 %), $F(1,31) = 5.9, p < 0.05$.

Response Time. Mean RT to the color task (655 ms) did not differ significantly from mean RT to the letter task (640 ms), $F(1,31) = 2.0, p > 0.1$. Thus, we were successful in choosing tasks that were roughly equal in difficulty. Furthermore, task type did not interact significantly with other experimental factors of interest (i.e. expectancy and repetition). Mean RT to expected tasks (633 ms) was faster than mean RT to unexpected tasks (662 ms), $F(1,31) = 14.4, p < 0.01$. Mean RT to task repetitions (633 ms) was faster than mean RT to tasks switches (662 ms), $F(1,31) = 27.4, p < 0.001$. Note that the main effect of task expectancy (29 ms) was roughly equal to the main effect of task repetition (31 ms).

Of primary interest is the interaction between expectancy and repetition. As shown in Figure 4, these variables produced a significant underadditive interaction on RT, $F(1,31) = 6.2, p < 0.05$; the effect of expectancy was 39 ms for task repetitions but only 18 ms for task switches. Before interpreting this interaction, however, we will first attempt to replicate it.

Experiment 2

(Four Stimuli and Four Responses)

The primary goal of Experiment 2 was to replicate the results of Experiment 1 and ensure that those results were not due to any special properties of the response mappings. Whereas in Experiment 1 the four stimuli from each task were mapped onto two different responses (e.g. participants pressed one key for the letters 'I' or 'S' and another for the letters 'O' or 'X'), in Experiment 2 the four stimuli were mapped onto four separate response keys.

One specific motivation for changing the response mappings was to evaluate the hypothesis that participants are sometimes able to bypass the normal response selection process when there are only two possible responses (as in Experiment 1). For example, a participant might compare the current stimulus category to the one encountered on the previous trial. If the category is the same, then the participant would simply initiate the same response; if the category is different, then the participant would initiate the opposite response. Because this comparison is based only on the memory of the previous trial, we refer to this possibility as the "cache memory" hypothesis (on analogy to computer memory).

Use of this shortcut strategy might substantially reduce RT, provided that comparisons with the previous trial can be performed faster than the usual process of response selection. However, this strategy would only be effective on task-repetition trials; on task-switch trials comparisons with the previous stimulus category would always yield a mismatch. Thus, use of this shortcut strategy could account for the task-repetition effect observed in Experiment 1. Note that on unexpected-repetition trials participants were not anticipating a task repetition and therefore might not have stored the information from the previous trial necessary to use the cache-memory strategy. Thus, this hypothesis can also explain the underadditive interaction between the effects of task expectancy and task repetition observed in Experiment 1.

Due to the use of four stimulus and response categories in Experiment 2, however, this cache-memory strategy should have been much less effective. When the exact same stimulus repeated (25% of all task-repetition trials), participants could have known to repeat the same response. These stimulus repetition trials were

excluded from our main analyses, however, so a speed-up on these trials would not have influenced the measured task-repetition effect. On the remaining trials (75%) the stimulus did not repeat; a comparison of the current stimulus category with that of the previous trial would always yield a mismatch. Participants would have known to change their response, but that knowledge would not have been very useful since there would still be three responses left in the running. Thus, the measured task-repetition effect in Experiment 2 should have been affected very little by use of a cache memory strategy. If we nevertheless find a substantial task-repetition effect, similar to that of Experiment 1, then we can reject the cache memory hypothesis and focus on other possibilities.

Method

Except where noted, the method of Experiment 2 was identical to that of Experiment 1.

Participants. Thirty-two students participated.

Procedure. The four stimulus letters were mapped onto the four response keys {M, <, >, /} in the order {ISOX} or the order {OXIS}. The four stimulus colors were mapped onto the same response keys in the order {red, green, blue, yellow} or the order {blue, yellow, red, green}. These stimulus-response mappings were counterbalanced across participants.

Results and Discussion

Error Rates. Participants made 4.0 % errors. Participants made more errors to the color task (5.0 %) than the letter task (3.0 %), $F(1,31) = 5.5, p < 0.05$.

Response Time. Mean RT in Experiment 2 (778 ms) was greater than the mean RT found in Experiment 1 (648 ms). Because the main difference between experiments was the number of different stimulus-response mappings, the increase in RT can presumably be attributed primarily to an increase in response-selection time.

Despite the difference in overall response times between experiments, the pattern of results was very similar. Once again, mean RT to the color task (787 ms) did not differ significantly from mean RT to the letter task (769 ms), $F(1,31) = 1.2, p > 0.1$. Task type also did not interact with task expectancy or task repetition. Mean RT to expected tasks (766 ms) was 24 ms faster than mean RT to unexpected tasks (790 ms), $F(1,31) = 12.7, p < 0.01$. Mean RT to task repetitions (755) was 46 ms faster than mean RT to task switches (801 ms), $F(1,31) = 34.4, p < 0.001$.

One specific goal of Experiment 2 was to determine whether the cache-memory strategy can account for the task-repetition effects we found in Experiment 1. Thus, we used a design in which participants should not have benefited from the use of cache-memory. Experiment 2 nevertheless produced substantial effects of task repetition (46 ms) that were even larger than those found in Experiment 1 (31 ms). It appears that something other than the use of a cache-memory strategy is responsible for the task-repetition effect.

The interaction between task expectancy and task repetition, shown in Figure 5, did not reach significance in this experiment, $F(1,31) = 1.0, p > 0.1$. However, the data showed a trend towards underadditivity similar to that observed in Experiment 1; the effect of expectancy was 31 ms for task repetitions but only 17 ms for task switches. As will be seen, an underadditive data pattern was also observed in Experiment 3.

Experiment 3

(Three Tasks)

The primary goal of Experiment 3 was to determine whether the results from Experiments 1 and 2 generalize to task-switching scenarios with more than two tasks. For example, it is possible that participants can simultaneously maintain at

most two task-sets in short-term memory. If so, the present three-task experiment might produce a qualitatively different pattern of results. In addition, two-task experiments might be a special case, because when switching away from one task there is no choice but to switch back to the other task; in other words, a task switch always involves returning to the only other task. With three possible tasks, however, a task switch requires that participants choose which of the remaining tasks to perform. This type of task-switching might be more representative of real-world task-switching scenarios, where it is often necessary to select from among several possible tasks.

The use of three tasks, in sequences such as AABBC, also allows us to examine a type of unexpected-switch condition not present in Experiments 1 and 2. In those experiments, unexpected-switches were always trials in which participants anticipated a task repetition but instead received a task switch (e.g. AABA). The present three-task experiment includes this type of unexpected-switch condition, but also adds a different type of unexpected-switch condition (e.g. AABBA) in which participants anticipated a switch to one task (C) but instead switched to a different task (A). These two types of unexpected-switch conditions might require different types of task-set disengagement. In the former type of unexpected-switch condition (e.g., AABA or AABC), the task-set to be disengaged has been performed recently and therefore might be well instantiated. In the latter type of unexpected-switch condition (e.g. AABBA), however, the task-set to be disengaged (Task C) has not been performed recently and therefore might not be well instantiated. A comparison of these two types of unexpected-switch conditions, therefore, should tell us whether the task-repetition effect depends on the difficulty of task-set disengagement.

Method

Except where noted, the method of Experiment 3 was identical to that of Experiment 1.

Participants. Forty-eight students participated.

Procedure. In addition to the color and letter tasks, this experiment included a dots task. Either 2, 3, 4, or 5 dots were presented, arranged as they appear on standard dice. Half of the participants were instructed to press the 'M' key if the total number of dots was even, and to press the '<' key if the total was odd. The remaining participants were given the opposite instructions.

Because there were three tasks, the stimulus grid consisted of six cells (two for each task) arranged as shown in Figure 3b. As in Experiments 1-2, the stimuli rotated through the six cells in the clockwise direction.

The two practice blocks and twelve experimental blocks contained 56 trials each (6 warm-ups plus 50 experimental trials). As in the previous experiments, the two practice blocks contained no unexpected trials. The subsequent blocks contained 8 unexpected trials and 48 expected trials, for an overall task-cue validity of 85.7% (only slightly lower than the task-cue validity in Experiments 1 and 2, which was 86.7%).

Results and Discussion

Error Rates. Participants made 3.6 % errors. An analysis of error rates revealed no significant effects.

Response Time. Mean RTs to the color task (683 ms), letter task (684 ms), and dots task (706 ms) did not differ significantly, $F(2,94) = 2.4, p > 0.05$. As with all the other experiments reported in this paper, the effects of task did not interact with the effects of expectancy or repetition. Mean RT to expected tasks (666 ms) was 49 ms faster than mean RT to unexpected tasks (715 ms), $F(1,47) = 60.5, p < 0.001$. Mean RT to task repetitions (663 ms) was 56 ms faster than mean RT to task switches (719 ms), $F(1,47) = 111.1, p < 0.001$.

Mean unexpected-switch RT was 737 ms when participants anticipated a task repetition (i.e. when the prepared task had been performed recently) and was 739 ms when participants anticipated a task switch (i.e. when the prepared task had not been performed recently); this difference was not significant, $F < 1$. Thus, the present data provide no evidence that the task-repetition effect depends on the difficulty of task-set disengagement (i.e. from a task-set that has recently been carried out versus from a task-set that has not recently been carried out); in fact, the slight trend went in the direction opposite to that predicted. This tentative finding suggests either that the duration of task-set disengagement does not depend whether the to-be-disengaged task was performed recently, or that performance of an unexpected task simply does not require any task-set disengagement. The latter hypothesis is plausible; because our tasks were very distinct, it might have been unnecessary to disengage one task set in order to perform another.

The interaction between task expectancy and task repetition, shown in Figure 6, again showed an underadditive trend; the effect of expectancy was 61 ms for task repetitions and 51 ms for task switches. Although this underadditive trend was not significant, $F(1,47) < 1.2$, $p > 0.05$, it is consistent with the results of Experiments 1-3. Experiment 4 will attempt to replicate this effect and determine its generality.

Experiment 4

(Neutral Trials versus 100 % Valid Trials)

To further test the generality of the observed underadditivity between task expectancy and task repetition, Experiment 4 probed a different range of levels of expectancy. Whereas Experiments 1-3 compared the performance of expected tasks (85.7-86.7% likely to occur, depending on the experiment) and unexpected tasks (13.3-14.3% likely to occur), Experiment 4 compared performance on tasks that were 100% likely to occur (certain

tasks) to performance on tasks that were 50% likely to occur (neutral tasks).

Experiment 4 used the same pair of tasks and the same 2X2 stimulus grid used in Experiments 1-2. However, the assignment of tasks to cells was different, as shown in Figure 3c. The upper-left cell was assigned to one of the two tasks (the color task for half of the participants and the letter task for the other half) and the lower-right cell was assigned to the other task. The same task was presented every time a stimulus appeared in one of those two cells (i.e. 100% cue validity). In the remaining cells (upper-right and lower-left) the task was chosen randomly on each trial.

One specific goal of this experiment was to determine whether the previous results occurred only because the unexpected condition was a special case. Because the unexpected task was unlikely to occur, it might have produced an initial “surprise” reaction, during which no progress was made on the task. The neutral condition of the present experiment was unlikely to generate a surprise reaction, so it should now be possible to measure the effects of expectancy without contamination from the effects of surprise.

The present design had one additional advantage over that used in Experiments 1-3. In those experiments, participants might have been unwilling to fully commit to the expected task if doing so resulted in extremely poor performance on the unexpected-task trials. This lack of commitment to the expected task could have greatly weakened the effects of the expectancy manipulation. In Experiment 4, however, the expected (i.e. “certain”) task was presented on 100% of the trials in the upper-left and lower-right cells, encouraging participants to prepare for that task exclusively.

Method

Except where noted, the method of Experiment 4 was identical to that of Experiment 1.

Participants. Thirty-two students participated.

Procedure. As in Experiment 1, the four stimuli of each task were mapped onto two responses. However, the assignment of tasks to the four cells of the 2X2 stimulus grid was different. As shown in Figure 3c, the upper-left cell was assigned to one task (color or letter) and the lower-right cell was assigned to the other task (counterbalanced across participants). The validity of this cue was 100 % (i.e. there were no unexpected task presentations in those cells). In the other two cells the task was selected at random, with the restriction that half of these trials within each block be task switches and half be task repetitions. Participants were instructed to maintain a neutral task expectancy on these trials.

Results and Discussion

Error Rates. Participants made 3.9 % errors. An analysis of error rates revealed no significant effects.

Response Time. Once again, the mean RT to the color task (610 ms) did not differ significantly from the mean RT to the letter task (624 ms), $F(1,31) = 2.0$, $p > 0.1$. Also, the type of task (color versus letter) did not interact with expectancy or repetition. Mean RT to expected tasks (605 ms) was 25 ms faster than mean RT to unexpected tasks (630 ms), $F(1,31) = 23.1$, $p < 0.001$. Mean RT to repeated tasks (603 ms) was 29 ms faster than mean RT to task-switches (632 ms), $F(1,31) = 85.0$, $p < 0.001$.

The interaction between task expectancy and task repetition, shown in Figure 7, was very close to additive, $F(1,31) < 1$; the effect of expectancy was 26 ms for task repetitions and 25 ms for task switches. Thus, the underadditive trend observed in Experiments 1-3 was not found here. The next section explores a possible explanation for the apparent discrepancy in results.

Task Repetition versus Task Recency: The Effects of Lag

The apparent empirical discrepancy between Experiment 4 and Experiments 1-3 prompted us to more closely examine our independent variables. It became readily apparent that the task repetition variable as traditionally employed (comparing switch trials versus repetition trials) does not perfectly capture the more abstract variable of theoretical interest, namely task recency. In brief, the task repetition variable is binary, whereas task recency varies continuously. As will be discussed next, this insight led to a principled reanalysis of the data using a more appropriate measure of task recency. The reanalysis indicated that there was no empirical discrepancy after all.

It is useful to consider in more fine-grained detail the nature of the task repetition and task switch trials. Task repetitions are trials in which the current task was also performed on the previous trial; thus, the task lag (the number of trials since the task was last performed) on task repetition trials is always equal to 1. Task switches, on the other hand, are trials in which the current task was not performed on the previous trial; thus, the task lag is always greater than 1. However, not all task switch trials have the same lag. On task-switch trials, the new task might have been performed two trials previously (a lag of 2), or it might have been performed five trials previously (a lag of 5). The binary task repetition variable, however, lumps all these switch trials together into the task-switch condition.

Variations in task lag might be harmless, if not for the fact that lag tends to be correlated with task expectancy. In Experiments 1-2, for example, expected switches (e.g. AABBA) could have a lag of 3 or 4, whereas unexpected switches (e.g. AABA) always had a lag of only 2. If RT

increases with increasing task lag, as might be expected, then this confound could have biased our results. Specifically, it would have artificially reduced the measured effect of expectancy on task switch trials (while having no effect on task repetition trials). The net result would be just the sort of underadditive interaction between task expectancy and task repetition that we observed in Experiments 1-3.

To determine whether RT increases with the task lag we conducted a control experiment. We presented the color and letter tasks in a random order, producing task lags as low as 1 and as high as 15 or more. Figure 8 shows how mean RT in the unexpected-switch condition varied with lag, plotted on a log scale. Note that we pooled data from adjacent lags where the data was relatively sparse. Two aspects of these data are important for the present purposes. First, note that lag had a substantial effect (about 100 ms). Second, mean RT increased roughly linearly with the log of the lag. This relationship reflects a “diminishing returns” pattern such that as the lag became larger the effect of an increment to the lag became smaller. As will be seen, a similar effect of lag was observed in the Experiments 1-4.

Given that RT increases with the task lag, it is clearly necessary to reanalyze the present data. Our approach was straightforward: instead of treating all task switch trials the same (as in our original task-repetition analyses), we coded them according to their task lag. In a sense, we simply replaced the binary task repetition variable with the continuous task recency variable, operationally defined as the task lag. Figure 9 shows the data from Experiments 1-4 plotted as a function of task lag. In this figure, task repetitions correspond to lag 1, while task switches correspond to lags greater than 1. To facilitate visual apprehension of the data patterns and provide a basis for interpolation between lags, we used a log transformation of the lag axis. As can be seen in Figure 9, this transformation linearized the

relation between RT and lag, just as it did in the control experiment described above.

The key question now becomes whether task lag (scaled logarithmically) and task expectancy have additive effects on RT. Because the expected and unexpected conditions have different sets of lags, the data cannot be analyzed using a simple ANOVA on mean RT. Instead, we used two alternative approaches. In the first approach, we determined the slopes relating RT to lag for both the expected and unexpected conditions. If the effects of recency and expectancy are additive, then these slopes should not differ. Indeed, ANOVAs showed that the difference in slope between the expected and unexpected conditions did not even approach significance in any of the experiments (all F 's < 1). In the second approach we fit⁴ the data with a model in which the effects of recency (i.e. $\log[\text{lag}]$) and expectancy are assumed to be additive (see the solid lines in Figure 9). Table 1 shows the corresponding parameter values. As can be seen from Figure 9, the deviations from the additive model were very small (generally less than 3 ms). Furthermore, there was no systematic trend toward either overadditivity or underadditivity. Thus, both approaches converge on the conclusion that recency and expectancy had additive effects.

An important consequence of using task lag as a measure of task recency is that it removes the discrepancy in outcomes between Experiments 1-3 and Experiment 4⁵. Despite several significant methodological differences (in the number of responses, the number of tasks, and the levels of expectancy), all four experiments now show additivity. The remarkable consistency of these results adds to our confidence that the reanalysis was appropriate and necessary. The theoretical implications of the additive relationship between task expectancy and task recency

will be considered in the General Discussion.

General Discussion

Task-switching studies have shown that performance depends upon both top-down influences (e.g. task expectancy; Meiran, 1996; McCann et al., 2000) and bottom-up influences (e.g. task repetition; Rogers and Monsell, 1995). The purpose of the present paper was to determine how these top-down and bottom-up factors interact. In the Introduction we outlined a wide variety of candidate hypotheses, all of which had high a priori plausibility. To evaluate these hypotheses, we measured the interaction between task expectancy and task recency using a combination of the alternating-runs paradigm of Rogers and Monsell (1995) and the task-cuing paradigm of McCann et al. (2000).

In Experiment 1, participants performed two different tasks (A and B), generally in runs of two (AABBAABB...). We occasionally violated this task sequence by presenting the task that participants were not expecting, resulting in either an unexpected task repetition (e.g. ...AABBB) or an unexpected task switch (e.g. ...AABA). This design allowed us to measure the interaction between task expectancy and task repetition. Experiment 2 was similar to Experiment 1, except that each task required a more difficult stimulus-response mapping. Experiment 3 involved three tasks instead of just two. Experiment 4 explored a different range of levels of expectancy. Despite these methodological differences, all four experiments (when analyzed properly) showed the same basic result – an additive interaction between task expectancy and task recency (discussed in more detail below).

The tasks studied in this paper were relatively simple and had distinct stimulus sets. Furthermore, only one stimulus was presented on each trial. Thus – as is the case in many real-world task-switching scenarios – there was little potential for participants to perform the wrong task on any given stimulus. We refer to this as a “single-affordance” condition, because each

stimulus had affordances for only one task or the other. As a consequence, participants presumably never had to inhibit processing of the inappropriate task. Also, both tasks used the same response buttons; thus, there was no opportunity for participants to benefit from task expectancy by priming one set of effectors over another. Nevertheless, we observed substantial main effects of both task expectancy and task recency in all four experiments. The difference in RT between expected tasks and unexpected tasks ranged from 24 to 49 ms across experiments. The difference in RT between task repetitions and task switches (often called the residual switch cost) ranged from 29 to 56 ms. Because all exact stimulus repetition trials were removed from the principal analyses, the observed task-repetition effect cannot be attributed to the benefits of processing the exact same stimulus on two trials in a row. Furthermore, the task-repetition effect was large even when four stimuli were mapped onto four separate responses (Experiment 2), so the task-repetition effect also cannot be attributed to the use of a “cache memory” strategy – bypassing response selection and instead responding based solely on whether the current stimulus category does or does not match that of the previous trial (see Experiment 2 for more discussion of this argument).

Stimulus Repetition

To obtain a relatively pure measure of the effects of task repetition, our main analyses excluded exact stimulus repetitions. It is useful, however, to directly study stimulus repetition effects and determine how these effects interacted with those of task expectancy. Separate ANOVAs conducted on each experiment revealed that mean RT to task repetitions was faster when the stimulus repeated than when the stimulus changed ($p < 0.01$ in each experiment). The average effect of a stimulus repetition, across all four experiments, was 83 ms. The

stimulus repetition effect was about as large for expected tasks (85 ms) as it was for unexpected tasks (81 ms). Thus, the effects of stimulus repetition were approximately additive with the effects of task expectancy.

The observed additivity between stimulus repetition and task expectancy indicates that participants benefited from stimulus repetitions even when they had no reason to expect the stimulus to repeat; therefore, the effects of stimulus repetition may be in some sense “automatic.” According to the additive factor method (Sternberg, 1969), factors that affect the same stage should interact and those that affect distinct discrete stages should have additive effects. Therefore, the additivity between stimulus repetition and expectancy also suggests that these factors do not influence the same stage(s) of mental processing. Thus, by determining which stage is affected by stimulus repetition (a variable about which quite a bit is already known), we can then rule out that stage as the likely locus for task expectancy effects. Previous studies have generally shown that the effects of stimulus repetition interact with the effects of factors that influence the duration of response selection (e.g. Bertelson, 1963; Kornblum, 1969), but not with the effects of factors that influence the duration of perceptual processing stages (e.g. Hansen & Well, 1984). The straightforward interpretation of these results is that stimulus repetition primarily influences response selection (see Pashler & Johnston, 1989); it follows that task expectancy should not influence response selection (given the aforementioned additive interaction between these factors). This conclusion is interesting in its own right, and will come in handy later when we attempt to interpret our main result (additivity between the effects of task expectancy and task recency).

Response Repetition

Do participants respond more rapidly when the required physical response matches the one from the previous trial? Because each

of our tasks used the same response keys, the effects of response repetition can be measured on task-switch trials as well as task-repetition trials. We observed a consistent, small cost of response repetition, both on task switch trials (13 ms) and on task repetition trials (10 ms).⁶ Rogers and Monsell (1995) also observed a small cost of response repetition on task-switch trials (15 ms); however they observed a substantial benefit of response repetition on task-repetition trials (52 ms).

This empirical discrepancy might stem from differences in the tasks being studied. In Rogers and Monsell (1995), each stimulus was assigned to a response based on its membership in a well-learned category (e.g. odd or even numbers). It seems likely that participants first categorized the stimulus, then selected a response based on this mediating categorical representation. The benefit for response repetitions reported by Rogers and Monsell might therefore reflect heightened activation at the categorical level that was left over from the previous trial. In contrast, our stimuli were assigned more-or-less randomly to response keys. In Experiment 1, for example, the letters ‘I’ and ‘S’ were mapped onto one response key while ‘O’ and ‘X’ were mapped onto the other response key. Because the pair of stimuli mapped onto a common response did not belong to the same well-learned, response-relevant category, participants might not have used categorical membership (e.g. the “I/S” category versus the “O/X” category) as a mediating representation. Instead, they might have directly mapped each stimulus onto the corresponding response. As a result, participants should not have benefited from heightened activation at the categorical level in our study. Consistent with the preceding arguments, response repetition in single-task studies benefits have usually been found only when the two stimuli

belong to the same well-learned category and that category is relevant to the response (see Campbell & Proctor, 1993 and Pashler & Baylis, 1991b).

The preceding argument can explain why there was no benefit of physical response repetition, but it needs to be supplemented in order to explain why there was actually a small cost of physical response repetition (~12 ms overall). This cost might have resulted from an overall bias to repeat a response whenever the stimulus repeated and to change a response whenever the stimulus changed. Such a response bias would have been helpful on most trials, and thus could have reduced overall RT. This bias would have been especially useful on trials where the required response changed, because all such trials necessarily involved a stimulus change as well. However, this bias should not have been useful on trials where the required response repeated; specifically, the bias would have been unhelpful on the subset of trials where the response repeated even though the stimulus changed (e.g. when the stimulus 'S' was followed by the stimulus 'I'). Other things being equal, therefore, this response bias should have slightly elevated RT in the response-repetition condition relative to the response-change condition. For more discussion of response-repetition effects, see Campbell & Proctor (1993), Kleinsorge (1999), Kleinsorge and Heuer (1999), and Pashler & Baylis (1991b).

The Interaction between Expectancy and Recency

The most important result of the present research is the test for the interaction between the effects of task expectancy and task repetition on RT. In Experiments 1-3 this interaction showed an underadditive trend, whereas in Experiment 4 the interaction was additive. These analyses are suspect, however, because the binary task repetition variable does not by itself fully capture the abstract variable of interest, namely task recency. Specifically, the task repetition variable lumps together all task

switch conditions, even though these conditions differ in terms of the number of trials since the task was last performed (i.e. task lag).

To deal with this problem, we reanalyzed our data using the continuous task-lag variable as our measure of task recency (see Figure 9). Note that lag 1 corresponds to the task-repetition condition, while lags greater than 1 correspond to task-switch conditions. The picture that emerged from this new analysis was strikingly simple. All of the present experiments, despite several key methodological differences, showed very close to equal effects of expectancy at all task lags. In other words, task recency and task expectancy had additive effects.

Implications of Factor Additivity

In the Introduction we outlined several candidate models of how top-down and bottom-up factors interact in task-switching paradigms. One model that seemed especially appealing, a priori, was the Activation Suppression model. According to this model, participants suppress the activation of a recently performed task if it is not expected to occur on the upcoming trial. This strategy would help to prevent interference from irrelevant tasks by increasing the relative accessibility of the task-set needed on the upcoming trial. A consequence of this model is that participants should benefit from task repetition when it is expected, but not when it is unexpected, resulting in an underadditive interaction between expectancy and repetition. The present data, however, showed consistent additivity between these factors, arguing against the Activation Suppression model⁷.

The data also argue against the "Dual-Purpose Configuration" model. According to this model, task-set reconfiguration is needed whenever the task is either not expected or not repeated. When

a task is both not expected and not repeated, it is assumed that participants still need to perform only a single (dual-purpose) task-set configuration. The straightforward prediction from this model is that the effects of task expectancy and task recency should interact underadditively, contrary to our results.

We also considered the Multiplicative model, in which both expectancy and recency influence the rate parameter of the same stage of mental processing. This model predicts an overadditive interaction between task expectancy and task recency, contrary to our results. The Redundant Routes model, which proposes that expectancy and recency are simply two different routes to achieving basically the same state of task preparation, also incorrectly predicts an overadditive interaction.

In summary, four of the five candidate models considered in the Introduction – which a priori seemed very plausible – appear to be ruled out by the present data. This constitutes substantial progress in pruning the space of a priori viable models. In the next section we discuss the only model of the five that is consistent with the present data, the Separate-Stages model.

Separate-Stages Model. The observed additivity between task expectancy and task recency supports the Separate-Stages model, which asserts that expectancy and recency selectively influence distinct stages of mental processing (see Sternberg, 1969). Exactly which stages are affected by task expectancy and task recency? Although the available data are not sufficient to definitively answer this question, they do provide several important clues. These clues, along with some general considerations, allow us to make an educated guess.

One general consideration is that for our tasks the most difficult and time-consuming process was response selection. The tasks required relatively simple, well-learned perceptual judgments (e.g. color and letter identification under nearly ideal viewing

conditions) and relatively simple output processes (button pressing). In contrast, the tasks required an arbitrary mapping of stimuli onto responses that participants had never performed before. It seems reasonable to hypothesize, therefore, that the response-selection stage will be relatively slow and will be sensitive to many experimental manipulations. For more evidence supporting this conclusion, see Pashler and Baylis (1991a), Pashler and Johnston (1989), Ruthruff, Johnston, and Van Selst (2001), and Van Selst, Ruthruff, and Johnston (1999).

This consideration, a priori, makes the response-selection stage the most plausible locus for task-repetition effects. In addition, task repetition can be thought of as a weak form of stimulus repetition, a variable that has previously been shown to primarily influence the response-selection stage (see Bertelson, 1963; Kornblum, 1969; Pashler & Johnston, 1989). This hypothesis is also supported by specific aspects of the present data. Experiment 2, which required an especially difficult response mapping, produced a larger task-repetition effect than did Experiment 1 (46 ms versus 31 ms). According to additive factor logic, this overadditive interaction supports the hypothesis that task repetition influences the same stage as response mapping difficulty. Because response mapping difficulty clearly influences the response-selection stage, the obvious conclusion is that task repetition also influences the response-selection stage.

Whereas the task repetition variable appears to influence the response-selection stage, there are reasons to believe that the task expectancy variable does not. First, the additivity between task expectancy and task recency indicates that these variables have different processing loci, so they should not both be assigned to the response-selection stage. Furthermore, as noted earlier, we observed an additive interaction between

task expectancy and stimulus repetition; because stimulus repetition appears to influence the stage of response selection, task expectancy should not influence that same stage. Moreover, if task expectancy influenced response selection, then expectancy effects should have been larger in Experiment 2 (difficult stimulus-response mapping) than in Experiment 1 (relatively easy stimulus-response mapping). Contrary to this prediction, the expectancy effect in Experiment 2 (35.6 ms; see Table 1) was nearly identical that of Experiment 1 (35.7 ms).

To summarize so far, the overall pattern of factor interactions paints a very consistent picture. The effects of task repetition, stimulus repetition, and response-mapping difficulty appear to interact with one another, suggesting that they all affect the same stage. The most obvious candidate stage is the response-selection stage. On the other hand, the effects of these three variables appear to be additive with the effects of task expectancy. Additive factor logic, therefore, indicates that task expectancy does not influence response selection.

What, then, is the locus of the task expectancy effect? One logical possibility is that task expectancy influences a perceptual stage, such as stimulus identification. This conjecture seems unlikely, however, because the present stimuli were always highly discriminable and were presented under nearly ideal viewing conditions. Even under challenging viewing conditions, Los (1999) found that expectations regarding perceptual categories had little effect on RT. A more appealing explanation, shown in Figure 2, is that participants perform a task-set reconfiguration when presented with an unexpected task (see McCann et al., 2000). This reconfiguration might involve loading the rules for the task to be performed, or it might instead simply involve moving a mental “pointer” to the memory location of the appropriate task set. Presumably, this reconfiguration process would occur sometime after participants determine that the

present stimulus does not belong to the expected task, but before they perform demanding central operations such as response selection.

Note that at this point there is no compelling reason to postulate a task-set disengagement operation prior to the task-set reconfiguration, at least not in single-affordance designs where there is little competition between tasks. Experiment 3 (the three-task design) provided evidence that participants could shift as easily from a task that they had just performed as from a task that they anticipated but had not just performed.

How long does the putative task-set reconfiguration take on unexpected trials? The reconfiguration time causes the expectancy effect, which was 25-50 ms. As will be explained, however, the size of the expectancy effect should greatly underestimate the true reconfiguration time. Presumably, the reconfiguration is triggered soon after stimulus onset and then operates in parallel with stimulus identification (as well as any other mental operations that must be completed before response selection can begin). This processing overlap is shown in Figure 2. The task-set reconfiguration will only delay RT on unexpected trials (producing a task expectancy effect) to the extent that it lasts longer than stimulus identification. Hence, the reconfiguration could be much larger than the expectancy effect of 25-50 ms, perhaps as large as 100-200 ms. Even so, one might argue that 100-200 ms is not enough time to complete a task-set configuration; other studies have shown that participants need a few hundred msec to optimally ready themselves for an upcoming task switch (e.g. Rogers & Monsell, 1995). However, the reconfiguration might take much less time when cued by an actual stimulus for that task (as in our unexpected condition) than it would when there is no

perceptual support (as in the typical task-switch condition used to measure reconfiguration time).

Putting these pieces of the puzzle together, we arrive at the hypothesis that task expectancy affects the time required to complete the programming of upcoming central mental operations, whereas task recency affects the time required to execute those central operations once they have been programmed (see Figure 3). This “Configuration-Execution” model provides a consistent and parsimonious account of a wide range of data. Further experiments are needed, however, to test this model more directly.

The Ready, Set, Go! model of Fagot (1994). After arriving at the Configuration-Execution model, we became aware of the “Ready, Set, Go!” model proposed in an unpublished doctoral dissertation by Fagot (1994)⁸. Fagot presented a series of 12 experiments designed, among other things, to divide the task-switching cost into its basic components. In particular, he tried to determine which components were reduced in size by deliberate task preparation and which were not. To explain his set of data, Fagot hypothesized that the response selection mechanism must be set for a particular task before it can begin. The setting of the mechanism can be accomplished during the inter-trial interval but cannot be interrupted once started. Meanwhile, the “readiness” of the response selection mechanism for a particular task depends on task recency, but does not depend on deliberate task preparation. Although Fagot’s model was primarily directed towards dual-affordance conditions, whereas our Configuration-Execution model was meant only to cover single-affordance conditions, they share certain core assertions. Both models emphasize a key distinction between “set” or “configuration” of a task (which is a function of expectancy) and “readiness” or “priming” of that task (which is a function of recency).

Bottom-up and Top-down Influences on Task Preparation

Having discussed how to explain the main results of our experiments, we now return to the more general issue of the interaction between top-down, executive control processes and bottom-up, stimulus-driven processes. In the present experiments, both top-down and bottom-up factors had consistent effects of roughly the same magnitude. In addition, when we pitted these effects against one another, we found that unexpected-repetitions were about as fast as expected-switches. More importantly, these factors had additive effects, indicating that they influence separate components of the cognitive architecture. Specifically, the data are consistent with the hypothesis that top-down control serves to set-up, or “program,” central operations (e.g. response selection), whereas bottom-up factors serve to modulate the rate at which these central operations proceed once they have been programmed (e.g. via priming).

Folk, Remington, and Johnston (1992) reached a related conclusion regarding top-down control in a spatial cueing paradigm. In their experiments, probe displays were preceded by a spatial pre-cue that was never positively correlated with the position of the target (so that it was not in the participant’s interest to attend to the location of the pre-cue). Folk et al. (1992) found that the pre-cues were nevertheless capable of automatically capturing attention, but only when the properties of the pre-cue matched the properties of the target stimulus. For example, if the task was to identify a green target, then a green pre-cue captured attention, but a sudden onset stimulus did not; if the task was to identify a sudden-onset target, then a sudden-onset pre-cue captured attention but a green pre-cue did not. Folk et al. (1992) concluded that the

role of top-down control was to pre-program the attention control system prior to stimulus presentation. They argued that once the stimulus was presented, the attention-control system operated in a stimulus-driven, bottom-up way, with no direct online intervention from top-down processes.

Thus, there is convergence toward a general hypothesis about how goals are able to exert top-down control in tasks that require rapid responses to stimuli. The hypothesis is that top-down control serves to prepare critical mental mechanisms to carry out certain mental operations, but has little or no online involvement during the actual execution of those mental operations. In other words, top-down effects are “out of the loop”. This hypothesis is very appealing given that top-down (i.e. conscious) control over mental processing appears to be relatively slow. Shifts of covert spatial attention, for instance, proceeds much more slowly when directed by a central cue than when automatically captured by a peripheral cue (Posner, 1980). Due to the inherent sluggishness of top-down control, activities that require rapid and accurate responding (e.g. speeded RT tasks, hitting a baseball) might best be performed “automatically;” that is, using autonomous mechanisms that, once pre-programmed, require no further top-down control. This hypothesis is consistent with the lore from athletics that one should not attempt to think online about the mechanics of a rapid action (e.g. swinging a golf club or a baseball bat). Note, however, that this hypothesis might not apply very early in practice, when the critical mental operations still require oversight from top-down mechanisms.

Summary

Once task lag (the number of trials since a task was last performed) was properly accounted for in our data analysis, task expectancy and task recency were found to have additive effects on response time. Additivity was observed in several experiments involving different stimulus-response mappings, numbers

of tasks, and levels of expectancy. The robust additivity argues against many different candidate models of task preparation that, a priori, seemed very plausible. For example, additivity argues against the Activation Suppression model, in which participants deliberately suppress the activations of recently performed tasks when those tasks are not expected to occur on the next trial. The data also argue against any model in which task expectancy and task recency are assumed to have the same basic effect (e.g. they both influence some generic state of task readiness). The present data, however, support the Separate-Stages model, which asserts that expectancy and recency have qualitatively different effects on different stages of mental processing. One specific proposal, consistent with most details of the present and previous task-switching results, is that task expectancy affects the time required to prepare upcoming central mental operations whereas task recency affects the time required to execute those central operations once they have been programmed (see also, Fagot 1994). This “Configuration-Execution” model embodies a more general principle that top-down control is often needed to preprogram rapid mental operations, but has little online involvement during the execution of those mental operations.

References

- Allport, A., Styles, E. A., & Hsieh, S. (1994). Shifting intentional set: Exploring the dynamic control of tasks. C. Umiltà & M. Moscovitch (Eds.), *Attention and performance XV: Conscious and nonconscious information processing*. Cambridge: MIT Press.
- Balota, D. A. (1983). Automatic semantic activation and episodic memory encoding. *Journal of Verbal Learning and Verbal Behavior*, **22**, 88-104.
- Balota, D. A., Black, S. R., & Cheney, M. (1992). Automatic and attentional priming in young and older adults: Reevaluation of the two-process model. *Journal of Experimental Psychology: Human Perception and Performance*, **18**, 485-502.
- Bertelson, P. (1963). S-R relationships and reaction time to new versus repeated signals in a serial task. *Journal of Experimental Psychology*, **65**, 478-484.
- Campbell, K. C., & Proctor, R. W. (1993). Repetition effects with categorizable stimulus and response sets. *Journal of Experimental Psychology: Learning, Memory and Cognition*, **19**, 1345-1362.
- Epstein, W. & Rock, I. (1960). Perceptual set as an artifact of recency. *American Journal of Psychology*, **73**, 314-328.
- Fagot, C. (1994). Chronometric investigations of task switching. Unpublished dissertation, University of California, San Diego. (Available on the world-wide web at http://pima.ucsd.edu/pashler/default/pashler_manuscript.htm)
- Folk, C. L., Remington, R.W., & Johnston, J.C. (1993). Involuntary covert orienting is contingent on attentional control settings. *Journal of Experimental Psychology: Human Perception and Performance*, **18**, 1030-1044.
- Gottsdanker, R. (1980). The ubiquitous role of preparation. In G. E. Stelmach and J. Requin (Eds.), *Tutorials in Motor Behavior*, **3**, 301-312.
- Hansen, R. S. & Well, A. D. (1984). The effects of stimulus sequence and probability on perceptual processing. *Perception & Psychophysics*, **35**, 137-145.
- Jersild, A. T. (1927). Mental set and shift. *Archives of Psychology*, Whole No. 89.
- Kleinsorge, T. (1999). Response repetition benefits and costs. *Acta Psychologica*, **103**, 295-310.
- Kleinsorge, T., & Heuer, H. (1999). Hierarchical switching in a multi-dimensional task space. *Psychological Research*, **62**, 300-312.
- Kornblum, S. (1969). Sequential determinants of information processing in serial and discrete choice reaction time. *Psychological Review*, **76**, 113-131.
- Los, S. A. (1999). Identifying stimuli of different perceptual categories in pure and mixed blocks of trials: Evidence for stimulus-driven switch costs. *Acta Psychologica*, **103**, 173-205.
- Marcel, T. and Forrin, B. (1974). Naming latency and the repetition of stimulus categories. *Journal of Experimental Psychology*, **103**, 450-460.
- McCann, R. S., Remington, R., & Folk, C. L. (2000). Task-switching bottlenecks and PRP bottlenecks are not the same. Presented at the annual meeting of the Psychonomic Society, New Orleans.
- McClelland, J. L. (1979). On the time relations of mental processes: A framework for analyzing processes in cascade. *Psychological Review*, **86**, 287-330.
- Miller, J. O., van der Ham, F., & Sanders, A. F. (1995). Overlapping stage models and reaction time additivity: Effects of the activation equation. *Acta Psychologica*, **90**, 11-28.
- Neely, J. H. (1977). Semantic priming and retrieval from lexical memory: Roles of inhibitionless spreading activation and limited-capacity attention. *Journal of Experimental Psychology: General*, **106**, 226-254.
- Pashler, H. & Baylis, G. (1991a). Procedural learning: 1. Locus of practice effects in

- speeded choice tasks. *Journal of Experimental Psychology: Learning Memory and Cognition*, **17**, 20-32.
- Pashler, H. & Baylis, G. (1991b). Procedural learning: 2: Intertrial repetition effects in speeded-choice tasks. *Journal of Experimental Psychology: Learning Memory and Cognition*, **17**, 33-48.
- Pashler, H. & Johnston, J.C. (1989). Chronometric evidence for central postponement in temporally overlapping tasks. *Quarterly Journal of Experimental Psychology*, **41A**, 19-45.
- Posner, M. I. (1980). Orienting of attention. *Quarterly Journal of Experimental Psychology*, **32**, 3-25.
- Roberts, S. & Sternberg, S. (1993). The meaning of additive reaction-time effects: Tests of three alternatives. *Attention and Performance XIV: Synergies in Experiment Psychology, Artificial Intelligence, and Cognitive Neuroscience – A Silver Jubilee*. Cambridge, MA: MIT Press.
- Rogers, R. D., & Monsell, S. (1995). Costs of a predictable switch between simple cognitive tasks. *Journal of Experimental Psychology: General*, **124**, 207-231.
- Ruthruff, E., Johnston, J. C., & Van Selst, M. V. (2001). Why practice reduces dual-task interference. *Journal of Experimental Psychology: Human Perception and Performance*, **27**, 3-21.
- Sohn, M., & Carlson, R. A. (2000). Effects of repetition and foreknowledge in task-set reconfiguration. *Journal of Experimental Psychology: Learning Memory and Cognition*, **26**, 1445-1460.
- Spector, A. & Biederman, I. (1976). Mental set and mental shift revisited. *American Journal of Psychology*, **89**, 669-679.
- Sternberg, S. (1969). The discovery of processing stages: Extensions of Donders' method. *Acta Psychologica*, **30**, 276-315.
- Sudevan, P. & Taylor, D. A. (1987). The cuing and priming of cognitive operations. *Journal of Experimental Psychology: Human Perception and Performance*, **13**, 89-103.
- Van Selst, M. A., Ruthruff, E., & Johnston, J. C. (1999). Can practice eliminate the Psychological Refractory Period effect? *Journal of Experimental Psychology: Human Perception and Performance*, **25**, 1268-1283.

Acknowledgements

This research was supported by grants from the National Research Council and the National Aeronautics and Space Administration. We thank Julie Creagh, Amelia Hunt, Pierre Jolicoeur, Ray Klein, Valerie Lewis, Rob McCann, Nachshon Meiran, and Beth Siegel for useful comments and discussions. Please send correspondence to Eric Ruthruff, NASA – Ames Research Center, Moffett Field, CA 94035, or eruthruff@mail.arc.nasa.gov.

Table 1: Best-fitting parameter values for a model that assumes the effects of expectancy are additive with the effects of task lag, and that response time increases linearly with the log of the task lag. This model has two free parameters: (a) the intercept difference [in msec] between the expected and unexpected conditions (i.e. the size of the task expectancy effect), and (b) the slope [in msec/log(lag)] of the functions relating RT to lag (i.e. the size of the task recency effect).

	Intercept Difference (the task expectancy effect)	Slope (the task recency effect)
Experiment 1	35.7	85.2
Experiment 2	35.6	108.7
Experiment 3	59.4	94.7
Experiment 4	30.6	58.7
Average	40.3	86.8

Figure Captions

Figure 1: Possible results for the interaction between task expectancy and task repetition on response time. (A) Underadditive, (B) Overadditive, (C) Additive.

Figure 2: The Configuration-Execution model. Response selection does not begin until (a) the required input processing has been completed and (a) the task-set has been configured. Task expectancy influences the duration of the task-set reconfiguration stage. Task recency influences the duration of response selection.

Figure 3: Stimulus grids used in Experiments 1-4. The first stimulus of each block appeared in the upper-left cell; each subsequent stimulus appeared in the cell located immediately clockwise from the location of the previous stimulus. Numbers in parentheses (not present on-screen during the experiment) indicate the percentage of trials in which the expected task was presented in those cells. (A) Experiments 1-2, (B) Experiment 3, (C) Experiment 4.

Figure 4: Results of Experiment 1 as a function of task expectancy and task repetition.

Figure 5: Results of Experiment 2 as a function of task expectancy and task repetition.

Figure 6: Results of Experiment 3 as a function of task expectancy and task repetition.

Figure 7: Results of Experiment 4 as a function of task expectancy and task repetition.

Figure 8: Response time as a function of the task lag, plotted on a log scale.

Figure 9: Results of Experiments 1-4 as a function of task expectancy and task lag (plotted on a log scale). Lag 1 corresponds to the task repetition condition; lags greater than 1 correspond to task switch conditions. The lines represent the best fitting model in which task expectancy and task lag have additive effects. (A) Experiment 1, (B) Experiment 2, (C) Experiment 3, (D) Experiment 4.

Figure 1A

(A) Underadditivity

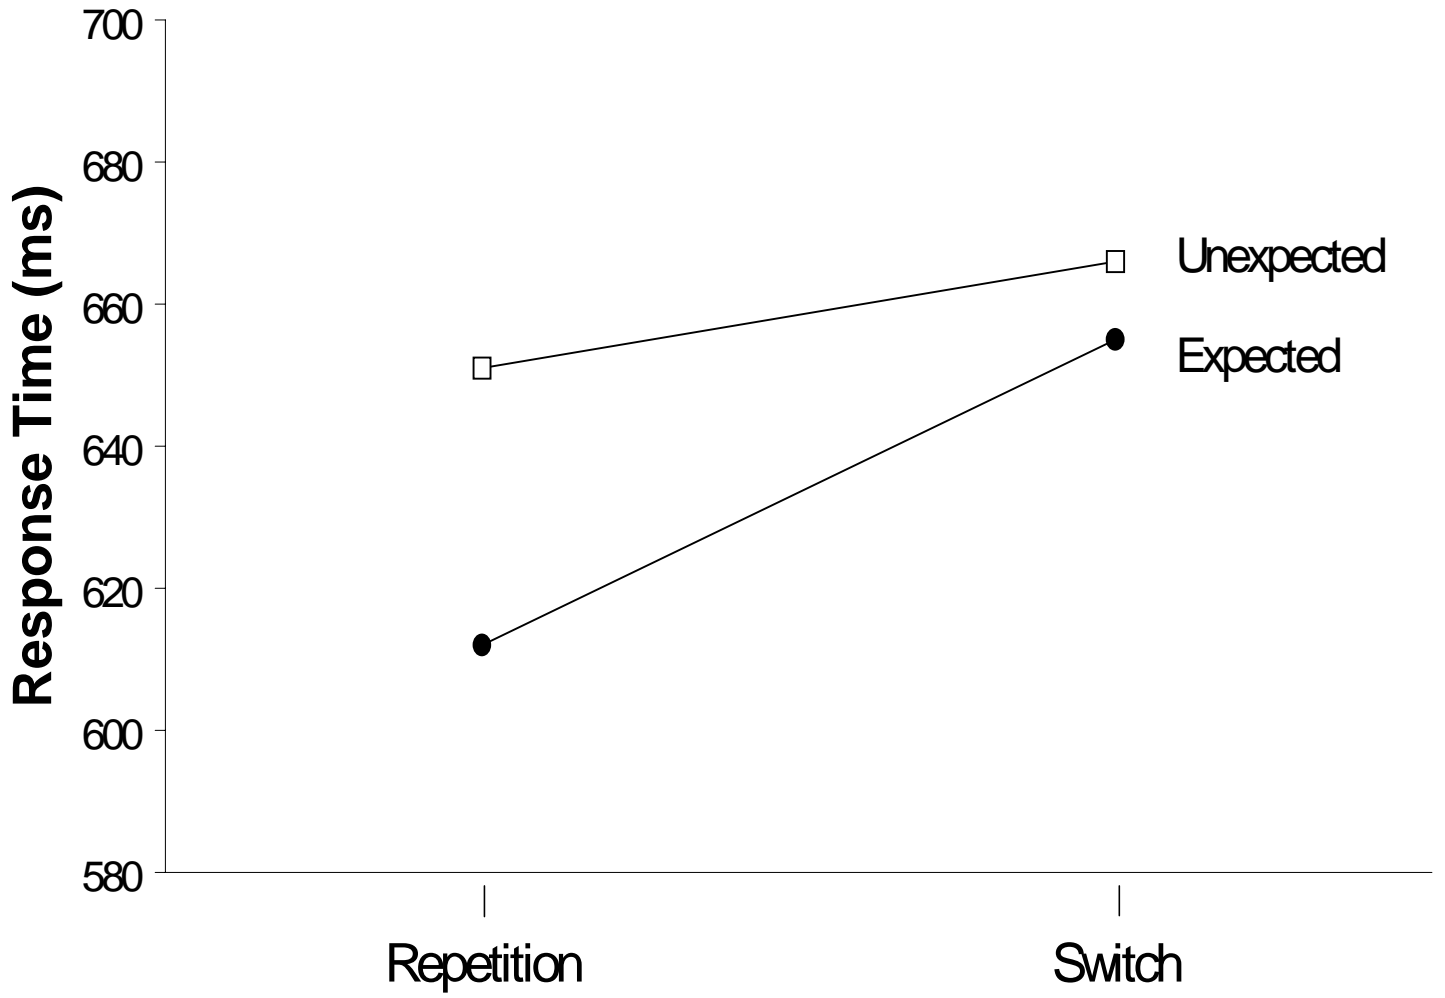
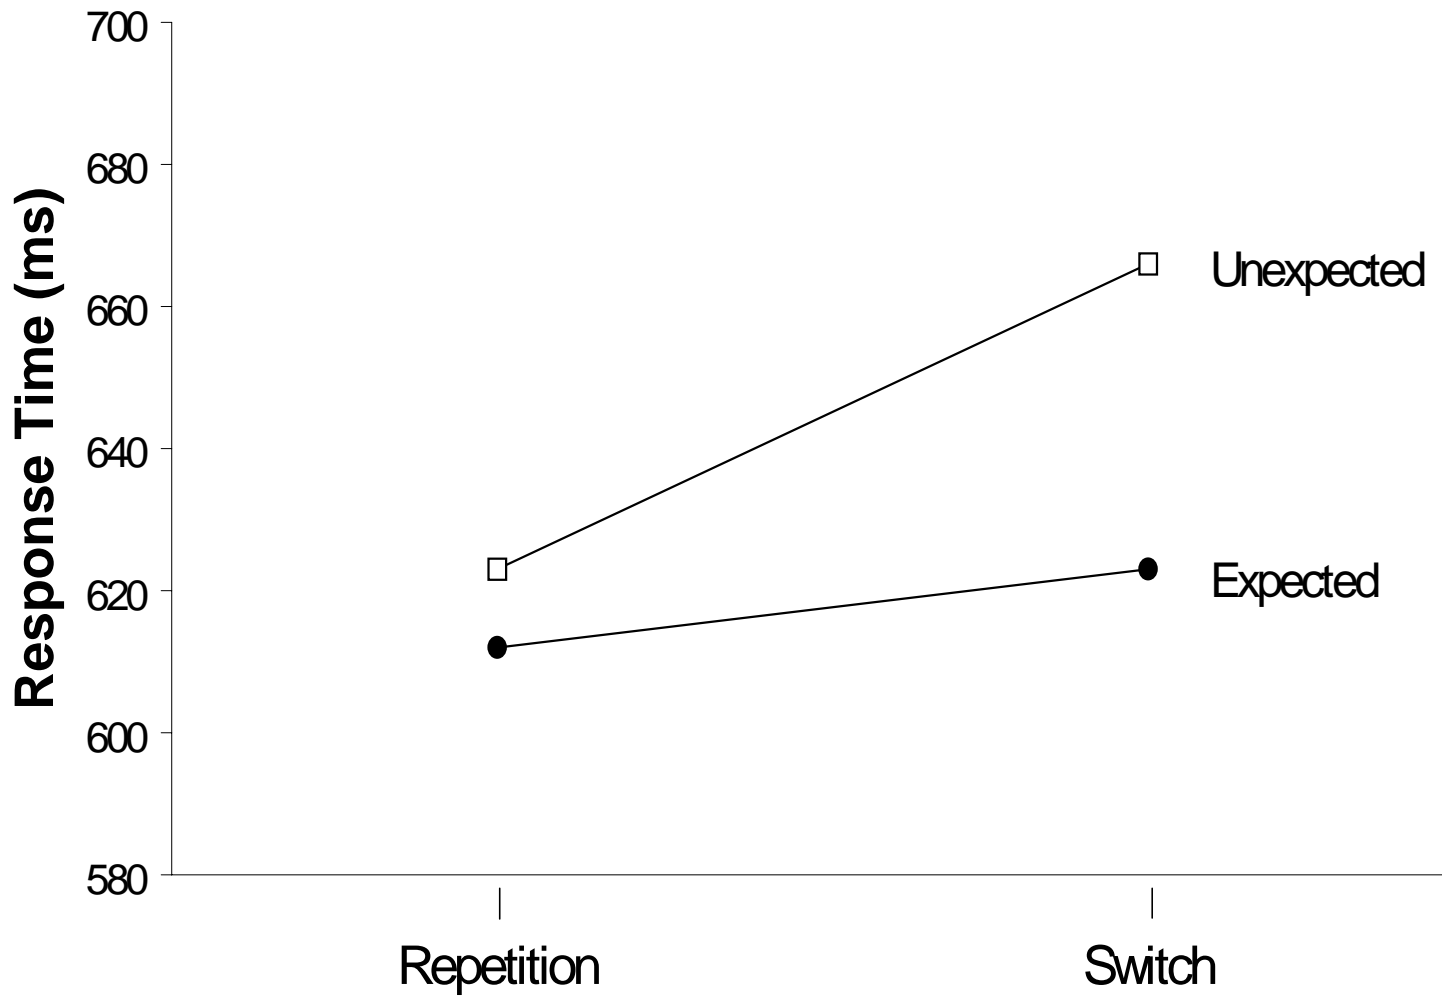


Figure 1B

(B) Overadditivity



(C) Additivity

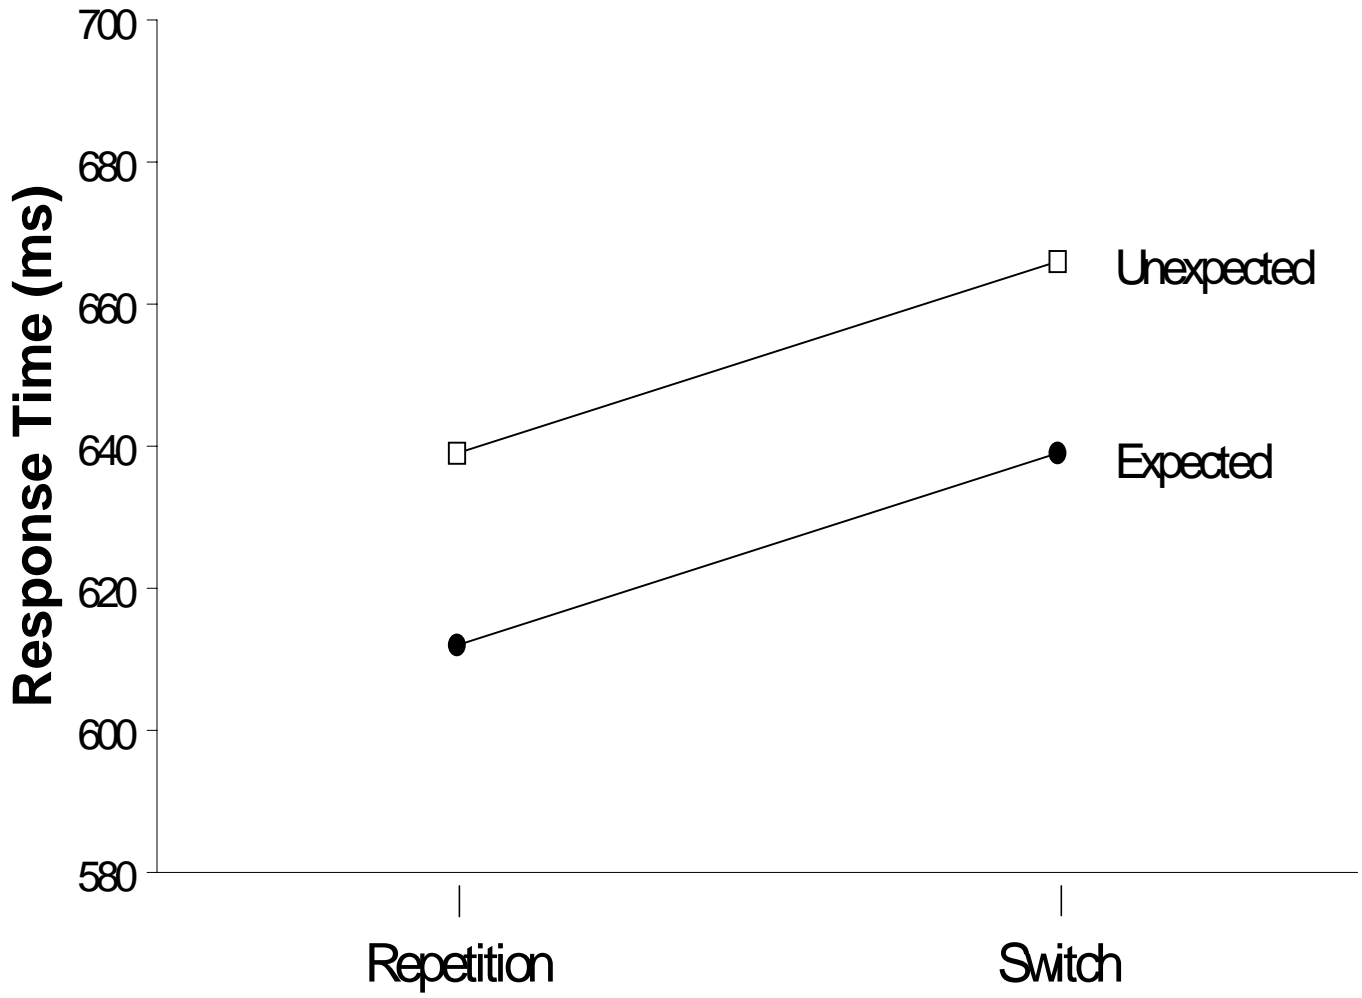
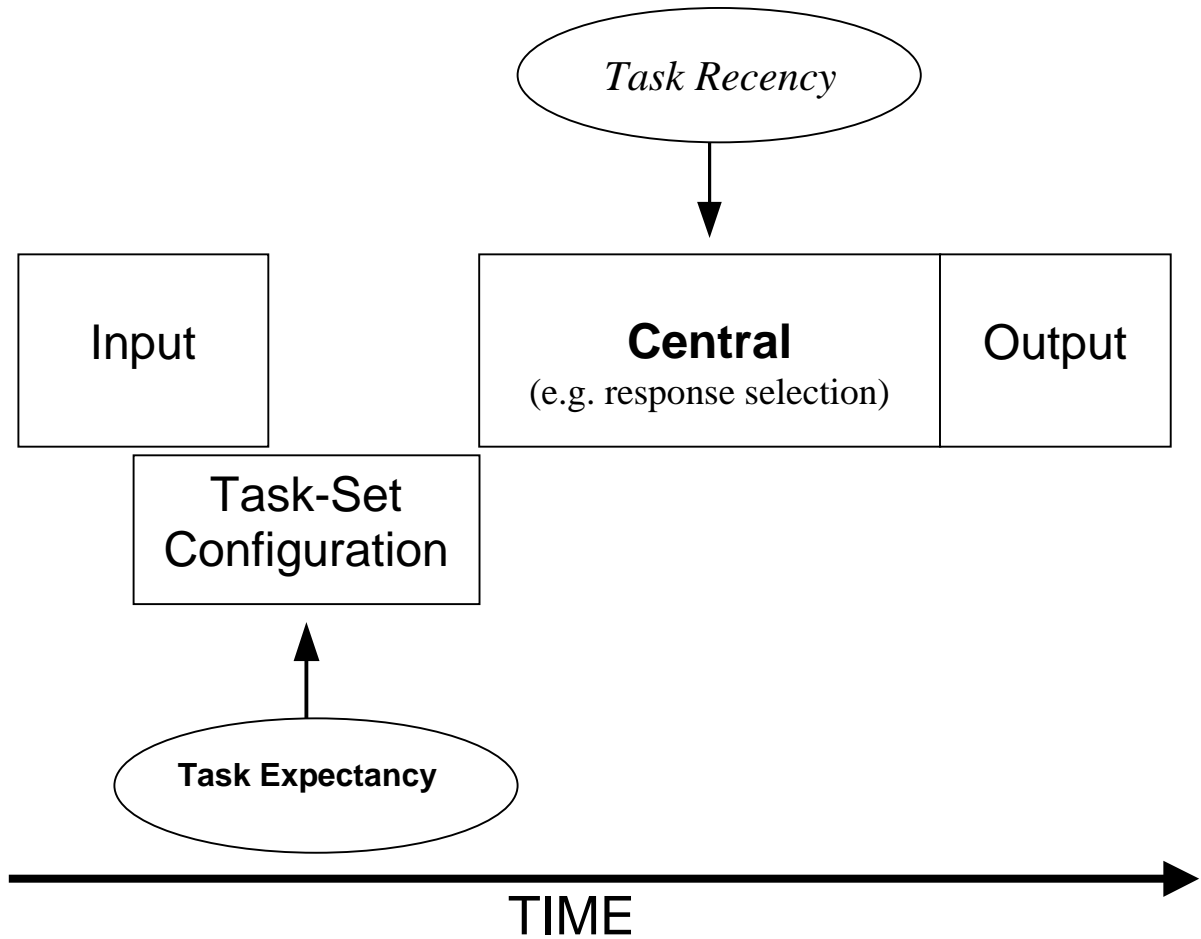


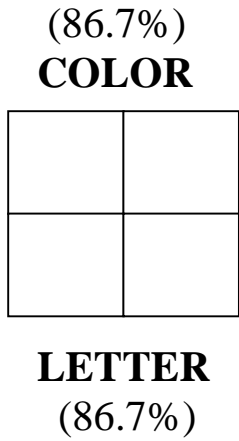
Figure 2

Configuration-Execution Model

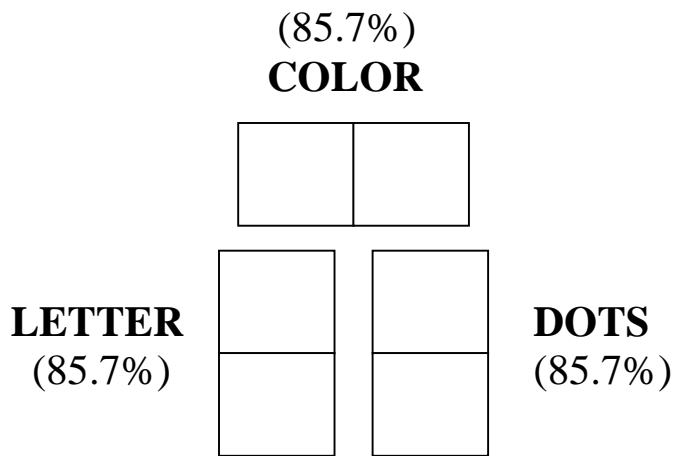


(A) Experiments 1-2

Figure 3



(B) Experiment 3



(C) Experiment 4

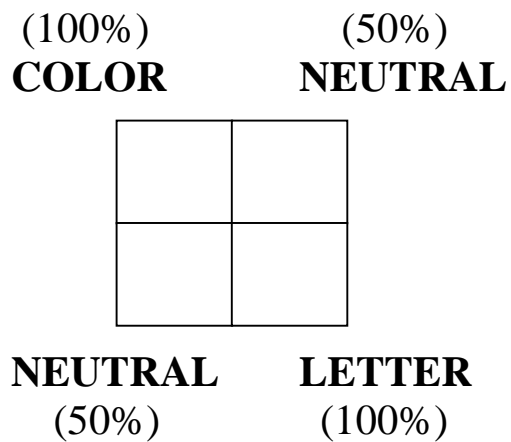


Figure 4

Experiment 1

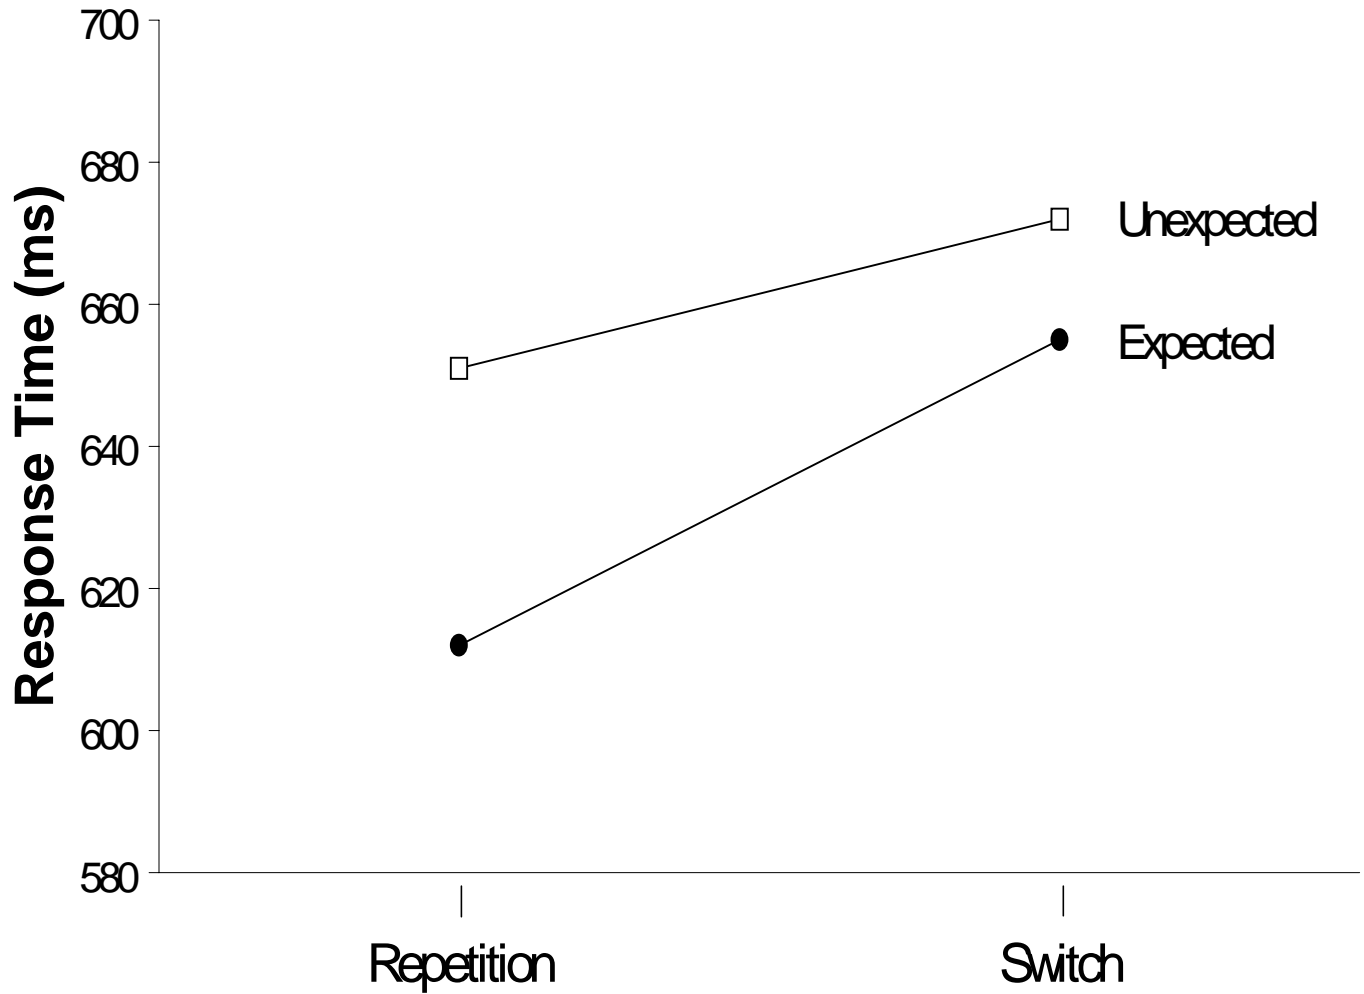


Figure 5

Experiment 2

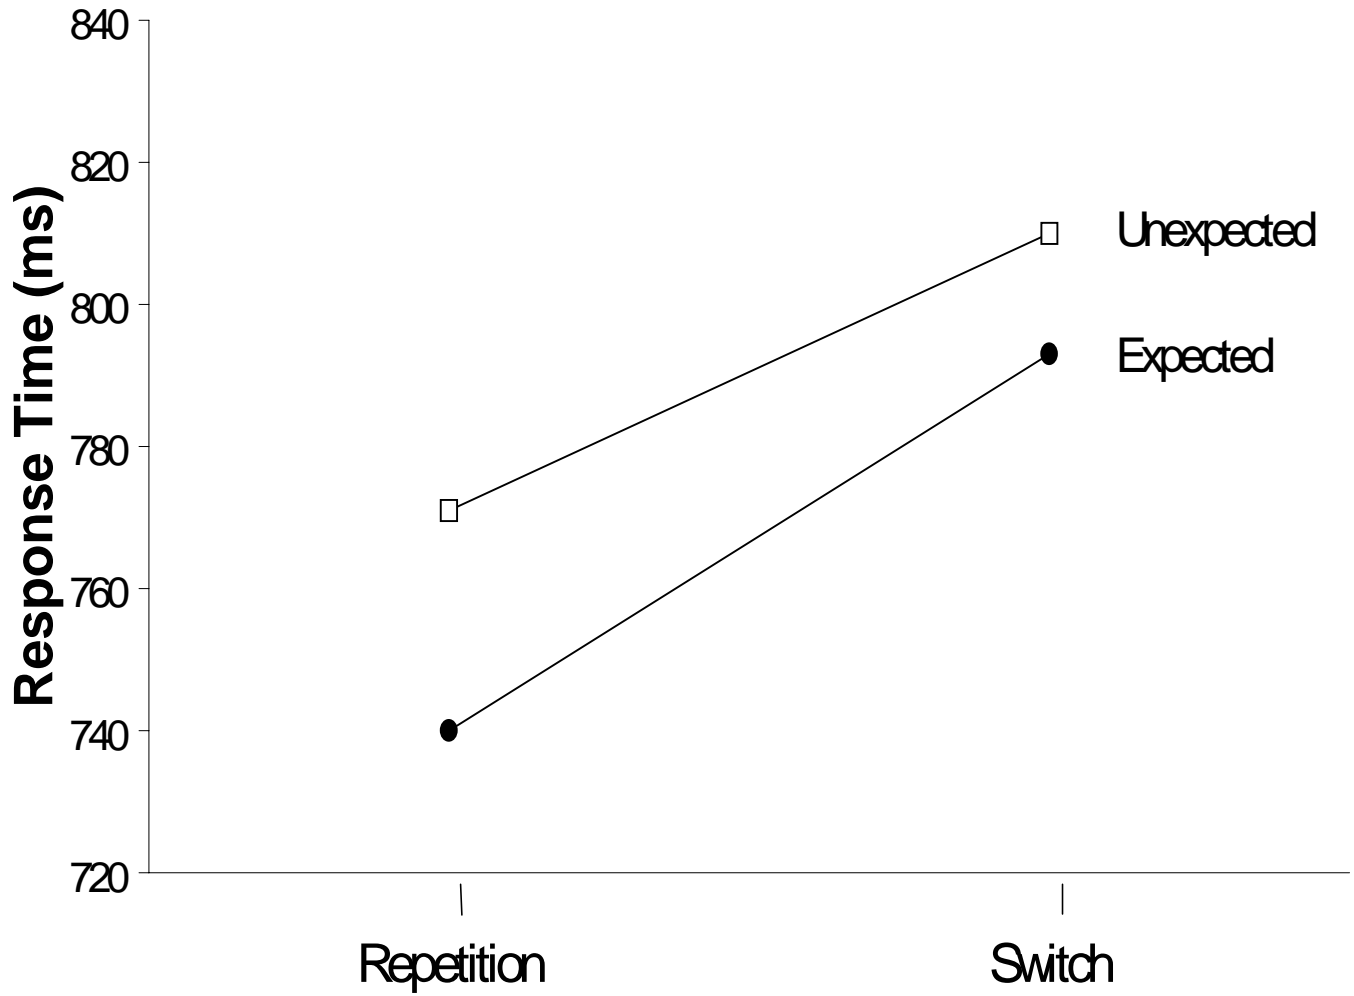
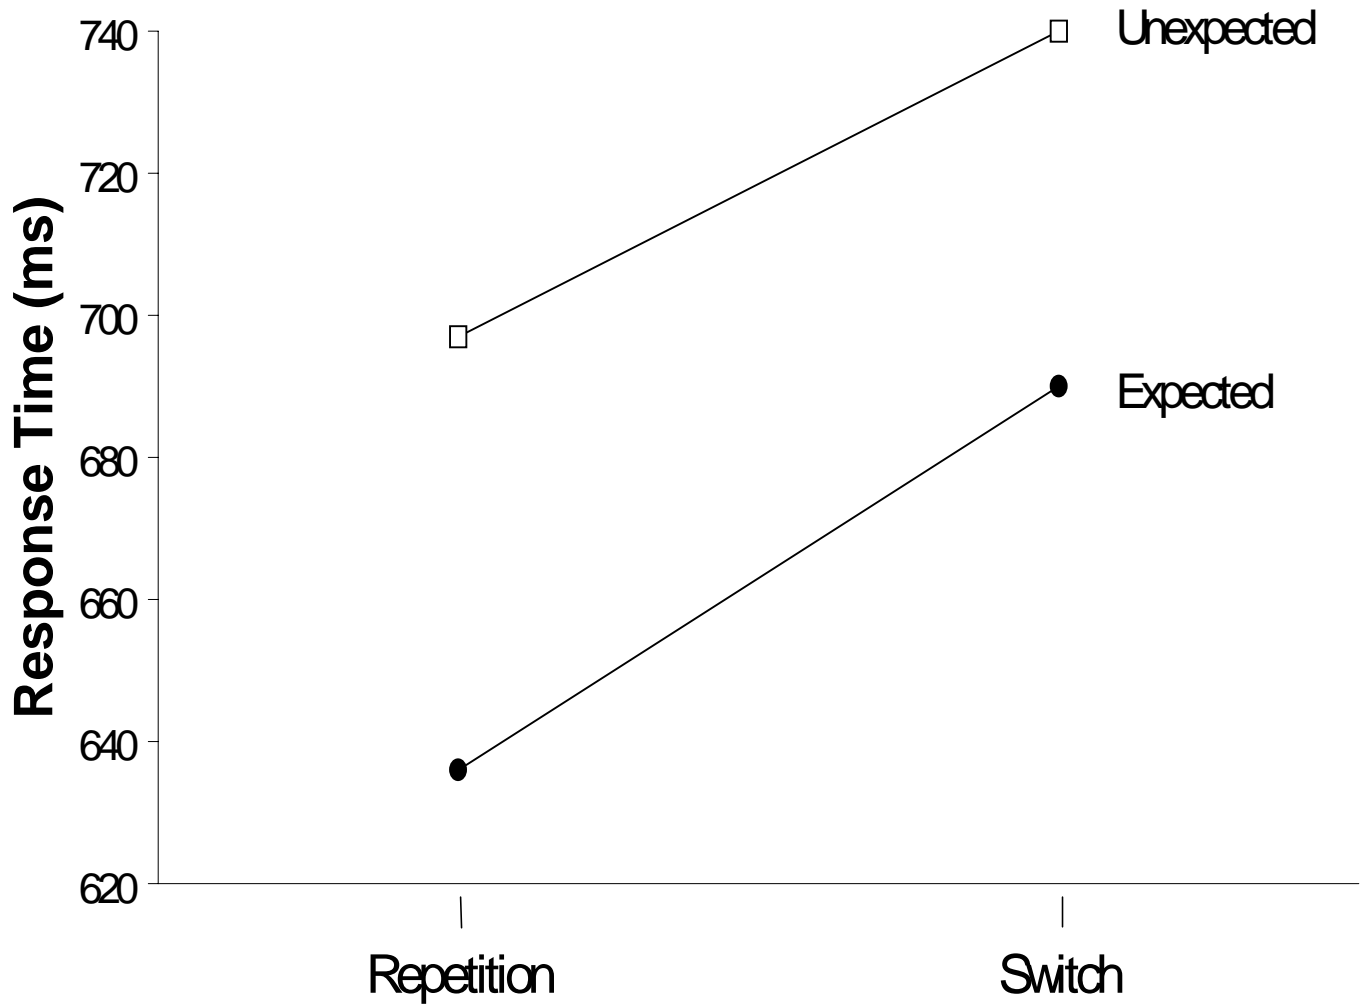


Figure 6

Experiment 3



Experiment 4

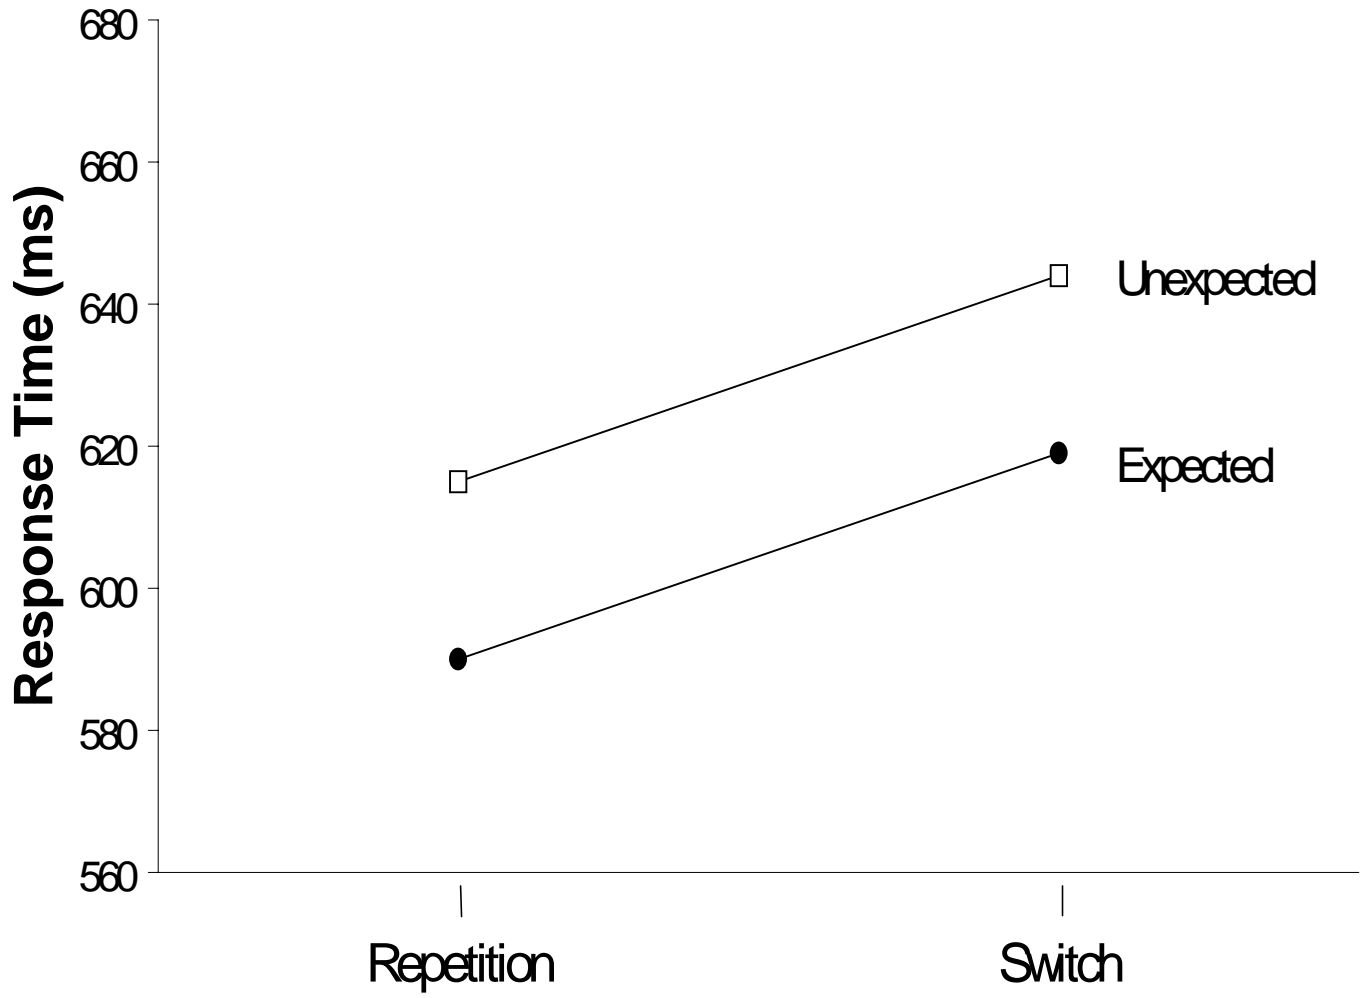


Figure 7

Figure 8

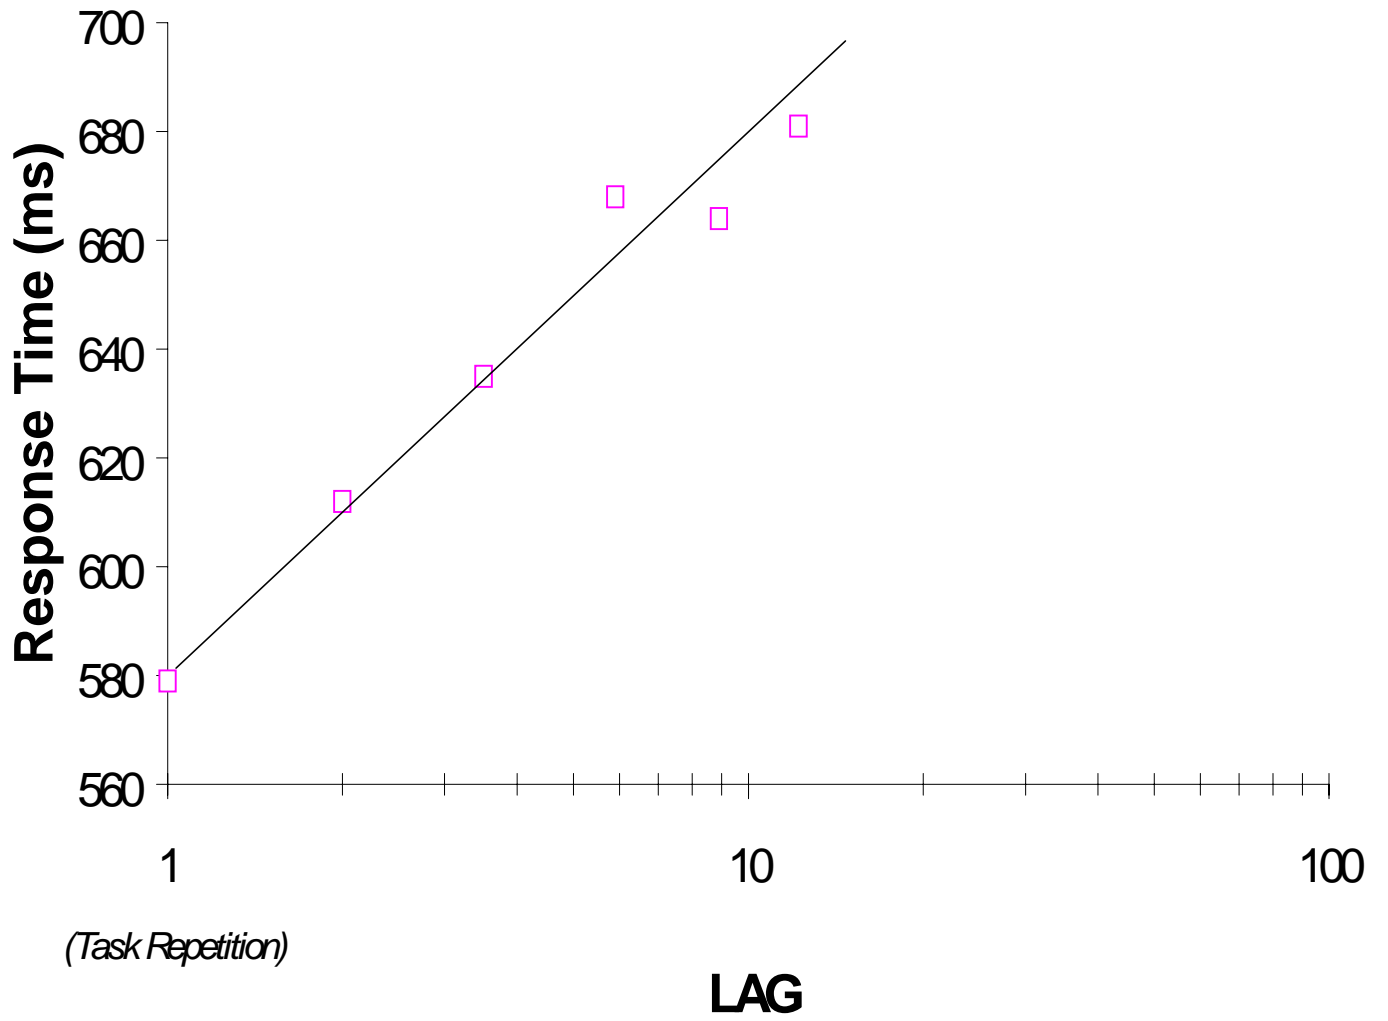


Figure 9A

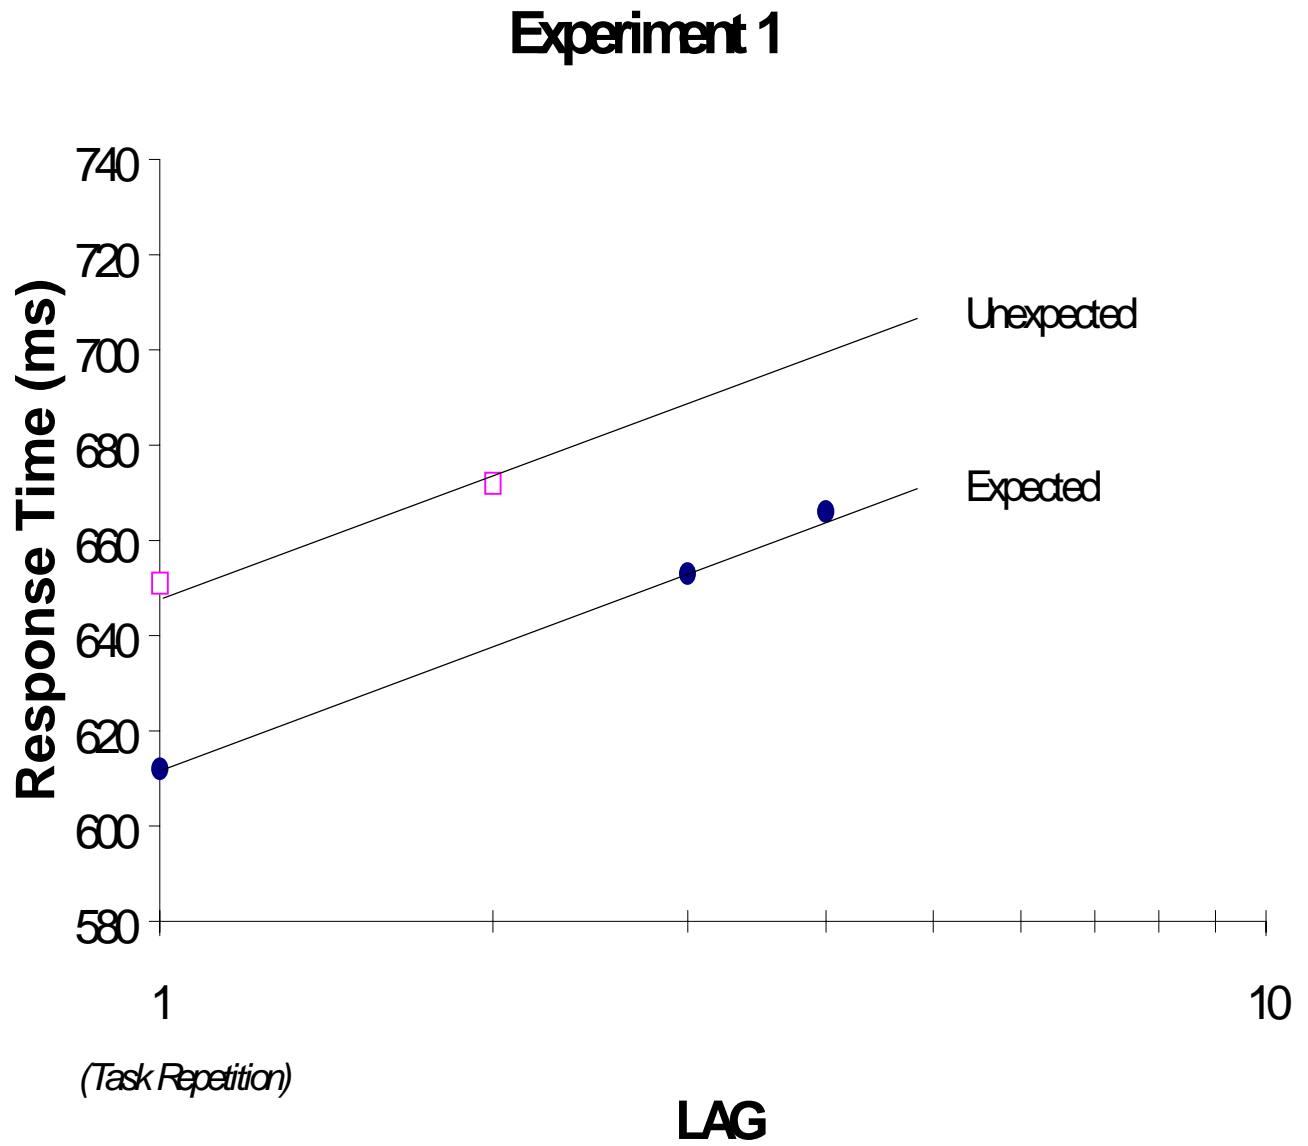


Figure 9B

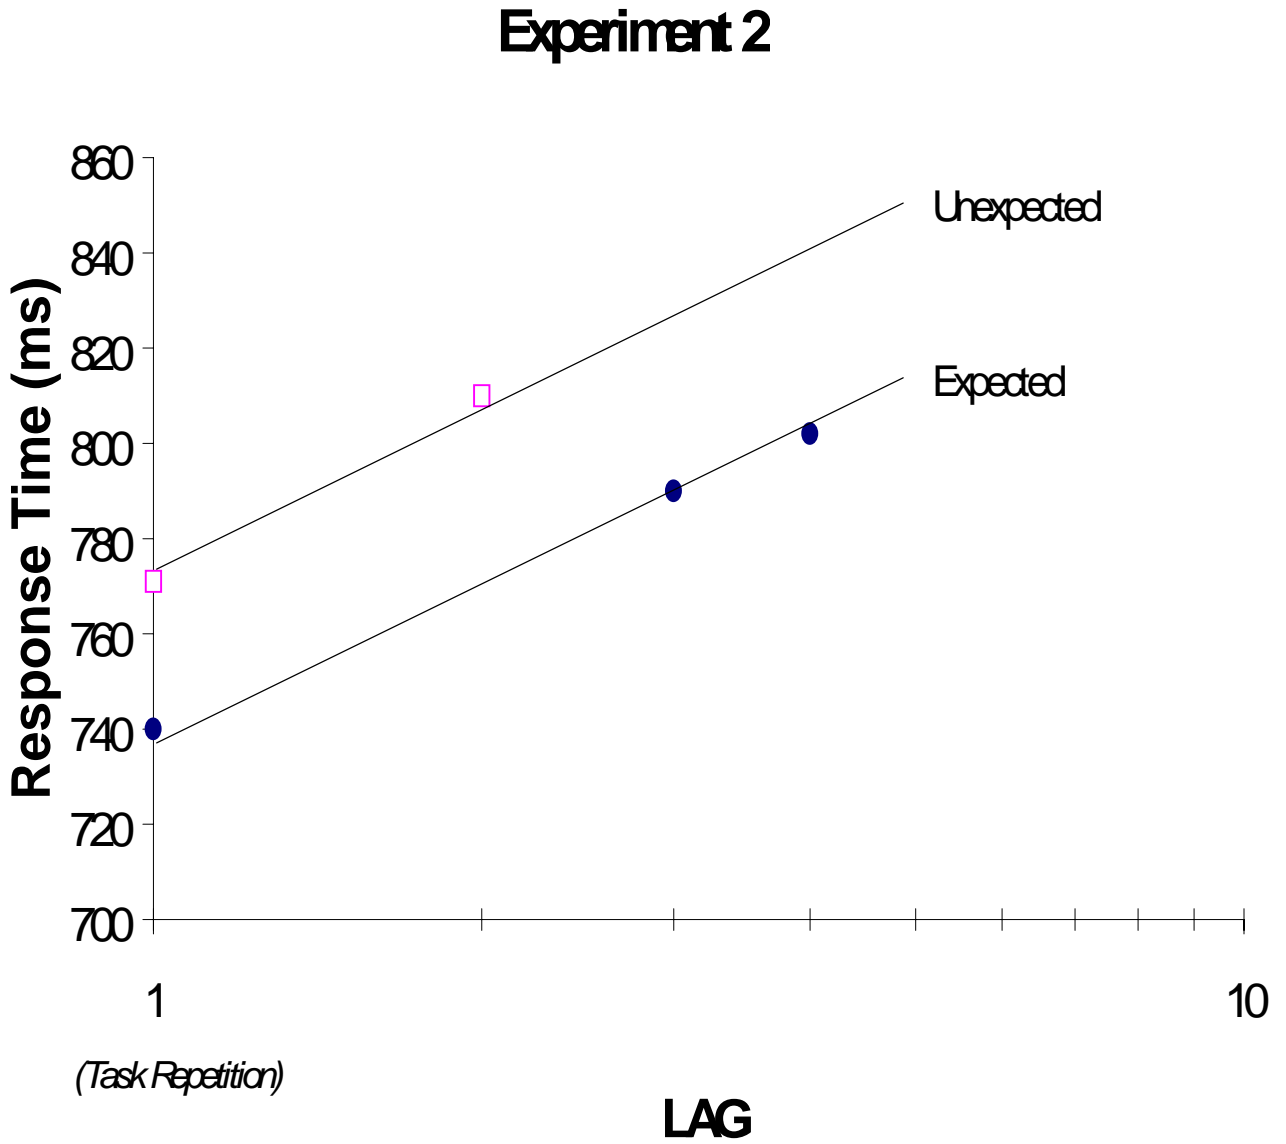


Figure 9C

Experiment 3

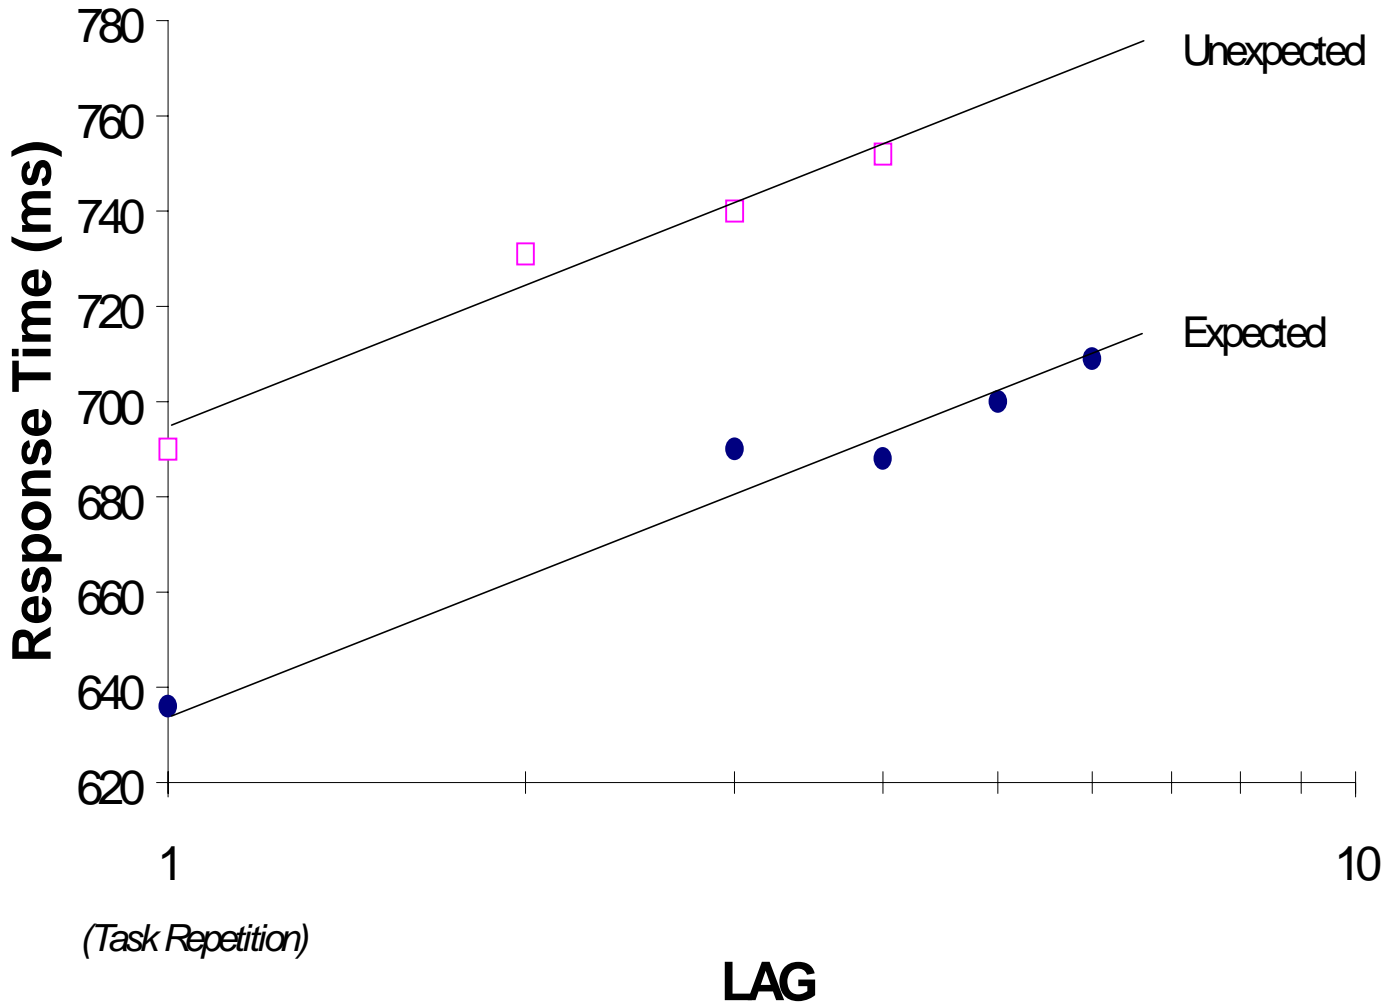
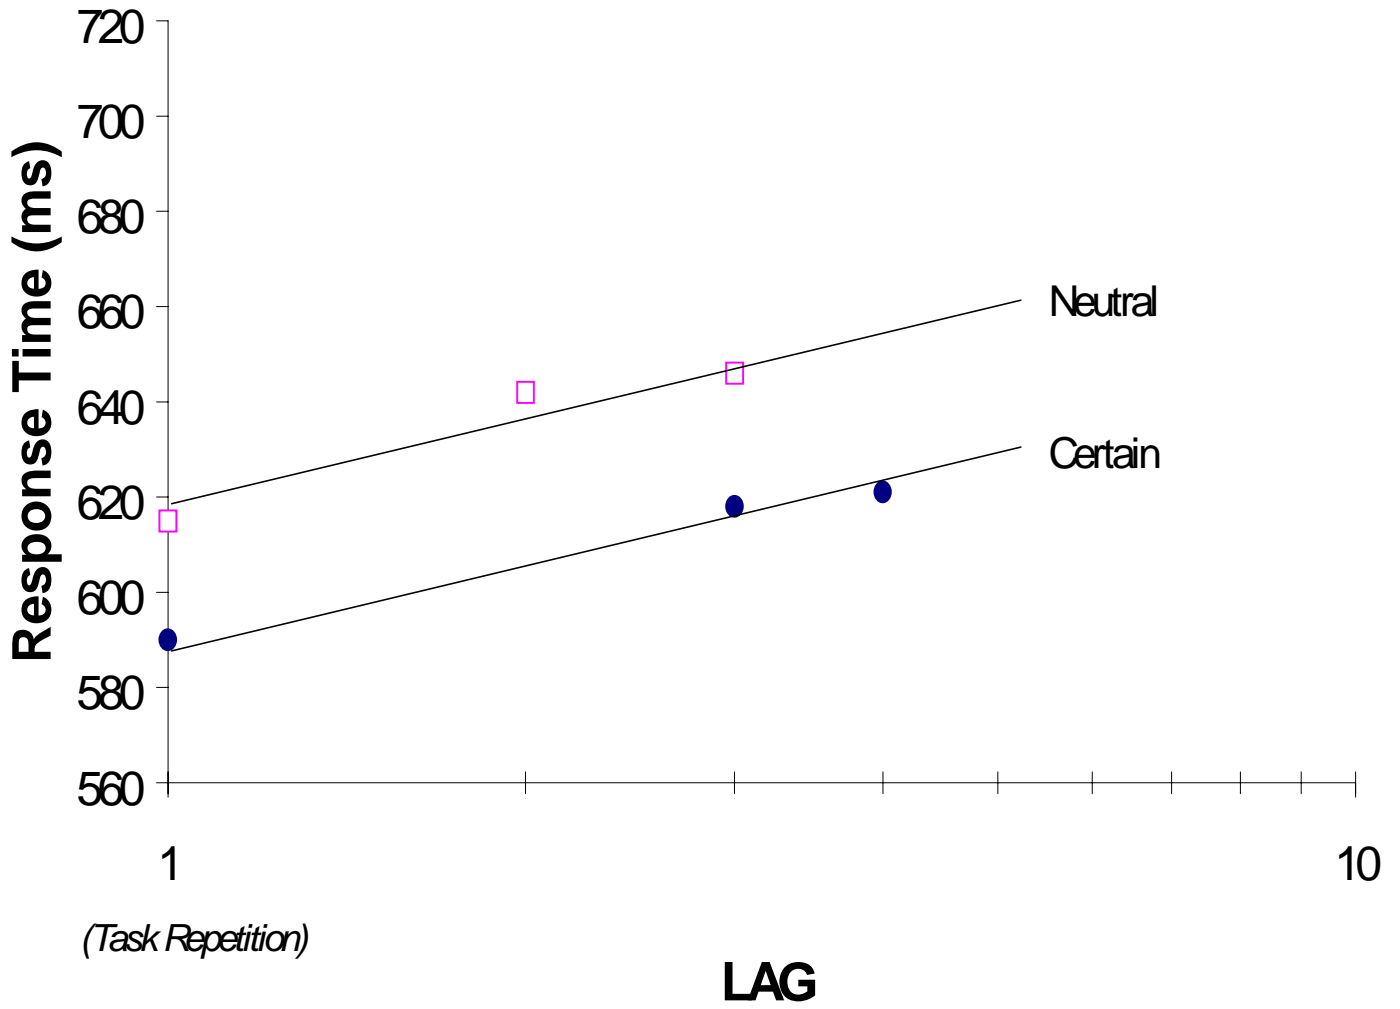


Figure 9D

Experiment 4



Footnotes

¹ Marcel and Forrin (1974) touched on the related question of how category repetition interacts with category expectancy. Using a character-naming task with letters and digits, they found faster responses to category repetitions than category switches. In a follow-up experiment, they found that this category repetition effect was large for unexpected categories but nearly absent for expected categories (an overadditive interaction). Because they studied character naming, it is unlikely that the different categories (letter vs. digit) constituted different tasks. Thus, their study is only weakly related to the present goals. Marcel and Forrin proposed that their effects are due to stronger inter-item associations within than between categories. Because they used a naming task, inter-item response associations occurred between responses (i.e. the response “one” is more highly associated with “two” than with “a”) and we believe that is the most likely locus of their effects. Note that we used the same response codes for both tasks; thus response association strength could not have played a role. This analysis might help explain why, as will be seen, our results differed so strikingly from theirs.

² Very late in the process of publishing the present work, we became aware of a paper by Sohn and Carlson (2000) on the interaction between task repetition and task foreknowledge. On the critical trials of their study, the foreknowledge group could infer what the upcoming task would be, but the no-foreknowledge group could not. Foreknowledge effects on RT were found to be roughly additive with the effects of task repetition. Although Sohn and Carlson’s foreknowledge manipulation resembles our task expectancy manipulation, there are several important differences. Whereas our stimulus displays were single-affordance, theirs were dual-affordance; each display contained both a letter and a digit, with color used to indicate which was relevant on that trial. More importantly, their foreknowledge manipulation combined expectancy effects with other variables. Specifically, the no-foreknowledge condition forced participants to determine stimulus color and use it to decide which task should be performed, but the foreknowledge condition did not. It is possible that the majority of their foreknowledge effect (222-252 ms) was due to this extra processing step, rather than to task expectancy per se. Consequently, we cannot determine in their paradigm how task expectancy, by itself, would have interacted with task repetition.

³ Additive factor logic appears to hold fairly well even when stages are not strictly serial. McClelland (1979) and Miller, van der Ham, and Sanders (1995) have shown that certain continuous-flow models (e.g. the Cascade model) – in which a sequence of contingent mental processes overlap somewhat in time – also predict approximately additive effects between factors that influence distinct mental processes.

⁴ Data were fit using a recursive parameter-search procedure that attempted to minimize the sum of squared errors.

⁵ It appears that the discrepancy in the original analyses was due simply to a weaker confounding of lag with expectancy in Experiment 4 than in Experiments 1-3.

⁶ In Experiment 4, the four stimuli of each task were mapped onto four separate response keys. Therefore, on task repetition trials, it was impossible to examine the effects of response repetition unconfounded by the effects of stimulus repetition. It was possible, however, to measure the response repetition effect on task switch trials: responses were 34 ms slower when the response repeated than when it changed.

⁷ As pointed out to us by a reviewer, it is possible that there is a temporary suppression of activation of the recently performed task when that task is not expected to repeat. Because the suppression is only temporary (e.g. during the deliberate task preparation that takes place during the inter-trial interval) in this case, it might not still be in place when the next trial begins. Thus, it is logically possible that some activation suppression occurred even though we observed a benefit of an unexpected task repetition on the subsequent trial.

⁸ We are grateful to Nachshon Meiran for pointing this out.