

Erik M. Altmann

Functional decay of memory for tasks

Accepted: 18 March 2002 / Published online: 15 August 2002
© Springer-Verlag 2002

Abstract Correct performance often depends on remembering the task one has been instructed to do. When the task periodically changes, memory for the current task must decay (lose activation) to prevent it from interfering with memory for the next task when that is encoded. Three task-switching experiments examine this decay process. Each shows *within-run slowing*, a performance decline occurring as memory for the current task decays. In experiment 1, slowing is attenuated when memory for the task is optional, suggesting that memory is indeed causal. Experiment 2 finds slowing despite a flat hazard rate for task instructions, suggesting that slowing is not an artifact of instruction anticipation. Experiment 3 finds slowing in the familiar alternating-runs paradigm (Rogers & Monsell, 1995), suggesting that it may lurk elsewhere. A process model of activation explains within-run slowing and relates it to switch cost and “restart cost” (Allport & Wylie, 2000) in functional terms.

Functional decay of memory for tasks

Correct performance often depends on remembering what one has recently been instructed to do. For example, one might be asked to fetch something from another room, in which case it is necessary, having achieved that room, to remember what the target thing was. In a more controlled setting, a participant in a psychological experiment might encounter an event sequence like $I_A A A A I_B B B B B$, where I_A and I_B are instructions indicating which task to perform, and A and

B are trials. Seeing I_A , the participant would perform task A for several trials, then switch to B on seeing I_B . The stimulus on each trial might be a simple character, “4,” for example, to be classified according to the current task. If the task were “odd or even,” then the correct response would be “even,” and if the task were “low or high” (relative to 5), then the correct response would be “low.” Memory for the current task (the task cued by the most recent instruction) is essential here because the stimulus is ambiguous (Spector & Biederman, 1976) or bivalent (Fagot, 1994) in affording either task.

In recent years, there has been growing interest in studying the processes that encode a new task cognitively such that it replaces the previous task and assumes control of performance. These task-encoding processes (and perhaps others) are reflected in what is commonly known as “switch cost,” or the latency increase on the first trial governed by a new task. The view has been that this switch cost “might seem to offer an index of the control processes involved in reconnecting and reconfiguring the various modules in our brains, so as to perform one task rather than another given the same input” (Monsell & Driver, 2000a). To date, however, the proliferation of empirical results on switch cost (Monsell & Driver, 2000b) has outpaced the development of integrative theory.

No task-switching work to date has focused on the question of how the current task is *retained* in memory. The importance of retention is shown by the event sequence above. If the instructions (I_A and I_B) are the only cues to the task – that is, if no additional cues appear on trials – then memory for the latest cue is the only source of information about what task to perform. Therefore, a memory trace for that cue must remain accessible until the next cue comes along.

This article tests a theory of retention of tasks in memory, and relates it to the processes of establishing a task set. Functional decay theory proposes that interference is a central constraint on cognitive control. However, contrary to the view that interference causes switch cost directly (Allport, Styles, & Hsieh, 1994;

E.M. Altmann
Department of Psychology,
Michigan State University, East Lansing,
MI 48824, USA
E-mail: ema@msu.edu
Tel.: +1-517-353-4406
Fax: +1-517-353-1652

Allport & Wylie, 2000), the argument is that interference drives cognitive adaptations, and that these adaptations are what we see reflected in behavior. One such adaptation is decay of memory for the current task, and the parameters of this decay are tested in three experiments. In the General discussion, an activation model is introduced to offer a process account of this decay. This activation model relates decay to initial task encoding, addressing some of the sources of behavioral variance often lumped together under the designation “switch cost.”

Functional decay theory

One way to frame the question of how a task is retained in memory is to try to predict the time course of performance over several trials under that task. For example, the sequence $I_A A A A A I_B B B B B$ begins with an instruction to do task A, which is followed by a run of four A trials. (In general a run could be of any length, and runlength need not be predictable). Does performance change in any systematic way over these trials? There appear to be roughly three possible answers. First, one could predict a practice effect (e.g., Newell & Rosenbloom, 1981) that would cause performance to improve within a run. Second, one could predict that the “modules in our brains” cited earlier, once “reconnected and reconfigured” to represent a new task, remain stable until control processes intervene again, producing no systematic change (Rogers & Monsell, 1995). The third and seemingly unlikely alternative is that performance *worsens* within a run.

The third prediction, that performance worsens within a run, is developed here, because, unlikely as it seems, it follows rather directly from an analysis of functional constraints on the cognitive system. The logic begins with a simple signal-detection analysis of memory for a task, as illustrated in Fig. 1. The abscissa in each panel shows the activation of a memory trace. A memory trace here represents an instructional cue, and for convenience will be referred to as a *task set*. The ordinate in each panel shows the probability of a task set being at a given activation level. Because memory is a noisy system, activation fluctuates transiently about an expected value, so is represented by a probability density function. Time progresses from the top panel to the bottom. In the top panel, a first task set has just been encoded, and in the bottom panel a second has just been encoded. Each has the same initial activation level, on the assumption that cue processing does not diverge in any systematic way over successive cues. In the middle panel, representing the interval between the two encodings, the first task-set decays (loses activation), such that its density function as a whole shifts gradually to the left along the activation axis. This decay meets a central functional constraint on the cognitive system: There must be a positive d' (in signal detection terms) separating the two activation densities in the bottom panel.




Fig. 1 A signal detection analysis of memory for task sets. A task set codes one event of being exposed to an instructional cue. Task set 1 is encoded in the top panel and decays in the middle panel, such that d' is positive when task set 2 is encoded in the bottom panel

Without a positive d' , the system would be unable to “detect” the later task set; in other words, there would be catastrophic interference among task sets, crippling the ability to be goal-directed.

The basic behavioral prediction of this signal detection analysis comes from the middle panel of Fig. 1. There, the task set encoded in the top panel decays gradually. Because this is the *current* task set – it represents the latest instruction – performance governed by that task set should reflect this decay. Specifically, response latency should increase gradually from trial to trial as the current task set loses activation and becomes harder for the system to recall. When the next cue appears, the system will encode a new, full-strength task set and performance will return to its pre-decay level. Thus, in terms of the three possible trends introduced earlier, performance should indeed worsen, rather than improve or stay the same. This worsening performance, as measured in terms of response latency in particular, will be referred to as *within-run slowing*.

Preliminary evidence for within-run slowing exists already (Altmann & Gray, 2002), and the goals here are

to strengthen the connection between data and theory, and to generalize the effect. Experiment 1 aims to demonstrate that memory is causal by manipulating the need for it – when memory is unnecessary, the effect should be attenuated. Experiment 2 addresses an alternative account of within-run slowing, namely instruction anticipation. The “aging” hazard function governing runlength in most task switching studies (including experiment 1) could let the system predict arrival of the next instruction and slow down in anticipation of having to shift mental gears. Experiment 3 aims to generalize the effect by searching for it in the familiar alternating-runs paradigm of Rogers and Monsell (1995). This experiment is also an opportunity to clarify how interference manifests under functional decay theory as compared to the *task-set inertia* model (Allport et al., 1994; Allport & Wylie, 2000). Two remaining theoretical issues are addressed in the General discussion. First, it is important to address the tarnished reputation of decay as a theory of forgetting. Second, a process model is introduced to account for decay and to relate it to run-initial effects like switch cost.

Experiment 1

Experiment 1 used sequences of the form $I_A A A A I_B B B B B$ in the context of a memory manipulation. In one condition, the stimulus sets for the two tasks were completely overlapping, and in the other they were completely non-overlapping. In the overlapping (*Necessary*) condition, memory for the most recent instructional cue was strictly necessary to disambiguate each of the stimuli that followed. In the non-overlapping (*Optional*) condition, the task could be inferred from the stimulus itself, so memory was not strictly necessary. The *memory contingency* prediction is that within-run slowing will be attenuated in the *Optional* condition. The prediction is not that within-run slowing will be eliminated, because short of disabling the memory system one cannot prevent its use. Indeed, memory could be a more efficient source of information about the task than an inference process run anew on every stimulus. Nonetheless, given that the system generally chooses stochastically from among available strategies (e.g., Hommel, 2000), making memory optional should attenuate within-run slowing in aggregate data.

Materials and methods

Experiment 1 was replicated twice, as experiments 1A and 1B, with some variety of tasks and stimuli. Forty-six Michigan State University undergraduates participated for \$10 or course credit, 22 per replicate and 11 per condition with two excluded for falling below 90% accuracy overall.

Stimuli were single digits or letters, 7 mm wide×10 mm high, presented in the center of a dark screen in a fixed-width sans serif font. A stimulus was red, blue, or white, with red and blue serving as instructional cues. Digits were 1 through 9, except 5, and letters were G, K, M, R, A, E, I, and U in experiment 1A and A, B, E, F,

T, U, X, and Y in experiment 1B. Stimulus presentation and response recording were controlled by software developed in Macintosh Common Lisp 4.3, running on Power Macintosh computers (240–500 MHz) under MacOS 8.6, 9, or 9.1.

Each replicate had two between-subjects conditions, *Necessary* and *Optional* (as described above). Four tasks were used in all: classifying the stimulus as odd or even (*OddEven*), lower or higher than five (*LowHigh*), consonant or vowel (*ConsVow*), or nearer the start or nearer the end of the alphabet (*StartEnd*). In experiment 1A the *Necessary* tasks were *LowHigh* and *OddEven* (both performed on digits) and the *Optional* tasks were *LowHigh* and *ConsVow* (one performed on digits and the other on letters). In experiment 1B, the *Necessary* tasks were *ConsVow* and *StartEnd* (both performed on letters), and the *Optional* tasks were *ConsVow* and *OddEven*. In the *Necessary* conditions, the instructional cue was a blue or red stimulus, which was followed by a run of white stimuli leading up to the next cue. In experiment 1A, red indicated *OddEven* and blue indicated *LowHigh*, and in experiment 1B, red indicated *StartEnd* and blue indicated *ConsVow*. In the *Optional* conditions, all stimuli were white, with no redundant color cues for the instruction. The assumption was that avoiding redundancy was more important than controlling for color cue because redundancy could open doors for strategic variation in cue processing.


The *runlength*, or number of trials following an instructional cue, ranged from one to eight, uniformly distributed. A block contained 12 runs. After the block, a feedback display appeared giving accuracy for that block, as well as the number of the next block to indicate progress through the session. There were 36 blocks per session, which took roughly 45 min to complete. Response time was measured from stimulus onset to key downstroke. Responses were made with the C and M keys of a standard keyboard, with participants asked to use their index fingers. In experiment 1A, *Odd*, *Low*, and *Consonant* mapped to C and *Even*, *High*, and *Vowel* mapped to M. In experiment 1B, *Consonant*, *Start*, and *Odd* mapped to C, and *Vowel*, *End*, and *Even* mapped to M. C and M also dismissed the feedback display. The response-stimulus interval (RSI) was zero, and there was no error feedback other than on feedback displays.

Participants were tested individually. Each was read a description of the tasks and responses and told the possible range of runlengths and the total number of blocks. Each was told, “It’s important that you try to work as accurately as you can from start to finish,” and was encouraged to use the feedback screen as a self-paced rest period. The participant completed a practice block of 50 trials broken into ten runs of varying lengths, with 25 trials under each task and five task switches. The color/task and response/key mappings were visible on the screen for the practice block and then removed. Participants were encouraged to memorize the mappings during the practice block. The experimenter was present during the practice block, and stayed longer if necessary until the participant was able to perform both tasks correctly and switch between them.

Results

Results are shown in Fig. 2. Latencies and standard deviations are means of participant medians and standard deviations, respectively, on correct trials. Standard deviations are examined to assess whether within-run slowing reflects a shift in location or an increase in variability of the latency distribution, and errors are examined to test whether within-run slowing simply reflects a speed-accuracy tradeoff. To control for task effects, only the task common to the two conditions was included. The first six blocks of a session were excluded to remove the period of most rapid learning as a source of variance, and the first run of a block was excluded as a warm-up after the feedback break. Data were analyzed

Fig. 2 Data from experiment 1. *Top panels:* within-run slowing is present in the Necessary condition and attenuated in the Optional condition. Markers with no lines are standard deviations. *Bottom panels:* error also increases within a run, in both conditions. Error bars are $\pm$ RMSE for the Position effect



with 7×2 repeated-measures analyses of variance (ANOVAs) on Position (P2 to P8) and Memory (Necessary, Optional). P0 was excluded because it was the stimulus/cue combination, so stimulus and cue processing were directly confounded. P1, the first non-cue trial of a run, is much slower than later trials (Gopher, Armony, & Greenspan, 2000; Kramer, Hahn, & Gopher, 1999), and was excluded to avoid obscuring the effect of interest.

Experiment 1A

For latencies, the main effect of Position was reliable [$F(6,120) = 8.0$, $P < 0.001$]. A contrast by the method of orthogonal polynomials showed a linear trend [$F(1,120) = 43.1$, $P < 0.001$], which accounted for 90% of the variance due to Position. No higher-order trends were significant. The effect of Memory was reliable [$F(1,20) = 5.7$, $P < 0.03$], and there was a Position $\times$ Memory interaction [$F(6,120) = 3.2$, $P < 0.007$].

To probe the interaction, separate one-way ANOVAs were conducted on Position for each level of Memory. In the Necessary condition, there was a main effect of Position [$F(6,60) = 8.4$, $P < 0.001$], and a linear trend [$F(1,60) = 44.8$, $P < 0.001$], which accounted for 89% of the variance due to Position. There was also a marginal cubic trend [$F(1,60) = 3.9$, $P = 0.052$], which accounted

for 8% of the variance due to Position, but no other trends were significant. In the Optional condition, there was no main effect of Position ($P > 0.1$), but there was a linear trend [$F(1,60) = 4.5$, $P < 0.04$], which accounted for 69% of the variance due to Position. No higher-order trends were significant.

For standard deviations, no main effects, trends, or interactions were significant.

For errors, there was a main effect of Position [$F(6,120) = 3.5$, $P < 0.004$], and a linear trend [$F(1,120) = 13.9$, $P < 0.001$], which accounted for 65% of the variance due to Position. No higher-order trends were significant. There was a main effect of Memory [$F(1,20) = 6.0$, $P < 0.03$], but no Position $\times$ Memory interaction ($F < 1$).

Experiment 1B

For latencies, there was a main effect of Position [$F(6,120) = 5.7$, $P < 0.001$], and a linear trend [$F(1,120) = 28.7$, $P < 0.001$], which accounted for 85% of the variance due to Position. No higher-order trends were significant. There was no effect of Memory ($P > 0.1$), but there was a Position $\times$ Memory interaction [$F(6,120) = 3.4$, $P < 0.005$].

To probe the interaction, one-way ANOVAs were conducted on Position for each level of Memory. In the Necessary condition, Position was reliable [$F(6,60) = 7.2$,

$P < 0.001$], as was the linear trend [$F(1,60) = 37.5$, $P < 0.001$], which accounted for 86% of the variance due to Position. There was again a marginal cubic trend [$F(1,60) = 3.6$, $P = 0.063$], accounting for 8% of the variance due to Position. In the Optional condition, Position was not significant ($F < 1$), nor were any trends.

For standard deviations, Memory was reliable [$F(1,20) = 5.1$, $P < 0.04$], but no other effects were.

For errors, there was a marginal main effect of Position [$F(6,120) = 2.1$, $P = 0.054$], and a marginal linear trend [$F(1,120) = 3.5$, $P = 0.065$], which accounted for 27% of the variance due to Position. No higher-order trends were significant. There was no main effect of Memory ($F < 1$), but there was a Position $\times$ Memory interaction [$F(6,120) = 14.5$, $P < 0.009$].

To probe the error interaction, one-way ANOVAs were conducted on Position for each level of Memory. In the Necessary condition, Position was reliable [$F(6,60) = 3.8$, $P < 0.004$], as were all trends except linear and quintic ($P < 0.05$), reflecting the influence of the outlier at P8. Statistical variation seems the most likely explanation for this outlier, particularly as errors are rare to begin with and grow more so for later positions (there were 134–205 observations per data point for P2 but only 12–35 for P8). With P8 omitted, the main effect of Position and the linear trend were significant ($P < 0.02$ and $P < 0.006$, respectively), with the linear trend accounting for 52% of the variance due to Position and no reliable higher-order trends. In the Optional condition, Position was not significant ($P > 0.1$), but there was a linear trend [$F(1,60) = 7.0$, $P < 0.03$], accounting for 75% of the variance due to Position. No other trends were significant.

Discussion

In both replicates, Necessary latencies increased gradually but steadily within a run, but Optional latencies increased barely at all, supporting the memory contingency prediction. In the logic of causality, these results are a critical demonstration of both *modus ponens* and *modus tollens* – the need to remember the latest cue implies within-run slowing, and attenuated within-run slowing implies no need for memory.

Errors increased in both conditions, supporting the theory in an equally critical way – had errors decreased within a run, one could have explained within-run slowing as half of a speed-accuracy tradeoff (e.g., Wickelgren, 1977). Finally, standard deviations showed no trends across trials, suggesting that within-run slowing does not reflect a growing influence of outliers, which might have implicated a gradual within-run adjustment of arousal or vigilance.

Experiment 2

Experiment 1 leaves open an alternative account of within-run slowing, which is that Necessary participants

were anticipating the appearance of the color cue in a way that caused a gradual trial-by-trial slowdown and increase in error. An anticipation account of within-run slowing seems intuitively plausible, and is the primary focus of experiment 2, but does bear several burdens that are important to identify at the outset. First, per experiment 1, anticipation would have to be linked to the format of the Necessary instructional cue, not to a pending switch per se (which also occurred in the Optional condition). Second, to match functional decay theory in terms of explanatory adequacy, anticipation would have to play a functional role, providing some kind benefit to the system in the context of task switching. Third, anticipation is a notoriously flexible construct. For example, if experiment 1 were repeated with a redundant color cue in the Optional condition, and there were still no within-run slowing in that condition, the objection might be that the system only anticipates cues that are not redundant. Thus, anticipation as a model has handicaps of its own.

Experiment 2 implemented three changes to make anticipation more difficult and less necessary. First, runlengths were now exponential, not uniform as in experiment 1. An exponential density produces a flat hazard rate for the imperative event, affording no predictive information about the event's time of occurrence (e.g., Luce, 1986). Here, a flat hazard rate for the cue meant that the probability of the cue occurring after the current trial was equal for all trials. Of course, participants could still have made subjective predictions about cue arrival, but to explain within-run slowing these would have to conflict with the statistical structure of the environment, placing yet another explanatory burden on an anticipation account.

The second change involved the format of the cue, which was now temporally and structurally distinct from the trial stimulus. In experiment 1, the integration of the color cue and stimulus in the Necessary condition meant that to encode a cue the system had to first inhibit its usual stimulus processing. Failure to inhibit could cause the system to “overrun” the cue, producing an error on that trial, and on downstream trials if it meant that the cue was not properly encoded in memory. To remove the danger of the cue being overrun, it was now inserted temporally between two trials and located spatially in the same position as trial stimuli. Thus, the cue was now much harder to miss.

The third change involved randomizing cue order. In experiment 1, tasks alternated predictably, but now each cue was equally likely to indicate either task. This is referred to below as the Continuity variable, with levels Switch and Noswitch. On Switch runs, the cue at the start of the run switched the task from the previous run (breaking task continuity), whereas on Noswitch runs the cue simply continued the task from the previous run. Thus if within-run slowing in experiment 1 reflected a gradual shift of mental resources away from the current task to prepare for the other task, this strategy would now pay off only at chance levels.

The Continuity variable also affords a prediction for errors, which under functional decay theory can be caused by interference. When the stimulus is incongruent (i.e., when the correct response differs with the task), recalling the wrong task will produce the wrong response. Thus incongruent stimuli should produce more response errors, but only on Switch runs, because only then does the previous task set (task set 1 at the bottom of Fig. 1) code the wrong task. On Noswitch runs, the previous task set codes the correct task, so a memory error should not cause a problem.

Materials and methods

Twelve Michigan State University undergraduates participated for course credit.

Materials, design, and procedure were as in experiment 1, with the following changes. Runlengths were now sampled from an exponential rather than a uniform distribution. Participants were informed that there would be at least one trial in every run, but that “there’s no fixed limit so you might get some sequences that seem to go on for a while.” Mean runlength was 3.67 trials, minimum runlength was one, and maximum runlength ranged from 16 to 31 across participants. There were now 48 blocks, of which the first eight were excluded from analysis. The mapping from response category to key (again C or M) was now randomized between participants.

The two tasks were OddEven and LowHigh and stimuli were again 1–9, except 5. However, each instructional cue was now equally likely to indicate either task. The cue itself now consisted of a distinct word pair like “Odd Even” displayed in the center of the screen, rather than a color change in the stimulus. The order of words in the pair was consistent with the response mapping for that participant (e.g., if the pair were “Odd Even,” odd would map to C and even to M). Each pair was visible for 500 ms immediately after the preceding trial. No response was necessary, and any responses were ignored. After instruction offset, P1 began immediately, and RSI after trials was again 0.

Results

Within-run data are shown in Fig. 3, and were analyzed with $8 \times 2 \times 2 \times 2$ ANOVAs on Position (P2 to P9), Continuity (Switch, Noswitch), Congruency (Congruent, Incongruent) and Task (OddEven, LowHigh). P1 was omitted as before. P9 was the last trial position for which there was at least one datum per participant. In terms of Congruency and Task effects (not shown in Fig. 3), Incongruent and OddEven trials were generally slower and less accurate than their opposites.

For latencies, Position had a main effect [$F(7,77)=7.8$, $P<0.001$], and linear [$F(1,77)=46.2$, $P<0.001$] and quadratic trends [$F(1,77)=4.1$, $P<0.05$] explaining 86% and 7% of the variance due to Position, respectively. Neither Continuity nor Task had main effects ($P>0.1$), but Incongruent trials were 21 ms slower than Congruent [$F(1,11)=8.9$, $P<0.02$]. There was a marginal Task $\times$ Continuity $\times$ Congruency interaction [$F(1,11)=4.1$, $P=0.064$], but no other reliable effects.

For standard deviations, there was a main effect of Position [$F(7,77)=2.2$, $P<0.05$], and a decreasing linear trend [$F(1,77)=9.6$, $P<0.003$]. No other effects were significant.




Fig. 3 Data from experiment 2. Within-run slowing is present despite steps taken to minimize anticipation. Markers with no lines are standard deviations. Errors are more frequent on Switch runs. Error bars are $\pm$ RMSE for the Position effect

For errors, Position had a marginal main effect [$F(7,77)=1.9$, $P=0.084$] and a fourth-order trend [$F(1,77)=6.1$, $P<0.02$]. Switch runs had 1.7% more errors than Noswitch [$F(1,11)=10.1$, $P<0.01$], Incongruent trials had 1.3% more errors than Congruent [$F(1,11)=10.4$, $P<0.01$], and there was a marginal Continuity $\times$ Congruency interaction [$F(1,11)=4.7$, $P=0.054$]. OddEven trials had 0.9% more errors than LowHigh [$F(1,11)=9.9$, $P<0.01$]. No other effects were significant.

To probe the Continuity $\times$ Congruency interaction, $8 \times 2 \times 2$ ANOVAs were applied to each level of Continuity. On Switch runs, Incongruent trials had 2.2% more errors than Congruent [$F(1,11)=9.5$, $P<0.02$], and OddEven trials had a marginal 0.9% more errors than LowHigh [$F(1,11)=4.3$, $P=0.062$]. On Noswitch runs, Congruency was not reliable ($P>0.2$), but Odd-Even trials had 0.9% more errors than LowHigh [$F(1,11)=5.6$, $P<0.04$]. No other effects were reliable.

Run-initial latency data appear in Fig. 4, and were analyzed with a $2 \times 2 \times 2 \times 2$ ANOVA on Position (P1, P2), Continuity, Congruency, and Task. There were main effects of Continuity [$F(1,11)=7.6$, $P<0.02$], and Position [$F(1,11)=16.7$, $P<0.003$], and the two interacted [$F(1,11)=10.2$, $P<0.01$]. Continuity also had a simple




Fig. 4 Data from the first two trials of a run (P1 and P2) in Experiment 2. “Restart cost” is P1 Noswitch–P2 Noswitch, and “start cost” is P1 Switch–P2 Switch. Error bars are $\pm$ RMSE for the Continuity effect

effect at P1 [$F(1,11)=9.7$, $P<0.02$]. Incongruent trials were 40 ms slower than Congruent trials [$F(1,11)=15.9$, $P<0.003$], with no Congruency $\times$ Position interaction ($F<1$). OddEven trials were 65 ms slower than Low-High trials [$F(1,11)=20.7$, $P<0.002$]. Task interacted with Congruency [$F(1,11)=6.0$, $P<0.04$], Position [$F(1,11)=24.7$, $P<0.001$], and the two together [$F(1,11)=7.0$, $P<0.03$].

Discussion

Within-run slowing survived several steps taken to reduce the success of and need for anticipating a task switch. Trial position within a run provided no information about cue arrival, and inserting the cue as a distinct event between trials should have minimized the incentive to predict its arrival. Also, with switches randomized, preparing for a switch as such would not have paid off on average. These results, together with those of experiment 1, shift the burden to anticipation models to provide an improved account of within-run slowing.

The error data do not echo the within-run increase of experiment 1, but do support the predicted Congruency $\times$ Continuity interaction – incongruent trials produced more errors, but only on Switch runs. This finding suggests a way to probe the underlying decay model.

Curvilinear decay, for example (see Fig. 6, below), would seem to predict that d' should vary within a run and thus a Continuity $\times$ Congruency $\times$ Position interaction. Future studies with additional power may detect such an effect.

Experiment 3

The goal in experiment 3 was to generalize within-run slowing to another, preferably familiar task-switching paradigm. A natural candidate was the alternating-runs paradigm of Rogers & Monsell (1995). Their experiment 6 contained runs of four trials instead of the usual two, and, indeed, “Inasmuch as there was any trend in RT over Trials 2 through 4, it was a 10-ms increase rather than a decrease.” An analysis excluding P1 “confirmed that this was not significant, $F<1$,” but it is not reported whether this analysis included a linear contrast across the three relevant positions. Their experiment 6 (henceforth “the original study”) is also apt because it was meant to test an interference-based interpretation of switch costs, though interference effects were interpreted differently there than they are here. The distinction is addressed in the discussion below.

Materials and methods

Ten Michigan State University undergraduates participated for course credit.

Stimuli were single characters, presented in white on a dark background. (In the original study they were two characters, but became single characters in Monsell, Yeung, & Azuma, 2000, and for simplicity this change was adopted here.) Each stimulus appeared in one of eight slices of a pie marked on the screen by light-gray lines, rotating clockwise one slice per trial at a radius of 2.8 cm from center of screen to center of character. Digit stimuli were 1 through 9, except 5, and letters were G, K, M, R, A, E, I, and U. RSI was 500 ms (increased from 450 ms in the original study), timed from key downstroke on the previous trial. An error was signaled immediately by a 100-ms visual flash (changed from an auditory signal in the original study), followed by a 1500-ms recovery interval and then the next stimulus. A block contained 20 runs of four trials and a session contained 20 blocks. Each block was followed by a display giving accuracy feedback and the number of the next block, as before. Instructions once again emphasized accuracy and did not mention speed.

The tasks were OddEven and ConsVow, with the current task indicated by which half of the screen the stimulus appeared in. The screen was divided either vertically or horizontally, with the dividing line being slightly thicker than the lines marking the other slices of the pie. The orientation of the dividing line (vertical or horizontal), the mapping from half of screen to task, and the mapping from response category to C or M were all randomized between participants.

Results

Results are shown in Fig. 5. The first six blocks of a session and the first two runs of a block were excluded, and error trials and trials after error trials were excluded from the latency measure. Data were analyzed with 3 $\times$ 2 ANOVAs on Position (P2 to P4) and Task (OddEven,




Fig. 5 Data from experiment 3. Within-run slowing is present in a paradigm modeled after Rogers and Monsell (1995), experiment 6. Markers with no line are standard deviations. Error bars are $\pm$ RMSE for the effect of Position (P2 to P4)

LowHigh). Latencies are again means of participant medians. For comparison with the original study, means of participant *means* for P1 to P4 were 826, 607, 622, and 634 ms.

For latencies, there was a main effect of Position [$F(2,18)=11.5$, $P<0.002$], and a linear trend [$F(1,18)=23.0$, $P<0.001$], accounting for 99% of the variance due to Position. No other trends were significant. There was an effect of Task [$F(1,9)=11.1$, $P<0.01$] (OddEven was 28.5 ms slower), but no Position $\times$ Task interaction ($F<1$). For standard deviations, there were no reliable effects. For errors, Position had no effect ($F<1$), but Task had a marginal effect [$F(1,9)=3.6$, $P=0.09$] (OddEven producing 1.5% more errors). The interaction was not reliable ($P>0.1$).

Discussion

Experiment 3 shows that within-run slowing can be found when memory for the task is optional – when stimulus set (letter or digit) and spatial location serve as cues on every trial. In the future, it will be important to ask what produced reliable slowing here when the effect was largely absent in the Optional conditions of experiment 1. Anticipation was at least possible here, as the timing of a switch was perfectly predictable, and spatial

cues were available to facilitate prediction. It may also be useful to replicate fully the stimuli of the original study, in which an irrelevant flanking character cued the wrong task, the correct task, or neither, and could be congruent or incongruent with the relevant character. This complex priming from the flanker could enhance within-run slowing, but could also attenuate it by amplifying the interference level in memory and thereby making spatial location more attractive as a source of information about the current task.

Whatever the cause of within-run slowing in the present experiment, the effect was highly reliable, and the changes in materials and design relative to experiments 1 and 2 suggest it might be rather general. Indeed, within-run slowing is easy to overlook, being not only small but also unexpected under the usual “reconfiguration” metaphor, with its tacit implication that mental states are stable once configured and would have no reason to evolve gradually over time. Perhaps, then, one might find it in other data on extended runs, by applying the appropriate trend analysis.

Experiment 3 also shows that within-run slowing survives an inter-stimulus interval. In the earlier experiments, within-run slowing could have been an artifact of inadequate preparation time between trials having a cumulative effect of some kind. Here, however, the 500 ms RSI was about as long as the system seems able to make use of in preparing to perform a task (Rogers & Monsell, 1995).

A decay account of slowing in experiment 3 suggests a very different view of interference than taken in the original study. The goal there was to test the task-set inertia hypothesis (Allport et al., 1994; Allport & Wylie, 2000) which, like functional decay theory, posits interference from previous task sets. However, unlike functional decay theory, task-set inertia has interference causing residual switch cost directly. Rogers and Monsell (1995) reasoned that on this view, switch cost should persist over several trials, declining gradually as the previous task set decays. Because switch cost was borne entirely by P1, they rejected interference as the cause of residual switch cost, and with it the task-set inertia hypothesis. Here, the proposal is that interference effects are mediated by other processes, notably decay (measured as within-run slowing). Another mediator is a strengthening process (introduced below) that may account for much of the latency cost at the start of a run. Thus, the view here is that interference is more than a nuisance reflected in switch costs – it is an architectural constraint reflected in adaptive processes like decay and strengthening.

General discussion

All three experiments showed within-run slowing and individually supported functional decay theory in specific ways. Slowing can be greatly attenuated simply by use of non-overlapping stimulus sets, showing that

memory moderates the effect (experiment 1). The effect survived procedural changes meant to limit the opportunity and the need to anticipate the next instructional cue (experiment 2). Finally, the effect transfers to at least one other, moderately distinct paradigm (experiment 3). Thus, indications are that within-run slowing is robust and general enough that task-switching theory will have to contend with it in some fashion. It is also important to note that slowing does not conflict with practice effects: Slowing within a run can coexist with speedup over a session (Altmann & Gray, 2002).


What of the sullied reputation of decay as a model of forgetting (e.g., McGeoch, 1932)? Cast as a process that affects *distractors* rather than targets, decay has actually appeared in other guises, such as “stimulus fluctuation” (Estes, 1955) and “unlearning” (Postman, 1971). Empirical evidence for a distinct decay process has also accumulated, in task switching (Koch, 2001; Meiran, 1996) and in memory studies generally (Baddeley & Scott, 1971; Reitman, 1974). Indeed, a pillar of interference theory – the probe-digit data of Waugh and Norman (1965) – contains evidence for gradual decay in the form of a presentation rate by serial position interaction (Altmann & Schunn, 2002). Overall, then, decay is quite as common as it is disreputable.

Functional decay theory addresses task retention, but what can it tell us about task encoding? In particular, can it tell us anything about switch cost, the focus of most task-switching studies to date? In terms of the signal detection analysis of Fig. 1, the question reduces to how a task set arrives at its initial activation level – the mean, say, of the density in the top panel of Fig. 1, then it must have activation to begin with, which must itself be the product of some encoding process. A model of this encoding process would be a candidate for explaining the substantial slowdown on P1, the initial trial of a run.

An activation model of encoding and decay appears in Fig. 6, which shows mean activation of a task set over time. Fig. 6 bears a direct relation to Fig. 1, which also showed task-set activation at different points in time. In Fig. 1 the emphasis was on activation variability and the ensuing signal-detection problem. In Fig. 6 the emphasis is on tracking (mean) activation over time as a function of an underlying process. The process is described by the equation $\text{activation} = \ln(r/\sqrt{t})$, adapted from ACT-R memory theory (Anderson & Lebiere, 1998). r is how often the task set has been retrieved over its lifetime. In Fig. 6, each marker on the curve is a retrieval (this is a finer-grained view of retrieval than is taken in other task-switching models; e.g., Mayr & Kliegl, 2000). t is the span of this lifetime, from trace encoding to present. Thus the model assumes an instance representation (Logan, 1988), in which each task set codes a single, episodic exposure to a cue.

Under this model, a task set passes through three phases during its lifetime: strengthening, use, and disuse (top of Fig. 6). During strengthening, activation builds up rapidly as the system “pays attention” to the cue, operationalized as a closely spaced sequence of retrievals (each indicated by a marker on the curve). This “massing” of retrievals is possible because the system is fully engaged with the cue; no other processing is called for at this point. Following Altmann and Trafton (2002), who develop an alternative to the ACT-R construct of self-activating goals and apply it to problem solving data,

Fig. 6 Task set activation over time, from onset of an instruction (Cue 1) through five trials (P1 to P5). Task-set retrievals are marked by filled circles. Activation = $\ln(r/\sqrt{t})$, with r = retrieval count and t = time since encoding. During strengthening, retrievals are massed (every ~100 ms). During use, retrievals are spaced (every ~600 ms), each boosting activation but not enough to offset the downward trend. During disuse, retrievals stop. Because of decay during use, d' is positive once the next task set (dashed ink) is fully encoded, as in Fig. 1



the assumption in Fig. 6 is that strengthening retrievals can occur every 100 ms. When strengthening is complete, the task set's (mean) activation is as high as it will ever be. This mean activation is also the mean of the density function in the top panel of Fig. 1.

Next, during the use phase, the task set is retrieved whenever the system needs to know what task to perform, once per trial, say, after the system has encoded the stimulus and begins response selection. Retrievals are now less frequent because there is other work to be done, like stimulus encoding and response selection; a trial takes roughly 600 ms, so the retrieval rate is now a sixth of what it was during strengthening. Each retrieval boosts activation slightly, but the abrupt shift from massed to spaced produces an overall downward trend, "serrated" across trials.

Finally, the task set falls into disuse, once it is superseded by a new task set. Now the old task set decays more rapidly. If it is retrieved again, this reflects a memory error caused by transient activation noise making it more active than the current task set (in terms of Fig. 1, such an error would correspond to a false alarm). This disuse phase overlaps in time with the strengthening phase of the new task set (in dashed ink in Fig. 6). At its peak activation, the new task set is separated from the old task set by the same d' that separates the densities in the bottom panel of Fig. 1.

This activation model has a number of attractive properties. First, it maps functional decay to a memory process, one that also accounts for other memory strategies like maintenance rehearsal (Anderson & Lebiere, 1998). Thus, no new mechanisms are necessary to meet the constraints of functional decay theory. Second, it reduces task set "reconfiguration" to a memory process also, one that may cause much of the variance on P1 – "switch cost," broadly defined. Indeed, time-extended strengthening offers one account of the otherwise-puzzling phenomenon of "restart cost" (Allport & Wylie, 2000), the marked slowing on P1 caused by a cue that does not actually switch the task. Fig. 4 illustrates restart cost with data from experiment 2, where a cue was as likely to "restart" the old task as it was to "start" a new one. Restart cost is P1 Noswitch minus P2 Noswitch, and is a substantial 110 ms (see experiment 2, Results, for analysis). The source of this cost, under the activation model, is the strengthening needed regardless of cue type. The previous task set will always have decayed, so a task set for the current cue must always be strengthened to maintain performance through the following run. Strengthening that cannot be completed while the cue is visible will spill over into P1, manifesting as restart cost in the Noswitch case and as start cost (P1 Switch minus P2 Switch) in the Switch case. Figure 6 shows this spillover schematically, with strengthening continuing into P1.

There are at least three limitations to this strengthening account of costs at the start of a run. First, other plausible accounts have yet to be ruled out. Restart cost and some amount of start cost could be due to the "task

switch" of shifting from processing a cue to processing a stimulus. They could also reflect a temporary shift in the response criterion to cautious responding, in those cases where P1 error decreases (Allport & Wylie, 2000). Second, as presented here the model does not explain why start cost should be greater than restart cost, as in Fig. 4. The difference between start and restart cost – a "pure" switch cost, net of task-neutral costs on P1 – might reflect prolonged strengthening in the Switch case, or foreshortened strengthening in the Noswitch case due to priming from the previous task set. A computational model incorporating the latter hypothesis has been implemented (Altmann & Gray, 2000) but awaits a fuller exposition. Third, the model does not explain why, in the general case, strengthening should spill from the instructional cue into P1. In experiment 2, the duration of the cue was only 500 ms, which may not have left enough time for strengthening to complete. On the other hand, longer preparation periods than this produce little extra reduction in "residual" cost on P1 (Rogers & Monsell, 1995). A strengthening account will clearly have to be more elaborate than the simple unitary stage depicted above if it is to account for this stubborn residual. One possibility is that strengthening and the failure-to-engage account of residual switch cost (De Jong, 2000) tell different parts of the same story. The offset of strengthening may be stochastic, like activation itself, and sometimes leave work to be done on P1 while a trace of the cue still lingers in perceptual memory. Under these conditions, a residual strengthening stage actually improves the cost/benefit ratio of strengthening overall (Altmann & Gray, 2000).

Despite the abstract level at which this activation model is specified, it carries a substantial theoretical advantage in being part of a larger, integrated set of mechanisms rather than a standalone account of a single effect. From this perspective, it makes a good starting point. Further development will depend on estimation of retrieval frequency and cycle time parameters, and elaboration to account for "pure" switch cost and residual processing on P1.

The signal detection model of memory (Fig. 1) has implications beyond those developed here, some interesting and some moderately awkward. One is that the decay rate (activation lost per trial) should vary inversely with runlength, if d' stays relatively constant. This effect has been found (Altmann & Gray, 2002), but a corollary is that the relationship is non-linear, escalating to very rapid decay indeed for full interleaving (ABABAB). Whether decay can be rapid enough to support such interleaving, or whether interleaving requires additional memory strategies, or whether a single alternative model can explain all within-run effects, will be important questions to address. Finally, other phenomena that seem like candidates for integration using the signal detection model include backward inhibition (Mayr & Keele, 2000), stimulus priming (Allport & Wylie, 2000), and mixing cost (e.g., Los, 1996).

In conclusion, conventional wisdom on executive control has “modules in our brains” being “reconnected and reconfigured” to produce fully formulated mental states (Monsell & Driver, 2000a). An alternative to this mechanical, all-or-none metaphor is that control arises from ordinary activation-based memory processes. As developed here, this memory approach captures the broad outlines of task switching behavior end to end, not stopping with the switch trial but continuing through the entire run. Filling in the details of this framework may quickly bear more fruit than complicating the job description of a mechanical homunculus to explain effects like within-run slowing.

Acknowledgements Thanks are due to the Royal Netherlands Academy of Arts and Sciences and to the organizers and participants of the Control of Cognitive Processes colloquium. Thanks also to Bernhard Hommel, Iring Koch, and an anonymous referee for most thoughtful reviews, to Kate Arrington and Dan Kuehner for their comments on the paper, and to Katie Frank, Kara DeVlieger, and Katie Syswerda for help with data collection.

References

- Allport, A., Styles, E. A., & Hsieh, S. (1994). Shifting intentional set: exploring the dynamic control of tasks. In C. Umiltà & M. Moscovitch (Eds.), *Attention and performance IV* (pp. 421–452). Cambridge, MA: MIT Press.
- Allport, A., & Wylie, G. (2000). Task-switching, stimulus-response binding, and negative priming. In S. Monsell & J. S. Driver (Eds.), *Control of cognitive processes: attention and performance XVIII* (pp. 35–70). Cambridge, MA: MIT Press.
- Altmann, E. M., & Gray, W. D. (2000). An integrated model of set shifting and maintenance. In N. Taatgen & J. Aasman (Eds.), *Proceedings of the 3rd international conference on cognitive modeling* (pp. 17–24). Riethoek, The Netherlands: Universal Press.
- Altmann, E. M., & Gray, W. D. (2002). Forgetting to remember: the functional relationship of decay and interference. *Psychological Science*, 13, 27–33.
- Altmann, E. M., & Schunn, C. D. (2002). In *Proceedings of the 24th annual meeting of the Cognitive Science Society*. Hillsdale, NJ: Erlbaum.
- Altmann, E. M., & Trafton, J. G. (2002). Memory for goals: an activation-based model. *Cognitive Science*, 26, 39–83.
- Anderson, J. R., & Lebiere, C. (Eds.). (1998). *The atomic components of thought*. Hillsdale, NJ: Erlbaum.
- Baddeley, A. D., & Scott, D. (1971). Short term forgetting in the absence of proactive interference. *Quarterly Journal of Experimental Psychology*, 23, 275–283.
- De Jong, R. (2000). An intention-activation account of residual switch costs. In S. Monsell & J. Driver (Eds.), *Control of cognitive processes: attention and performance XVIII* (pp. 357–376). Cambridge, MA: MIT Press.
- Estes, W. K. (1955). Statistical theory of spontaneous recovery and regression. *Psychological Review*, 62, 145–154.
- Fagot, C. D. (1994). Chronometric investigations of task switching. Ph.D. thesis, Psychology Department, University of California, San Diego.
- Gopher, D., Armony, L., & Greenspan, Y. (2000). Switching tasks and attention policies. *Journal of Experimental Psychology: General*, 129, 308–339.
- Hommel, B. (2000). The prepared reflex: automaticity and control in stimulus-response translation. In S. Monsell & J. Driver (Eds.), *Control of cognitive processes: attention and performance XVIII* (pp. 247–273). Cambridge, MA: MIT Press.
- Koch, I. (2001). Automatic and intentional activation of task sets. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 27, 1474–1486.
- Kramer, A. F., Hahn, S., & Gopher, D. (1999). Task coordination and aging: explorations of executive control processes in the task switching paradigm. *Acta Psychologica*, 101, 339–378.
- Logan, G. D. (1988). Toward an instance theory of automatization. *Psychological Review*, 95, 492–527.
- Los, S. A. (1996). On the origin of mixing costs: exploring information processing in pure and mixed blocks of trials. *Acta Psychologica*, 94, 145–188.
- Luce, R. D. (1986). *Response times: their role in inferring elementary mental organization*. New York: Oxford.
- Mayr, U., & Keele, S. W. (2000). Changing internal constraints on action: the role of backward inhibition. *Journal of Experimental Psychology: General*, 129, 4–26.
- Mayr, U., & Kliegl, R. (2000). Task-set switching and long-term memory retrieval. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 26, 1124–1140.
- McGeoch, J. A. (1932). Forgetting and the law of disuse. *Psychological Review*, 39, 352–370.
- Meiran, N. (1996). Reconfiguration of processing mode prior to task performance. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 22, 1423–1442.
- Monsell, S., & Driver, J. (2000a). Banishing the control homunculus. In S. Monsell & J. Driver (Eds.), *Control of cognitive processes: attention and performance XVIII* (pp. 3–32). Cambridge, MA: MIT Press.
- Monsell, S., & Driver, J. (Eds.). (2000b). *Control of cognitive processes: attention and performance XVIII*. Cambridge, MA: MIT Press.
- Monsell, S., Yeung, N., & Azuma, R. (2000). Reconfiguration of task-set: is it easier to switch to the weaker task? *Psychological Research*, 63, 250–264.
- Newell, A., & Rosenbloom, P. S. (1981). Mechanisms of skill acquisition and the law of practice. In J. R. Anderson (Ed.), *Cognitive skills and their acquisition*. Hillsdale, NJ: Erlbaum.
- Postman, L. (1971). Transfer, interference, and forgetting. In J. W. Kling & L. A. Riggs (Eds.), *Woodworth & Schlosberg's experimental psychology*. (3rd ed., pp. 1019–1132) New York: Holt, Rinehart, and Winston.
- Reitman, J. S. (1974). Without surreptitious rehearsal, information in short-term memory decays. *Journal of Verbal Learning and Verbal Behavior*, 13, 365–377.
- Rogers, R. D., & Monsell, S. (1995). Costs of a predictable switch between simple cognitive tasks. *Journal of Experimental Psychology: General*, 124, 207–231.
- Spector, A., & Biederman, I. (1976). Mental set and mental shift revisited. *American Journal of Psychology*, 89, 669–679.
- Waugh, N. C., & Norman, D. A. (1965). Primary memory. *Psychological Review*, 72, 89–104.
- Wickelgren, W. A. (1977). Speed-accuracy tradeoff and information processing dynamics. *Acta Psychologica*, 41, 67–85.