Using a Model to Compute the Optimal Schedule of Practice

Philip I. Pavlik, Jr. and John R. Anderson
Carnegie Mellon University

By balancing the spacing effect against the effects of recency and frequency, this paper explains how practice may be scheduled to maximize learning and retention. In an experiment, an optimized condition using an algorithm determined with this method was compared with other conditions. The optimized condition showed significant benefits with large effect sizes for both improved recall and recall latency. The optimization method achieved these benefits by using a modeling approach to develop a quantitative algorithm, which dynamically maximizes learning by determining for each item when the balance between increasing temporal spacing (that causes better long-term recall) and decreasing temporal spacing (that reduces the failure related time cost of each practice) means that the item is at the spacing interval where long-term gain per unit of practice time is maximal. As practice repetitions accumulate for each item, items become stable in memory and this optimal interval increases.

Keywords: memory, practice, paired-associate, efficiency, spacing effect

Since the end of the 19th century researchers have tried to describe the best way to practice to enhance learning and retention. This work began around the time of Ebbinghaus (1913/1885), who focused on how the history of practice for items controls the future strength or retrievability of memories formed. Ebbinghaus' research helped to establish methods for the scientific study of memory and demonstrated memory effects that are still studied today. Although many of his investigations demonstrated the effects of frequency and recency, first listed as principles of association by Thomas Brown in the early 19th century (Murphy & Kovach, 1972), he is also credited with uncovering the "spacing effect" because he discovered that by interspersing sleep periods between study sessions subsequent performance was improved compared to a contiguous session.

Since Ebbinghaus' (1913/1885) results, researchers have focused on the obvious educational implication of these findings, with particular emphasis on the spacing effect because of its apparent ability to produce more durable learning. Indeed, a central theoretical question in this paper, whether long temporal intervals between practices are always optimal for learning or whether initial practice requires shorter spacing, was first debated by Steffens (1900) and Jost (1897) according to a review on the literature by Ruch (1928). Although Steffens advocated wider spacing overall, Jost advocated narrower initial spacing, and the debate on this issue continues today. It is in the context of this more than century long effort to understand a general solution to the question of optimal practice scheduling that we will begin by introducing an understanding of the "spacing effect" and its implications.

The "spacing effect" is the often found and widely recognized learning advantage of having intervals of time between repetitions of a skilled performance, particularly when the skilled performance involves factual recall (Dempster, 1996). Although some authors such as Underwood, Kapelak, and Malmi (1976) have distinguished the spacing effect from the lag effect (which specifies that increasing the duration of the spaced lag between repetitions increase the benefit of spacing), this paper assumes that spacing effects are function of lag, particularly for fact learning. Whereas the spacing effect in fact learning has been well known since Ebbinghaus (1913/1885), more recent work has clarified how it functions by demonstrating that there are two important interactions to consider when applying spacing:

1. Several authors have shown that as retention interval increases, the benefit of spacing is larger (implying that spacing improves the durability of learning). This retention-interval-by-spacing interaction has been shown over short time scales ranging from seconds and minutes (Glenberg, 1976; Peterson, Wampler, Kirkpatrick, & Saltzman, 1963) and over days and months (Bahrick, 1979; Fishman, Keller, & Atkinson, 1969; Pashler, Zarow, & Triplett, 2003; Pavlik, 2005). Because the interaction of spacing and retention interval is often quite large in the above studies (e.g., Pashler et al., 2003), it is crucial to account for when scheduling practice.

Philip I. Pavlik, Jr., Human Computer Interaction Institute, Carnegie Mellon University, Pittsburgh; John R. Anderson, Department of Psychology, Carnegie Mellon University, Pittsburgh.

Correspondence concerning this article should be addressed to Philip I. Pavlik, Jr., Human Computer Interaction Institute, Carnegie Mellon University, 5000 Forbes Ave, Pittsburgh, PA 15213. E-mail: ppavlik@andrew.cmu.edu

2. Researchers have shown that spaced practice has a cumulative effect such that each additional spaced practice provides an additional advantage (spacing by practice quantity interaction) (Pavlik & Anderson, 2005; Underwood, 1970). This interaction with practice quantity highlights the continuing importance of spaced practice as practice accumulates.

The main effect of spacing and these interactions strongly suggest that spacing is important to take advantage of when scheduling practice, that spacing does continue to apply when frequency increases, and that these benefits are a function of the retention interval.

On the other hand, wide spacings can result in a significant slowing of the rate at which material can be learned because of the greater forgetting between repetitions. This is particularly true for practice procedures that include a test of the material being learned for each trial. For example, a typical practice procedure involves a test followed by a review presentation if the item cannot be recalled (referred to in this paper as a recall-or-restudy trial or a drill trial). At short spacing, the participant is likely to quickly recall the item and forgo the need for a review presentation (benefits of recency). This drill task should be an effective way to train items because testing has often been shown to be a particularly effective means of practice (e.g., Carrier & Pashler, 1992). Thus, it is necessary to have a method for finding the spacing that produces the most learning gains with the least slowing of learning rate.

Other researchers of optimal scheduling issues have tried to move forward on the problem by comparing alternative schedules in experiments to determine empirically what sorts of schedules are optimal for learning (Balota, Duchek, Sergent-Marshall, & Roediger, 2006; Cull, 2000; Cull, Shaughnessy, & Zechmeister, 1996; Landauer & Bjork, 1978; Rea & Modigliani, 1985). These studies provide data that are useful in helping us understand the optimal scheduling of practice, but they do not provide a systematic method of computing schedules such as is presented in this paper. For example, Karpicke and Roediger (2007) tested the theory that expanding spacing is optimal by comparing a condition with expanding spacing of practice (where practice is initially narrow but increases with each new repetition) to a condition with equally spaced practice. In these comparisons, they showed that the expanding spacing schedule only produced an advantage after a short retention interval of about 10 minutes, while the equal spacing condition produced superior recall after a long-term interval of 2 days. Although interesting, unfortunately, the result is hard to generalize because the authors compared only two contending spaced schedules (5–5–5 intervening trials, the even spacing condition and 1–5–9 intervening trials, the expanding spacing condition). Another limitation of prior experimental work is that it failed to control time on task (cost per trial) as a component of the learning rate (Balota et al., 2006; Cull, 2000; Cull et al., 1996; Landauer & Bjork, 1978; Rea & Modigliani, 1985). In all of these studies, the unit of analysis that defines the quantity of learning is the trial rather than time. Because of this lack of control for time on task, these studies cannot easily account for the speed advantage of narrower spacing.

To address this limited generalizability of an experimental approach, some researchers have taken a modeling approach. A modeling approach tries to create a general model to quantify the effects of prior practice through predictions of performance. A model such as this (if accurate) should allow one to search through the space of alternative practice schedules for each item being learned to choose the optimal schedule for each item. Of course, because practice often involves testing each item, a modeling approach can also use the information from these assessments to tune the model predictions as practice accumulates. Although such performance tracking adds considerable power to a modeling approach, predictions for the optimal schedule become more complex since the schedule will not be the same for each item but rather depend on the individual history of correctness for each item in addition to the recency, frequency, and spacing of prior successes or failures. Further, as will become clear when the ACT-R model is introduced, with a modeling approach we can predict the time costs (latency) of actions and therefore consider the time on task consequences of our scheduling decisions in a way that is difficult to do from experimental results.

During the 1960s and early 1970s cognitive psychologists were actively pursuing this model-based approach to optimal scheduling (Atkinson, 1972b; Atkinson & Crothers, 1964; Atkinson, Fletcher, Lindsay, Campbell, & Barr, 1973; Atkinson & Paulson, 1972; Dear, Silberman, & Estavan, 1967; Fishman et al., 1969; Groen & Atkinson, 1966; Karush & Dear, 1966; Laubsch, 1971; Lorton, 1973; Pimsleur, 1967; Smallwood, 1962, 1971). Although some of this work has persisted (e.g., the Pimsleur Method is still sold commercially), for the most part the potential of these systems was not realized and research declined in the early 1970s. After 1973, Wozniak and Gorzalanczyk (1994) is the next example of work that uses a model-based practice scheduling algorithm to tutor fact learning.

Indeed, this decline may be presaged by Smallwood (1962) in which the importance of considering the unit of analysis was explained. Despite the emphasis placed by Smallwood (and endorsed by Atkinson, 1972a) on the need to use time as the unit of analysis (rather than the trial), only Smallwood (1962, 1971) attempted to use models that captured the time on task cost of different scheduling decisions. Unfortunately, Smallwood used a very simple model that failed to capture important factors like spacing, so his results do not provide more than this surface insight that time on task must be accounted for.

Atkinson's Experiment

Of this prior research, Atkinson (1972b) is often cited as an example of the large potential benefit of practice schedule optimization. Therefore, it is useful to review this result carefully. In the paper, Atkinson compared the four procedures listed below to optimize the learning of 84 German-English vocabulary pairs over the course of 336 practice trials (percentage recall after 1 week in parentheses):

1. Participant self-selection of pairs to study (58%).

2. Random selection of pairs (38%).

3. Selection of the weakest pair to study according to a three-state Markov model with five parameters estimated for all pairs (54%).

4. Selection according to the same Markov model with the five parameters estimated for each pair (79%).

Unfortunately, there were problems with Atkinson (1972b), which may have inflated these results. The first problem was that in all conditions each trial began with a 10-s presentation of a cue list of 12 German words (Atkinson split the German words into 7 lists of 12 pairs each that cycled so repetitions were always spaced by at least six intervening trials). From each list, a word was selected depending on condition. In the self-selection condition, the test pair was selected by participants, in the random condition, the test pair was randomly selected, and in the optimized conditions, the test pair was selected according to the Markov model. After the 10 second viewing of the cue list in all conditions, the selected word was prompted, and the participant was expected to type the English translation. Each trial concluded with a presentation of the correct answer (regardless of participant response), which served as feedback when the answer supplied was incorrect. According to Atkinson, this procedure took approximately 20 seconds per trial, which seems like a rather slow procedure.

One problem with this procedure is that it is unclear what participants did during the 10-s viewing of the cue lists. Participants may have spent the 10 seconds covertly practicing the responses for any of the 12 cues that they currently remembered. Further, Atkinson's (1972b) optimization conditions (particularly the five parameters-per-pair condition) introduced the more difficult items earlier because they had different parameters and thus the participants in these conditions could focus their covert practice on these items early. Because the possibility of extra covert practice favored these most difficult items in the Atkinson (1972b) optimization conditions, it may have biased the results in favor of these conditions. The seven cue lists were each presented 48 times for a total of 3360 seconds of covert practice per participant so the bias may have been large.

Generally, the procedures Atkinson (1972b) used ignored concerns for efficiency. For example, the inclusion of the self-selection condition is problematic because its inclusion forced the rather unnatural process of previewing the lists in all conditions. If one wanted to maximize amount learned in a fixed time, one would have allocated this time for more practices. Similarly, allowing participants to proceed at their own speed by forgoing review after a correct response would have meant more recall-or-restudy practice opportunities in conditions with fewer study errors if each condition was allowed a fixed duration.

A final problem with Atkinson (1972b) is that the model used has no plausible representation for long-term forgetting. According to the Atkinson model, once a pair has been "permanently learned" it cannot be forgotten. In contrast, the ACT-R model compared in this paper can predict forgetting over intervals of months and possibly years (Pavlik & Anderson, 2005). For the reasons above, whereas the Atkinson (1972b) result is interesting, it is reasonable to suspect that one can produce both better learning conditions and a better test of the benefits of such conditions.

ACT-R Modeling System

We used the ACT-R (Adaptive Character of Thought–Rational; Anderson & Lebiere, 1998) modeling system to predict the effect of prior practice. Although the ACT-R model makes a fundamental distinction between production rules (procedures with an if-then structure) and declarative memory chunks (bits of learned information), this paper focuses only on using the set of equations from ACT-R that describe the strength of a memory chunk as a function of practice. The ACT-R declarative memory equations are particularly appropriate for our method because they capture both correctness and latency of performance as a function of prior practice. Both dependent measures are necessary for the modeling approach we will take since they are fundamental to our method for computing the expected efficiency of a proposed practice. In contrast, Atkinson's (1972b) Markov model (that controls a practice condition that we will compare with the new optimization method condition) only predicts correctness as a function of practice. Because of this, the Atkinson model as it stands could not be used with the method we explain here unless it was modified, perhaps by predicting latency as a function of expected correctness, to be adequate. This model substitutability (ability to substitute another model with the necessary dependent measures) of the method highlights that this paper is not about a particular model of practice; rather, this report is about how to apply any suitable quantitative model to the problem of practice schedule optimization.

The Experiment

We performed an experiment to test our approach to optimizing learning. In this experiment, 60 participants learned a set of 180 Japanese-English vocabulary words during learning sessions on a Monday, Wednesday, and Friday. Each participant was in one of three learning conditions for these three sessions. A final fourth session on the following Friday was used to assess correctness and latency effects of the three learning conditions.

This experiment accomplishes several goals. First, it provides a test of three ideas about how to optimize the schedule of learning. The first condition tests the schedule optimization method explained here using a version of the ACT-R memory model, the second condition tests Atkinson's (1972b) schedule optimization algorithm, and the third condition tests a simple flashcard procedure. This choice of conditions was based on prior results. In an early experiment, we had shown conditions with maximal spacing performed worse than the new method (Pavlik, 2007). Because this result showed maximal spacing was not optimal, the flashcard condition (Condition 3) was designed to provide a control condition that, whereas providing relatively wide spacing, did not maximize spacing in a way that we thought would almost certainly result in inefficient learning.

Second, the experiment provides a larger scale test of the ACT-R based practice-scheduling method than in Pavlik (2007). For the current experiment, the learning set included 180 Japanese-English word-pairs, which were learned over three training sessions and evaluated after a 1-week retention interval. Compared to Pavlik, this reflects nearly doubling the number of word pairs learned, tripling the number of learning sessions and adding 6 days to the retention interval. This scalability is particularly important when one considers the long durations of retention and large numbers of items that might need to be learned for some domains.

Third, the experiment allows the incorporation of additional components in the ACT-R based practice-scheduling algorithm, which were not included in (Pavlik, 2007). For instance, Pavlik assumed that there was no difference between test trials and

passive study trials of each pair despite the fact that test trials are typically found to be more effective (e.g., Carrier & Pashler, 1992). The current experiment, reflecting the results of Pavlik (2006), incorporated a model of this difference between study and test practice into the decision as to whether to deliver a study-only trial or recall-or-restudy trial.

Fourth, the experiment provides data to understand long-term memory processes better, which will aid in both model and theory development. In terms of model development, the experiment should allow for a refinement of the model's parameters to improve future performance. In terms of theoretical progress, this experiment tested an interesting hypothesis that the difficulty of the learning context (measured by overall percent correct during learning) may reduce the encoding or recallability of individual items being learned. In this experiment, we operationalized difficulty by equating it with the number of errors that a practice procedure produced. By extension, the difficulty of an item can be characterized by the quantity of errors with that item relative to other items practiced with the same schedule. Using this operational definition it becomes clear that the sources of item difficulty are manifold and include a lack of prior learning, presence of interference, complexity of the material, time pressure, poor contextual cueing, forgetting, reduced motivation, and student ability factors.

Whatever the cause, this effect of the context on learning deserves recognition because it has important implications for conclusions about within-subjects tests of different learning methods. For example, if one is comparing an easy procedure and a hard procedure using a within-subjects design, the easy procedure will tend to suffer a disadvantage from higher difficulty of the learning context compared to if it were tested alone, whereas the difficult procedure will gain an advantage from lower difficulty of the learning context compared to if it was tested alone. This confound seems to occur in Pavlik and Anderson (2005) and Pashler et al. (2003) for their within-subjects comparisons of wide and narrow spacing conditions. A better test of the independent utility of either procedure would compare effects between-subjects. To see whether this issue is significant, a between-subjects subcondition was included in the following experiment. An additional 28 word pairs were used in this subcondition.

Experiment Design

In the main between-subjects comparison of this experiment, three algorithms for optimizing learning were tested. The first condition tested a new method of scheduling based on an extended ACT-R model. This optimization condition used the model to derive decision criteria for when to present each pair for practice. An algorithm based on these criteria was then developed. This algorithm used a model of each student to keep a running estimate of the memory strengths of pairs during learning to determine whether they met the decision criteria for scheduling.

The second condition was a replication of Atkinson (1972b) in which he used a Markov model to schedule practice. Although this was a replication of the Atkinson Markov model and applied his same method to select items for practice, the testing of the algorithm had many differences from the original paper. For instance, in this replication there were a larger number of sessions, shorter trial durations, and recall-or-restudy trials instead of test-and-study

trials (Atkinson's procedure included feedback even when responses were correct). However, none of these differences should affect the applicability of the Atkinson model and method. In this condition, the first presentation of each word was a study-only presentation of the pair; subsequent presentations were always recall-or-restudy trials. It is useful to note that functionally the test-and-study procedure used by Atkinson and others is equivalent to the recall-or-restudy procedure used here because a test after successful recall has been shown to have no meaningful effect besides its time cost (Pashler, Cepeda, Wixted, & Rohrer, 2005; Pavlik, 2006).

The third condition consisted of a flashcard procedure in which the 180 word pairs of the learning set were split into six "decks" of 30 pairs each. This flashcard control condition was designed to match what a naive learner might do given the task of memorizing these word pairs. It seems somewhat intuitive to suppose that this learner would not cycle through the whole list of 180 word pairs, but would rather split the pairs into some manageable deck size and 30 pairs per deck seemed a manageable quantity. In this condition, participants were presented with pairs from each deck until each word pair in a deck was recalled correctly once. The six decks cycled so every word in each deck needed to be responded to correctly before the next deck began. After the last deck, the procedure began again with the first deck, cycling the decks as many times as possible during the three learning sessions. The first presentation of each word (during the first pass through each deck) was a study-only presentation of the pair; subsequent presentations were always recall-or-restudy trials.

Each of these three conditions involved three learning sessions, each of which was timed to last exactly 1 hour. These three sessions occurred at the same time on a Monday, Wednesday, and Friday. At the same time on the following Friday, the participants returned for an assessment session that did not vary by condition. This final session included the entire set of 208 words (180 words in the three conditions plus the 28 words in the within-subjects difficulty subcondition described below) delivered in random order twice. For the second session, the first 104 of total 208 items delivered in the first pass were rerandomized into the first 104 positions of the second pass, with an analogous procedure for the second half of the items. This insured that spacings did not vary as much as they would for randomizing each pass independently, whereas still preventing sequential order effects from being an issue. All of these trials were delivered as recall-or-restudy trials as described in the procedure section.

The difficulty of context subcondition described earlier was designed to look at how learning or recall might be affected by the overall context of learning. Specifically, the hypothesis was that a context with a high level of correctness would be more conducive to learning. To test this, 28 pairs were randomized into a presentation schedule that did not vary for the three training conditions. To do this, the 28 pairs were tested according to 28 prespecified patterns of study and recall-or-restudy practice that were independent of the scheduling of the remaining items according to condition. These trials were scheduled by randomizing the first (or only) practice for each pair for each session to occur between 10 and 40 minutes into the session with subsequent practices occurring according to the prespecified pattern.

There were also comparisons of item differences nested within each of the three training conditions. Within the ACT-R training

algorithm condition, a comparison was performed to determine how important it was to have prior estimates of item difficulty. To do this comparison the 100 Japanese-English pairs for which there existed prior data from Pavlik (2007) were randomly divided for each participant into two halves, one half had their schedules optimized with a prior parameter to estimate difficulty and the other half did not use prior estimates. For the Atkinson training condition, the nested item analysis compared the performance on the 100 old word pairs for which five individual Atkinson parameters were available for each pair with the performance on the 108 new word pairs, which must use the overall parameter set since it was impossible to compute prior parameters without data. For the flashcard condition, a nested item analysis was completed to compare learning for the 100 word pairs used in previous experiments and the 108 new pairs used in this experiment.

In all conditions, before the first learning session, the experimenter read a short (5-min) one-page description of the keyword mnemonic strategy (see Atkinson, 1975). Participants were encouraged to use this method if they wanted. One reason for this simple training was to reduce the variability among subjects because a reduction in variability should both increase the effect size of the results and improve the accuracy of the model (and therefore improve the efficiency of the optimization method).

Materials

The stimuli were 208 Japanese-English word pairs. English words were chosen from the MRC Psycholinguistic database such that the words had familiarity ratings with M of 534 (SD 48), and had imagability ratings with M of 481 (SD 59). These ratings were composed according to procedures described in the MRC Psycholinguistic Database manual (Coltheart, 1981). The overall MRC database means for familiarity and imagability are 488 (SD 120) and 438 (SD 99), respectively, so the words chosen had higher familiarity and imagability ratings than the database averages. Japanese translations (from the possible Japanese synonyms) were chosen to avoid similarity to common English words. English words averaged 4.17 (SD 0.6) letters, and Japanese words average 5.61 (SD 1.3) letters. Japanese words were presented using English characters. Word pair order of introduction and assignment to conditions was randomized individually for each participant.

Procedures

Participants were scored for motivational purposes, receiving 1 point for each correct response and losing 1 point for each incorrect response. Failing to provide a response, either by time-out or providing a blank response, resulted in a 0 score. Participants were paid $55 to $70 depending on their score.

The experiments exclusively used two different trial types: study-only trials and recall-or-restudy trials. The study-only trials (used for initial practice before testing) were cued with the prompt "Ready" for 0.5 seconds after which the pair was presented for 4 seconds for study by the participant. The recall-or-restudy trials also began with a 0.5-s prompt of the word "Ready" that was followed by presentation of the Japanese word on the left side of the screen. Participants typed the English translation on the right. If no response was made or if an in-process response was deleted by backspacing, the program timed-out in 6 seconds (in-process

responses were given unlimited time for completion). If correct, the response was followed by a 0.5-s presentation of the word "Correct" and the next trial began. If incorrect, a study presentation for the word (which was introduced by the word "Study" for 0.5 seconds) was given for 3 seconds. These study-only and recall-or-restudy trial parameters were fixed across conditions.

For the benefit of participants (to reduce fatigue and improve motivation), the three 1-hr learning sessions were delivered in blocks of 30 trials each. Similarly, the final testing session was split into two blocks of 208 trials each. Between blocks, participants continued by pressing the space bar when they were ready. Few participants paused at these opportunities for rest.

Participants

Participants were recruited from the Pittsburgh, Pennsylvania community with flyers and online postings. The participant population was primarily college students and participants were screened for English fluency and required to be less than 49 years old (average age 24 years). Sixty participants completed the experiment, 20 for each of the main conditions. There were 35 males and 25 females. Participants were assigned randomly to condition. Attrition affected all conditions with five participants failing to complete the optimization method condition, two participants failing to complete the flashcard condition, and five participants failing to complete the Atkinson condition. Participants with any prior knowledge of Japanese were excluded from participation.

Modeling Issues

The experimental design described above required the specification of both an ACT-R model and an Atkinson (1972b) Markov model.

ACT-R Declarative Memory Model

Because the ACT-R declarative memory model captures recency and frequency effects (Anderson & Lebiere, 1998) and produces predictions for both probability and speed of recall, it served as a starting model. Anderson and Schooler (1991) originally developed these equations by showing that memory strength for an item matches what would be optimal in the environment given the frequency and recency of usage of an item. A recent extension of the ACT-R memory model (Pavlik & Anderson, 2005) captures the spacing effect. Importantly, this extension also captures the spacing-by-practice interaction (that more practice leads to more effect of spacing) and the spacing-by-retention interval interaction (that longer retention intervals result in larger spacing effects), effects shown by Underwood (1970) and Bahrick (1979).

Further, the model has been extended to capture the fact that memory items are not equally difficult for each participant to learn (Pavlik, 2007). This extension allows the model to capture three kinds of item and individual differences. *Item differences* reflect that some items are more or less difficult than other items across participants. *Participant differences* reflect that some participants have greater overall facility in the task. Finally, *participant/item differences* reflect that there is consistent variation (given multiple tests) in how difficult a particular item is for a particular partici-

pant. These differences are captured by the β_i parameter (item differences), β_s parameter (participant differences), and β_{si} parameter (participant/item differences). For the full ACT-R model and specifics about how the optimality conditions were computed, see the Appendix.

Model-based optimal practice allocation policy. The model is applied to practice allocation by using it to find items that are at the point where a new recall-or-restudy trial will provide a maximal increase in future expected memory strength. Although the model is used to make these locally optimal decisions, the implications of these decisions must be assessed to insure that they do not in any obvious way result in less than optimal practice on a global level. The best example of a negative effect of such a local policy is the consequence of using probability correct as the measure to be optimized. Probability correct does not work well because of its sigmoid shape as a function of the number of practices (in which several recall-or-restudy practices occur before a first success after which many successes follow). This sort of shape is typical of what one sees in learning when recall-or-restudy repetitions are equally spaced. Because of this rapid shift from 0 to 1 recall probability after an arbitrary number of initial practices, the probability correct function fails to assign much value to the initial practices needed to get to the transition point (because they result in very little increase in probability correct), and ignores the value of overlearning when an pair is already being responded to correctly (again, because the effect on probability correct is minimal).

Because of this problem with using probability, one needs to consider other measures to optimize. Latency of recall (a more continuous, less categorical measure than percent correct) might be a useful measure to consider when trying to describe overlearning, but because failure latencies do not correlate with learning (a result replicated in the following experiment), this option will not help determine a utility function for learning before the first correct recall.

Therefore, the utility measure to be maximized in the following experiment was the model's estimate of the activation gain at final recall (see Appendix). The activation value for an item is a continuous real valued measure of the strength of the item in memory. Unlike probability correct, this measure values equally increases in activation that get one to the zone where probability of correct grows rapidly, increases within that zone, and overlearning that promotes long-term retention. The key is that activation gain for each trial is dependent on recency and spacing, while not being influenced significantly by the order in which the trials occur like probability correct. In other words, activation captures the retained learning from each trial whereas ignoring where the trial is in a sequence of repetitions. (Although the Atkinson model does not have an activation measure, we could transform the Atkinson model's probability correct estimates to create a continuous utility measure to optimize. So again, the learning efficiency method is not necessarily bound to the ACT-R model underlying it in this paper.)

Using activation as the utility measure for practice scheduling reveals interesting implications for the benefit of spaced practice. On the one hand, activation gain at final recall per trial increases monotonically as the spacing of each repetition increases (given an appropriately long retention interval), because wider spacing leads to less forgetting (Equation 3, Appendix). On the other hand, there are also costs associated with widely spaced (less recent) practice.

First, the time on task to successfully retrieve a pair for a recall-or-restudy trial (Equation 5, Appendix) increases monotonically with wider spacing. Second, wider spacing causes a decrease in recall probability (Equation 4, Appendix) during learning, which will necessitate additional time spent on restudy trials.

Knowing our utility function (activation gain) and cost function (time spent practicing including failure costs), we can use these values to compute when the learning rate for each item is maximal. As suggested by Smallwood (1962), the learning rate is the learning gain for the next trial/time cost for that next trial. The Appendix derives this learning rate function from the ACT-R model, which is equivalent to Equation 1. Maximizing this learning rate for each practice of each item is more useful than maximizing gain per trial (the numerator of Equation 1), because one is typically interested in how much learning occurs for a given amount of time rather than for just a given trial. For this reason, our model-based policy will make selections based on this learning rate statistic.

$$ learning\ rate_n = \frac{activation\ gain\ at\ retention\ test\ for\ item_n}{time\ cost\ now\ to\ practice\ item_n} $$

(1)

Finally, although it is true that one could just use this equation to select the pair for the next trial that results in the absolute largest gain in long-term activation per second, this policy would also be imperfect because it would sometimes select pairs because they resulted in better learning than any other pair, despite the fact that waiting longer might result in an even greater learning rate. This is another instance of a negative effect of a local policy, because selecting the pair with the maximum learning rate turns out to be not optimal when considered outside of the temporally isolated situation of the single decision. To resolve this issue, after an initial presentation, the algorithm "waits" for a pair to be spaced to the point where its own gain per unit of time is optimal. Essentially, pairs are practiced when the change in their learning rate (gain/time) is equal to 0. (Pairs are practiced when the second derivative of long-term learning with respect to time spent practicing equals 0.) The crux of this is that pairs are not being compared with each other to select the optimal pair; rather each pair is selected as close to its own optimal time for repetition as possible. This means that selection is much closer to the global optimum policy implied by the model (in which every pair is practiced at its exactly optimal spacing given its history of practice).

Determining optimality conditions. Using versions of Equation 1 (see Equation 7 and Equation 8 in the Appendix), it was possible to describe the practice efficiency functions for recall-or-restudy trials and study-only trials as a function of expected activation for the 9-day average retention interval in the experiment. See Figure 1, which shows our predictions that an expected activation of -0.33 was optimal for recall-or-restudy trials and that when expected activation dropped to -0.63 it became more advantageous to give a study-only trial. Figure 2 shows this selection algorithm graphically.

1. If the pair is above the study advantage point (-0.63 activation) but below the optimality point (-0.33 activation), then the pair should be drilled (i.e., given a recall-or-restudy trial) because waiting longer will result in less efficient practice. If multiple pairs are in the range, the weakest is chosen.
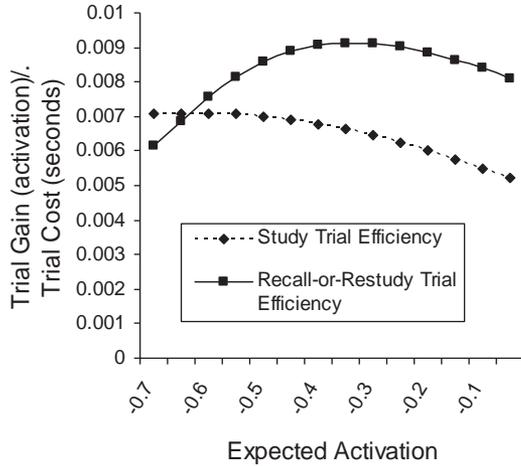
Figure 1. Efficiency functions for recall-or-restudy trials and study-only trials as a function of current activation for a retention interval of 9 days (the mean expected retention interval in the current experiment).

2. If there are no pairs between the points then one can:

   a. Present an old pair that dropped below −0.63 activation for a study-only trial.

   b. Present an unpracticed pair for initial study-only trial. (Unpracticed pairs are only introduced if no practiced pairs are below −0.33 activation.)

3. If the pair is above the optimality point (−0.33 activation), then it is better to wait for it to be forgotten more

before practicing it. If all pairs are above the optimality point then the lowest activation pair is selected.

Atkinson Model

To instantiate the Atkinson (1972b) control condition, the data from Pavlik (2007) were used to find parameters. In the Markov model that Atkinson used for performance estimation there are five parameters: x, y, z, g, and f. A pair starts out with probability g of being in state P (Permanently learned) and $1-g$ of being in state U (Unlearned) to represent prior knowledge participants bring into the experiment. Besides states P and U, pairs may also be in state T where they are considered to be temporarily learned for the duration of the current learning session. Each time a pair i is practiced the pair is transformed according to matrix A_i (see Table 1). Further, whenever *another* pair in the learning set is practiced matrix F_i applies to represent forgetting. These matrices indicate that when a pair is in the state listed on the left of the matrix it has the listed probabilities of transitioning to the states listed along the top of the matrix.

This Markov model needed to be used at the level of individual practices rather than for aggregate data. To make this aggregate model into a trial-by-trial model some things about the state of a pair as a function of prior practice outcomes were incorporated. With these additional considerations, the full model became:

1. If the pair is incorrectly recalled during a test then the pair was in the state $p(U) = 1$ before the trial.

2. If the pair is correctly recalled then the pair probabilities must have been as follows: $p(U) = 0$, $p(P) = p(P)/$
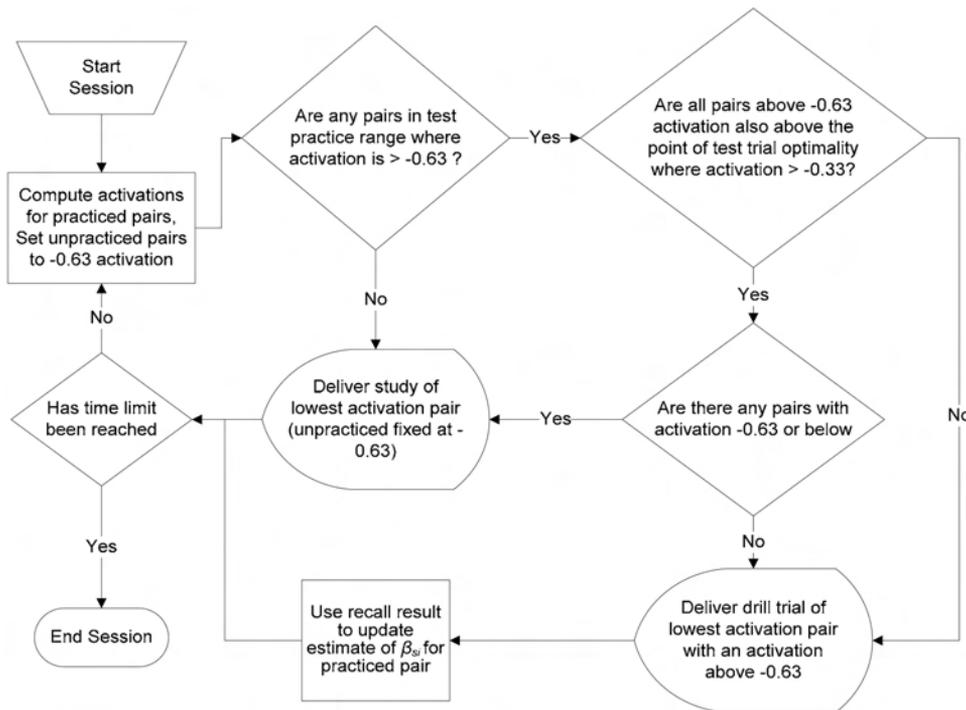

Figure 2. Schedule optimization algorithm flowchart.

Table 1
Atkinson (1972b) Markov Model

	Initial state		
Final state	Permanent	Temporary	Unlearned
	Learning matrix (A_i)		
Permanent	1	0	0
Temporary	x_i	$1-x_i$	0
Unlearned	y_i	z_i	$1-y_i-z_i$
	Forgetting matrix (F_i)		
Permanent	1	0	0
Temporary	0	$1-f_i$	f_i
Unlearned	0	0	1

$(p(P)+p(T))$, and $p(T) = p(T)/ (p(P)+p(T))$ because $p(P)+p(T)$ must sum to 1 if the item was recalled.

3. After any trial of a pair, learning from the feedback study or the correct retrieval lead to A_i being applied for that pair

4. the $F_i$matrix is applied to all other pairs when any pair is practiced.

5. When a pair survives a between-session retention interval it has been permanently learned, $p(P) = 1$, $p(T) = 0$, and $p(U) = 0$. (This model component also means a pair will no longer receive practice if the first test following a long-term retention interval succeeds.)

Table 2 shows the overall parameters determined for this model from the (Pavlik, 2007) data. To decide which pair to practice, Atkinson (1972b) used the Markov model to compute the pair that would have the greatest chance to shift into the permanent state if practiced next. This is also how the algorithm worked here for the control condition in this experiment; however, rather than introduce pairs with test-and-study trials as Atkinson did, the first trial for each word pair in the Atkinson control condition was a study presentation and subsequent presentations were recall-or-restudy trials. Further, because the schedule optimization in the 1972 paper used a minimum spacing of six intervening trials, the version here was likewise constrained to a minimum spacing of six between repetitions. However, pairs were not segregated into groups for rotation like Atkinson, rather, pairs were simply prevented from being selected if they had occurred in the last six trials and the next best pair was selected. (A bug affected 1.2% of trials, resulting in spacings shorter than six intervening trials.)

Results and Discussion

Figure 3 shows average performance for each session for each of three dependent measures. The left panels show Sessions 1 through 3 and the right panels provide Session 4 results with 95% confidence intervals. The main statistical question of interest was whether the schedule optimization condition had an effect on the dependent measures of correctness and latency. The graphs show that these effects were significant for Session 4 recall and latency, $F(2, 57) = 5.4$, $p = .0073$ and, $F(2, 57) = 10.3$, $p = .00015$, with all pairwise t test comparisons favoring the schedule optimization,

($ps < 0.05$). For correctness, these results show a Cohen's-d effect size of 0.796 SD compared to the Atkinson control and 0.978 SD compared to the flashcard control. For latency, these results show an effect size of 1.17 SD compared to the Atkinson control and 1.31 SD compared to the flashcard control. Failure latency was not significantly affected by condition. Before running the experiment, we simulated the predicted probabilities of recall according to the ACT-R model for the first trial of the assessment session for the three conditions and showed a similar result. The simulated values were 56% for the optimization condition, 31% for the flashcard condition, and 24% for the Atkinson condition. These are below the actual observed values 66%, 35%, and 40%, respectively. This suggests that, whereas the model was generally biased to underestimate performance, it seems to well capture the advantage of the optimized condition relative to the controls.

Recall that the design included different nested subcondition sets of pairs within each between-subjects condition. For the learning rate optimization method condition, the utility of using prior item β values was tested to see if they provided an enhancement to the schedule optimization. To do this, before the experiment began for each participant, half of the pairs with prior item β estimates were randomly designated to not use those estimates. This allowed a comparison of performance of the schedule optimization in conditions when the pairs had item βs and used them and when the pairs had item βs but did not use them. Although Session 4 results showed no effect of whether item βs were used, there was a significant advantage for item βs during learning ($F(1, 19) = 5.4$, $p = .032$, $d = 0.377$). This suggests that the item βs did have some benefit because they caused a significant reduction in errors during learning. However, because no effect carried through to Session 4, we probably should conclude that prior item βs are not very important to the overall goal of improving learning. Of course, this result depends on the pairs we used as stimuli, since if there were less homogeneity of difficulty among the pairs, item βs would have been larger and had a stronger effect on practice selection.
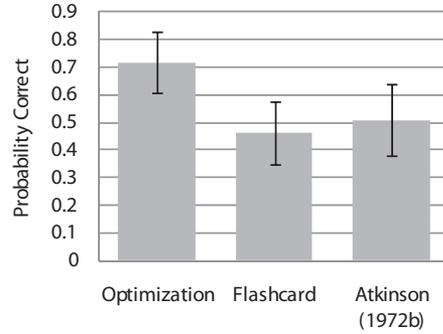
The results for the nested analysis performed in the flashcard control condition showed no serious differences for the new words in the set. In this case, retention of pairs that were new to the learning set (108 pairs) and the old learning set (100 pairs) was compared. There was a 1.3% difference in Session 1 thru 4 performance, which was statistically significant ($F(1, 19) = 9.3$, $p = .0065$, $d = 0.0745$). This difference reflected the fact that the 108 new word pairs were very slightly easier than the prior set of 100 word pairs. Although this difference was significant, there were no interactions and the effect size was very small, so it is implausible to suppose that the new words could have confounded the results of this experiment.

Table 2
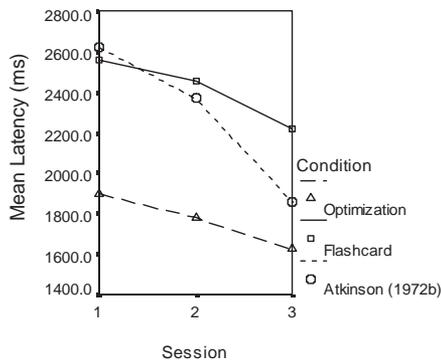*Overall Atkinson (1972b) Parameters for Pavlik (2007) Data*

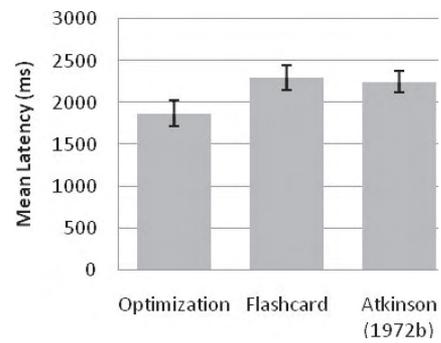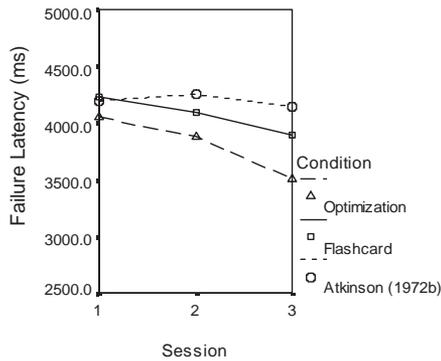| Parameter	Value
x	0.18
y	0.21
z	0.59
g	0
f	0.011
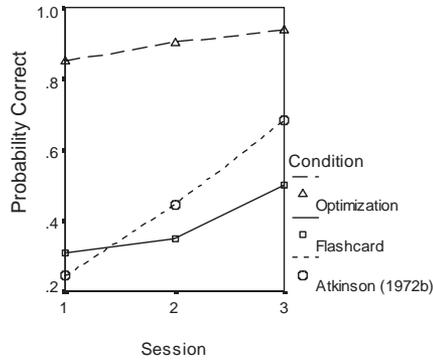
a) Session 1-3 : recall probability
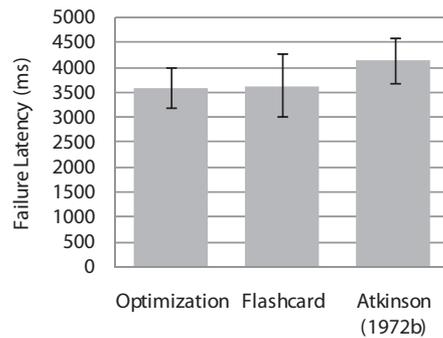
b) Session 4: recall probability

c) Session 1-3 : recall latency

d) Session 4: recall latency

e) Session 1-3 : failure latency

f) Session 4: failure latency

Figure 3. Main results for learning and performance.

Figure 4 shows the results of the nested analysis performed in the Atkinson (1972b) control condition. For this analysis, performance on five overall parameter pairs was compared with those pairs for which five individual parameters had been found. The graph shows a somewhat different pattern than expected. Based on Atkinson, one might have expected worse initial performance and better final performance for the individual parameter words. However, this did not occur because of the way the Atkinson routine

chose to introduce the overall parameter pairs as a cohort. This is illustrated by examining what happens to a pair responded to incorrectly in either the overall or the individual conditions. In the overall condition all pairs come up for practice at the same time, if they are responded to incorrectly, they must be in the U state and then have practice matrix A_i applied. This leaves the probabilities equal again, and the pairs will therefore come up for practice again at the same time (after the whole cohort of overall parameter failed
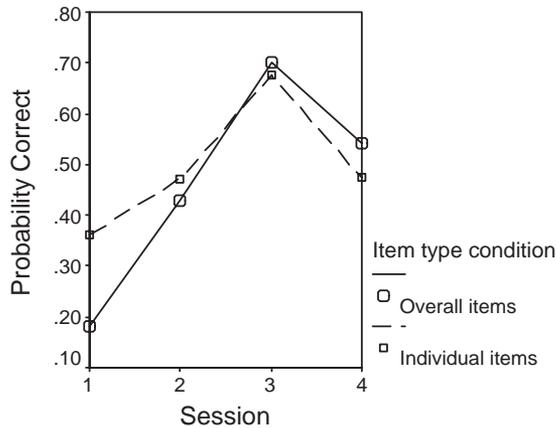
Figure 4. Five parameter overall compared to five parameter per pair subcondition results for the Atkinson (1972b) model condition.

pairs has cycled). In contrast, when an individual parameter pair has the highest transition probability to state P, and is then incorrectly responded to on the scheduled test, it will come up again at about the minimum spacing.

The results of the between-subjects overall difficulty subcondition were analyzed to determine what influence the three conditions might have had on pairs not being scheduled by the algorithms. The mean recall levels for all tests on these pairs was 0.567, 0.435, and 0.489 for the schedule optimization, flashcard, and Atkinson conditions, respectively. A first repeated measures ANOVA looked for significant difference because of condition. Although this overall ANOVA confirmed that differences exist, $F(2, 57) = 4.3$, $p = .019$, follow-up t tests showed that only the difference between the optimization and the flashcard conditions was significant ($p = .0052$, $d = 0.922$). However, a trend was observed for the optimization-Atkinson comparison ($p = .094$). A second ANOVA looked only at first session tests where recall probability during learning was maximally different, since the

greatest differences might be expected during this session. Indeed, the comparison was significant, $F(2, 57) = 5.3$, $p = .0075$ with pairwise t test comparisons significant for the flashcard condition comparison ($p = .016$, $d = 0.794$) and the Atkinson condition comparison ($p = .0032$, $d = 0.978$).

Modeling Questions

It is also possible to use the data from the human participants to ask questions about the models used to optimize performance.

Predictive validity of the ACT-R model during learning. One interesting question is how well the ACT-R model can predict recall given data about learning. To answer this question about prediction, it was possible to use the model with the prior overall parameters from Table 3 and the individual difference parameters (βs) determined during the three learning sessions from the Bayesian procedure described at the end of the Appendix. For each optimization condition participant, Session 4 Trial 1 probability of recall prediction was compared to the actual value the participant produced to determine the deviation of the model from the participant. Predictions for the 20 participants had an r^2 value of 0.86 with a mean absolute deviation of 0.9% (*SD* 9.6%).

Testing the Atkinson model long-term memory assumption. As was discussed earlier, one assumption of the Atkinson (1972b) model was that pairs in the permanently learned state cannot be forgotten. Because pairs are further assumed to lose all temporary strength during a long-term interval, an initial correct recall after a long-term interval necessarily implies a pair is in the permanent state with a probability of 1. Thus, pairs that are successfully recalled on the first opportunity after the long-term interval are no longer selected for practice by the algorithm.

This assumption of permanent memory seemed implausible, but it was possible to look at the data in the Atkinson condition to see how often these abandoned pairs were recalled during the assessment session. To do this the average assessment session recall was computed for the first trial of the assessment session for any pairs that were abandoned for further practice after being responded to

Table 3
Parameter Values and Usage

Parameter type	Symbol	Value	Used in simulation	Used to compute criteria	Used during optimization
Activation	a	0.177	Yes	Yes	Yes
	c	0.279	Yes	Yes	Yes
	h	0.0172	Yes	Yes	Yes
Recall	τ	-0.704	Yes	Yes	No
	s	0.0786	Yes	Yes	No
Latency	F	1.29	Yes	Yes	No
	s_2	0.75	Yes	Yes	No
Study	u	1.205	No	Yes	No
	v	0.000598	No	Yes	No
Fixed costs	*Failure cost*	8.2s	Yes	Yes	No
	Success cost	2.4s	Yes	Yes	No
	Study cost	4.5s	No	Yes	No
β participant	β_s	*SD*=0.283	Yes	No	No
β item	β_i	varied	Yes	No	Yes
β participant/item	β_{si}	*SD*=0.50	Yes	No	Yes
Difficulty	*Slope*	1.45	Yes	Yes	Yes
	Intercept	0.961	Yes	Yes	Yes

correctly for their first recall-or-restudy trial on either Session 2 or 3. The average probability of first trial assessment session recall for these pairs was 0.57 ($SD = 0.26$). This clearly indicates that the assumption of a "permanent" memory state fits the data poorly.

Again, though this reveals a weakness in the Atkinson model's fit, the model might be corrected to overcome this problem with its accuracy by simple increasing the number of states the memory could be in beyond just Unlearned, Temporary, and Permanent. For example, by splitting the Permanent state into two states (Long-term and Permanent) the model could be made to capture forgetting between sessions in an analogous manner as it currently captures forgetting between trials (by using a forgetting parameter that allows for only some items to become unlearned when there is another item tested). Of course, this revised model would have more complexity and parameters, but the feasibility of such four stage models has been shown in prior research (Atkinson & Crothers, 1964).

General Discussion

Ever since Ebbinghaus' (1913/1885) discovery of the spacing effect, researchers have considered how to best make use of this principle to improve how people learned. Although it is clear that some spacing of practice is almost always good, despite more than a century of research the jury appears to still be out on how much spacing is best (e.g., contrast this paper with Pashler et al., 2003). To resolve this longstanding theoretical debate, this paper employed an economic analysis instantiated by a computational model to compute the optimal schedule of practice for each item.

Theoretical Implications

Our results are theoretically important because they strongly qualify the proposition that wider spacing is typically the best practice (Atkinson, 1972b; Bahrick, 1979; Karpicke & Roediger, 2007; Pashler et al., 2003; Schmidt & Bjork, 1992). Our experiment and theoretical analysis shows that by sacrificing wide spacing benefits we can take advantage of recency effects and save time and thus provide significantly more practice. Figure 5 shows the significant effect of condition on the number of trials for the overall comparison ($F(2, 57 = 12.5, p = .000032$) with paired comparisons with the optimization method condition significant for both the flashcard ($p = .0000062, d = 2.21$) and Atkinson ($p = .0060, d = 0.807$) control conditions.

Our results are different and in conflict with current theory because our theory of spacing effect efficiency acknowledges that time costs per trial for spaced practice can easily grow faster than benefits per trial as spacing increases past a certain point. The theory captures this by saying that wide spacing leads to less forgetting, but also results in longer trial latencies. Given a specific retention interval, the modeling method that implements the theory allows one to compute the optimal spacing that maximizes this reduction in forgetting relative to the longer trial durations. The method suggests that initial spacing should be very short because the model predicted 99.2% recall during learning as optimal for the task here (in the minimum noise case). Although this prediction of high level of recall during learning is in general agreement with Skinner (1968) and his theory that errorless learning is optimal, unlike Skinner, the current research provided a mechanistic expla-
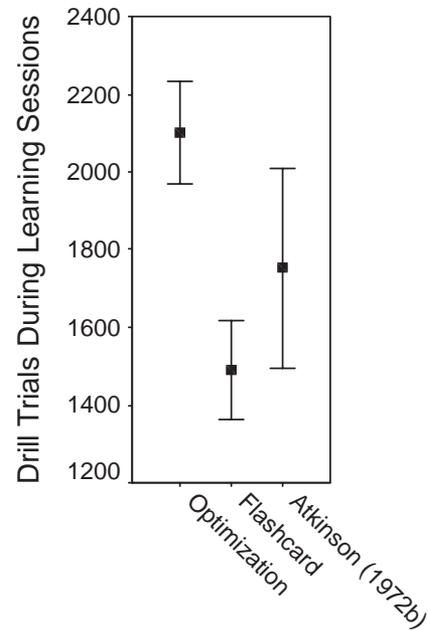

Figure 5. Total recall-or-restudy trials during learning by condition (95% CI error bars).

nation for why this should be so by appealing to cognitive constructs (e.g., the strength of a declarative memory), something that Skinner was unwilling to do. Further, this research provided a well-controlled experimental test of the notion that error minimization results in better learning.

Figures 6 and 7 show one way to visualize the data collected during learning. Because the vast majority of trials for each user were recall-or-restudy practice, the graphs plot correctness and spacing for recall-or-restudy trials only. The y-axes capture the correctness (averaged across items) for each of the first five recall-or-restudy trials, whereas the x-axes captures the spacing *following* each of these trials. To read the graphs, note that for each panel the first trial of each series is labeled with the count of observations, as is the last trial. This gives a perspective on the quantity of practice for each item in the various conditions. In addition, note that the left sides of Figures 6 and 7 show the overall average, whereas the right sides show performance conditional on the results of the first recall-or-restudy. These condition graphs show that there is a tendency in all conditions to schedule more repetitions when the first repetition is responded to incorrectly. We can see on the right side of Figure 6 that the optimization method reacts to correctness by increasing the spacing of the schedule and incorrectness by decreasing the spacing of the schedule and that average correctness is maintained at a relatively high and constant level in both cases. In contrast, the control conditions show a different pattern that gives much wider spacing after a correct answer, but which subsequently contracts to become narrower because second trial failures are very high.

The smoothly expanding spacing of practice produced by the method is in general agreement with the theory that expanding spacing of practice may be optimal (Cull et al., 1996; Landauer & Bjork, 1978; Rea & Modigliani, 1985). However, although these
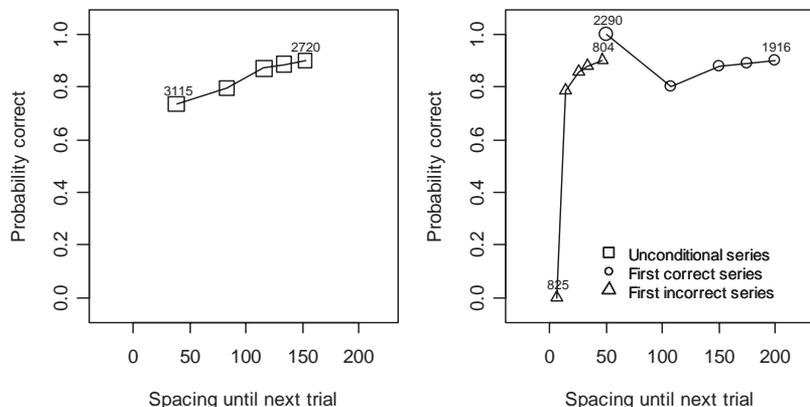
Figure 6. Graph of average recall-or-restudy recall probability (*y*-axis) and intervening trials of spacing before the subsequent repetition recall-or-restudy (*x*-axis) for the first 5 recall-or-restudy trials in the optimization method condition. The graph on the left is the overall average and the graph on the right is conditional on the results of the first recall-or-restudy.

studies only support expanding spacing for a test-only procedure (no review following errors), by attending to the costs of wider spacing our method has demonstrated that expanding spacing may also be optimal for recall-or-restudy trials. The method produces this expanding spacing because of the effect of frequency in the ACT-R model. As frequency increases, an item becomes more stable in memory because the model implements power function forgetting which results in strength from older accumulated practices decaying increasingly slowly as time passes. This increased stability with increased frequency allows more time between spaced practices as repetitions accumulate. Thus, our modeling approach has allowed us to derive that expanding spacing is the optimal solution to the scheduling problem by quantifying the theoretical relationships between recency, frequency, and spacing and their effects on final performance.

Practical Implications

Although the method clearly has implications for the learning of large sets of paired-associate items by young naïve participants, it is less obvious what this implies for different tasks, different populations of learners, or different materials. Take for example a different task like recognition. The interesting difference about recognition tasks is that there is no need for review in the case of an error because the presentation of the test probe itself can serve as the learning necessary for later judgments of recognition. Because there is therefore no need for review study opportunities, the cost of errors is much less. Therefore, because correctness would no longer be so important, recency would no longer provide an advantage and the peak learning rate spacing interval would be much longer than in our experiment here. In fact, unless there were other costs to consider, the method would probably suggest maximal even spacing of practice as optimal unless the set of items were extremely large.

It is similarly straightforward to speculate about other learning situations that may be amenable to the method we have described. Because procedural tasks are typically represented as if-then rules that might be collected into a set of items and trained using drill practice, they also seem applicable to the method here assuming an

accurate model could be found. Prior work on procedural learning (e.g., Carlson & Yaure, 1990; Mayfield & Chase, 2002) has compared blocked practice (a form of massed practice) versus mixed practice (which requires wider spacing of practice) and found that mixed presentation schedules caused an advantage to retention. These results are important because they provide evidence that rule-based items may behave according to a sort of spacing effect (any advantage to retention when using more distributed practice). Since the optimization method harnesses spacing effects in deciding practice schedules, knowing that the spacing effect applies to practice of rule application skills is important. Further, we can note that procedural tasks often have high failure costs. Together, these facts imply that procedural tasks may have an optimal expanding schedule similar to the one produced in the experiment here. However, because parameter values such as the forgetting rate (which controls the effect of recency) and the amount of learning per repetition might differ greatly, the magnitude of the spacing intervals might differ greatly.

The age of participants would also be likely to alter the optimal schedule. For instance, research by Balota et al. (1989, 2006) shows that although there tends to be no interaction of age and spacing effect, older individuals have generally worse recall. This recall deficit might be modeled as an increase in forgetting or a reduction in learning. Changes in the model parameters controlling these effects would tend to result in less wide practice schedules because narrower schedules would produce higher levels of recall during learning (and thus avoid costly review opportunities). However, the issue is complex because the schedule changes would depend upon which parameter(s) was used to capture the age related differences. For example, if deficient performance is caused by a higher level of memory strength variability, captured by the noise (*s*) parameter in the ACT-R model, the model's precision of recall probability prediction would decrease and the benefit of spacing would dominate the learning rate equation. This would imply wider spacing, or, in the case of very high noise, the model would predict maximal spacing at all times. This possibility is one reason why we tried to reduce overall variability during
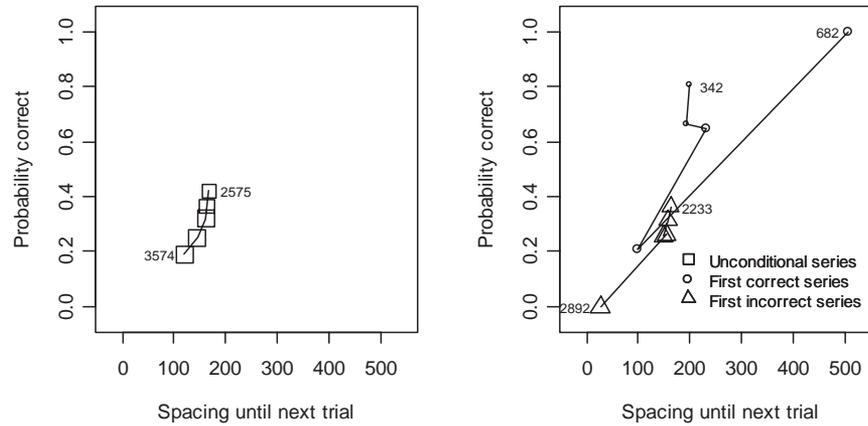
a) Flashcard condition schedule
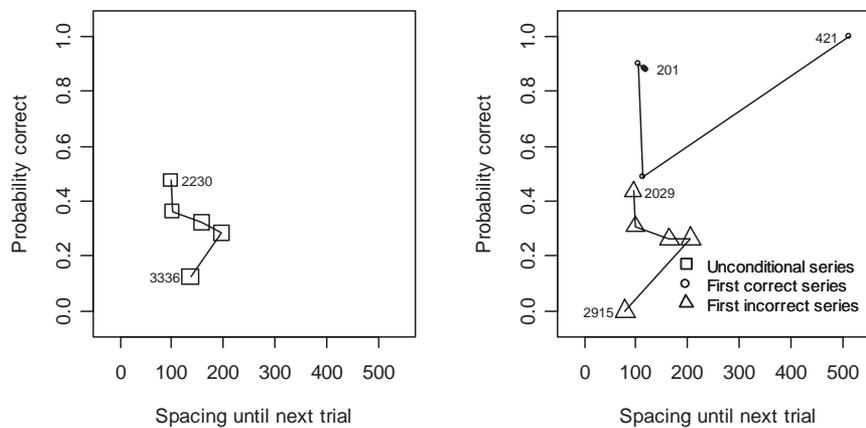
b) Atkinson condition schedule

Figure 7. Graph of average recall-or-restudy recall probability (*y*-axis) and intervening trials of spacing before the subsequent repetition recall-or-restudy (*x*-axis) for the first 5 recall-or-restudy trials in the Flashcard and Atkinson (1972b) control conditions averaged by item. The graph on the left is the overall average and the graph on the right is conditional on the results of the first recall-or-restudy.

stimuli design by choosing a relatively homogenous set of items and avoiding a language with many English cognates.

These examples, while showing limitations, suggest that the method we have developed should be useful in developing instructional programs where an accurate model of a task can be described. However, creating a model is clearly not trivial even for relatively simple domains such as the fact memorization we investigated. For more complex domains with multiple grain sizes (some items contain more information than others), dependencies (some items depend on other items) and transfer effects (learning some items transfers to other items) the modeling must take on additional complexity to produce an accurate model and enable the method. At the current time, it is unclear how issues such as these can be modeled, but it is clear that if they are not included in a model when they are strong effects then the method cannot be expected to work properly. Current work is exploring how issues such as grainsize, dependency, and transfer might be modeled and thus enable the method (Pavlik, Presson, & Koedinger, 2007).

Implementation

So long as one realizes the implications and limitations listed above, implementing the algorithm to train a collection of items is a straightforward process. Implementation involves the following steps:

1. Select a model of the effect of the history of practice. This model must characterize:

   a. Expected probability correct

   b. Expected latency of response

   c. Expected latency of failure

   d. Expected future learning utility measure (must be a continuous measure that values practices independent of frequency)

2. Estimate the model's parameters (using prior data)

3. Use the model before each drill trial to compute for every item:

   a. Expected learning rate = (expected future learning utility gain for a practice/expected time cost for a practice)

   b. Choose each item for practice at its maximum expected learning rate

   c. Repeat until stopping criterion for time practiced or expected performance

References

Anderson, J. R., Fincham, J. M., & Douglass, S. (1997). The role of examples and rules in the acquisition of a cognitive skill. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 23,* 932–945.

Anderson, J. R., & Lebiere, C. (1998). *The atomic components of thought.* Mahwah, NJ: Erlbaum Publishers.

Anderson, J. R., & Schooler, L. J. (1991). Reflections of the environment in memory. *Psychological Science, 2,* 396–408.

Atkinson, R. C. (1972a). Ingredients for a theory of instruction. *American Psychologist, 27,* 921–931.

Atkinson, R. C. (1972b). Optimizing the learning of a second-language vocabulary. *Journal of Experimental Psychology, 96,* 124–129.

Atkinson, R. C. (1975). Mnemotechnics in second-language learning. *American Psychologist, 30,* 821–828.

Atkinson, R. C., & Crothers, E. J. (1964). A comparison of paired-associate learning models having different acquisition and retention axioms. *Journal of Mathematical Psychology, 1,* 285–315.

Atkinson, R. C., Fletcher, D., Lindsay, J., Campbell, J. O., & Barr, A. (1973). Computer assisted instruction in initial reading: Individualized instruction based on optimization procedures. *Educational Technology, 8,* 27–37.

Atkinson, R. C., & Paulson, J. A. (1972). An approach to the psychology of instruction. *Psychological Bulletin, 78,* 49–61.

Bahrick, H. P. (1979). Maintenance of knowledge: Questions about memory we forgot to ask. *Journal of Experimental Psychology: General, 108,* 296–308.

Balota, D. A., Duchek, J. M., & Paullin, R. (1989). Age-related differences in the impact of spacing, lag, and retention interval. *Psychology and Aging, 4,* 3–9.

Balota, D. A., Duchek, J. M., Sergent-Marshall, S. D., & Roediger, H. L., III (2006). Does expanded retrieval produce benefits over equal-interval spacing? Explorations of spacing effects in healthy aging and early stage Alzheimer's disease. *Psychology and Aging, 21,* 19–31.

Carlson, R. A., & Yaure, R. G. (1990). Practice schedules and the use of component skills in problem solving. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 16,* 484–496.

Carrier, M., & Pashler, H. (1992). The influence of retrieval on retention. *Memory & Cognition, 20,* 633–642.

Coltheart, M. (1981). The MRC psycholinguistic database. *Quarterly Journal of Experimental Psychology: Human Experimental Psychology, 33,* 497–505.

Cull, W. L. (2000). Untangling the benefits of multiple study opportunities and repeated testing for cued recall. *Applied Cognitive Psychology, 14,* 215–235.

Cull, W. L., Shaughnessy, J. J., & Zechmeister, E. B. (1996). Expanding understanding of the expanding-pattern-of-retrieval mnemonic: Toward confidence in applicability. *Journal of Experimental Psychology: Applied, 2,* 365–378.

Dear, R. E., Silberman, H. F., & Estavan, D. P. (1967). An optimal strategy for the presentation of paired-associate items. *Behavioral Science, 12,* 1–13.

Dempster, F. N. (1996). Distributing and managing the conditions of encoding and practice. In R. A. Bjork & E. L. Bjork (Eds.), *Memory* (pp. 317–344). New York: Academic Press.

Ebbinghaus, H. (1913). *Memory: A contribution to experimental psychology (translated by Henry A. Ruger & Clara E. Bussenius; original German work published* 1885). New York: Teachers College, Columbia University.

Fishman, E. J., Keller, L., & Atkinson, R. C. (1969). Massed versus distributed practice in computerized spelling drills. In R. Atkinson & H. A. Wilson (Eds.), *Computer assisted instruction.* New York: Academic Press.

Glenberg, A. M. (1976). Monotonic and nonmonotonic lag effects in paired-associate and recognition memory paradigms. *Journal of Verbal Learning & Verbal Behavior, 15,* 1–16.

Groen, G. J., & Atkinson, R. C. (1966). Models for optimizing the learning process. *Psychological Bulletin, 66,* 309–320.

Karpicke, J. D., & Roediger, H. L., III. (2007). Expanding retrieval practice promotes short-term retention, but equally spaced retrieval enhances long-term retention. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 33,* 704–719.

Karush, W., & Dear, R. E. (1966). Optimal stimulus presentation strategy for a stimulus sampling model of learning. *Journal of Mathematical Psychology, 3,* 19–47.

Landauer, T. K., & Bjork, R. A. (1978). Optimum rehearsal patterns and name learning. In M. M. Gruneberg, P. E. Morris, & R. N. Sykes (Eds.), *Practical aspects of memory* (pp. 625–632). New York: Academic Press.

Laubsch, J. H. (1971). An adaptive teaching system for optimal item allocation. *Dissertation Abstracts International, 31,* 3961.

Lorton, P. V. (1973). Computer-based instruction in spelling: An investigation of optimal strategies for presenting instructional material. *Dissertation Abstracts International, 34,* 3147.

Mayfield, K. H., & Chase, P. N. (2002). The effects of cumulative practice on mathematics problem solving. *Journal of Applied Behavior Analysis, 35,* 105–123.

Metcalfe, J., & Kornell, N. (2003). The dynamics of learning and allocation of study time to a region of proximal learning. *Journal of Experimental Psychology: General, 132,* 530–542.

Murphy, G., & Kovach, J. (1972). *Historical introduction to modern psychology.* New York: Harcourt Brace Jovanovich.

Nelson, T. O., & Leonesio, R. J. (1988). Allocation of self-paced study time and the "labor-in-vain effect." *Journal of Experimental Psychology: Learning, Memory, and Cognition, 14,* 676–686.

Pashler, H., Cepeda, N. J., Wixted, J. T., & Rohrer, D. (2005). When does feedback facilitate learning of words? *Journal of Experimental Psychology: Learning, Memory, and Cognition, 31,* 3–8.

Pashler, H., Zarow, G., & Triplett, B. (2003). Is temporal spacing of tests helpful even when it inflates error rates? *Journal of Experimental Psychology: Learning, Memory, and Cognition, 29,* 1051–1057.

Pavlik, P. I., Jr. (2005). The microeconomics of learning: Optimizing paired-associate memory. *Dissertation Abstracts International: Section B: The Sciences and Engineering, 66,* 5704.

Pavlik, P. I., Jr. (2006). Understanding and applying the dynamics of test practice and study practice [Electronic Version]. *Instructional Science,* from http://dx.doi.org/10.1007/s11251–006–9013–2.

Pavlik, P. I., Jr. (2007). Timing is an order: Modeling order effects in the learning of information. In F. E., Ritter, J. Nerb, E. Lehtinen, & T. O'Shea (Eds.), *In order to learn: How order effects in machine learning illuminate human learning* (pp. 137–150). New York: Oxford University Press.

Pavlik, P. I., Jr., & Anderson, J. R. (2005). Practice and forgetting effects on vocabulary memory: An activation-based model of the spacing effect. *Cognitive Science, 29,* 559–586.

Pavlik, P. I., Jr., Presson, N., & Koedinger, K. R. (2007). Optimizing knowledge component learning using a dynamic structural model of practice. In R. Lewis & T. Polk (Eds.), *Proceedings of the Eighth International Conference of Cognitive Modeling.* Ann Arbor: University of Michigan.

Peterson, L. R., Wampler, R., Kirkpatrick, M., & Saltzman, D. (1963). Effect of spacing presentations on retention of a paired associate over short intervals. *Journal of Experimental Psychology, 66,* 206–209.

Pimsleur, P. (1967). A memory schedule. *The Modern Language Journal, 51,* 73–75.

Rea, C. P., & Modigliani, V. (1985). The effect of expanded versus massed practice on the retention of multiplication facts and spelling lists. *Human Learning: Journal of Practical Research & Applications, 4,* 11–18.

Ruch, T. C. (1928). Factors influencing the relative economy of massed and distributed practice in learning. *Psychological Review, 35,* 19–45.

Schmidt, R. A., & Bjork, R. A. (1992). New conceptualizations of practice: Common principles in three paradigms suggest new concepts for training. *Psychological Science, 3,* 207–217.

Skinner, B. F. (1968). *The technology of teaching.* East Norwalk, CT: Appleton-Century-Crofts.

Smallwood, R. D. (1962). *A decision structure for teaching machines.* Cambridge: MIT Press.

Smallwood, R. D. (1971). The analysis of economic teaching strategies for a simple learning model. *Journal of Mathematical Psychology, 8,* 285–301.

Thios, S. J., & D'Agostino, P. R. (1976). Effects of repetition as a function of study-phase retrieval. *Journal of Verbal Learning & Verbal Behavior, 15,* 529–536.

Underwood, B. J. (1970). A breakdown of the total-time law in free-recall learning. *Journal of Verbal Learning & Verbal Behavior, 9,* 573–580.

Underwood, B. J., Kapelak, S. M., & Malmi, R. A. (1976). The spacing effect: Additions to the theoretical and empirical puzzles. *Memory & Cognition, 4,* 391–400.

Wozniak, P. A., & Gorzelanczyk, E. J. (1994). Optimization of repetition spacing in the practice of learning. *Acta Neurobiologiae Experimentalis, 54,* 59–62.

Appendix: ACT-R Model

Activation Equation (Memory Strength Function). Practice effects and item/individual differences are captured by an activation equation (Equation 2), which represents the strength of an item in memory as the sum of the remaining learning from a number of individual memory strengthenings and values representing the item/individual differences. Each strengthening corresponds to a past practice event (either a memory retrieval or study event). Specifically, Equation 2 proposes that each time an item is practiced, the activation of the item, m_n (activation on the nth practice) receives a boost in strength that decays away as a power function of time. Each t_k in Equation 2 equals the age of trial k, whereas each d_k equals the decay rate for trial k. The β parameters are a new addition (Pavlik, 2007). They capture any deviation of the overall model from the data for any particular sequence of tests with an item for a participant. They represent differences at the level of the item (β_i), participant (β_s), and participant/item (β_{si}).

$$m_n(t_{1..n}) = \beta_s + \beta_i + \beta_{si} + \ln\left(\sum_{k=1}^{n} t_k^{-d_k}\right) \qquad (2)$$

To deal with the spacing effect Equation 3 was developed to estimate decay for the kth trial, d_k, as a function of the activation at the time it occurred. The implication of this function is that higher activation at the time of a practice will result in the benefit of *that* practice decaying more quickly. On the other hand, if activation is low, decay will proceed more slowly. It is important to note that *every practice has its own* d_k that controls the forgetting of that practice. In Equation 3 the decay rate d_k is calculated for the kth presentation of an item as a function of the activation m_{k-1} at the time the presentation occurred (e.g., the decay rate for the 7th trial (t_7) depends on the activation at the time of the seventh trial, which is a function of the prior six trials' ages and decay rates. It is important to note that since t_ks are ages, activation and decay depend on the *current time* as well as the number of practices).

$$d_k(m_{k-1}) = ce^{m_{k-1}} + \alpha \qquad (3)$$

In Equation 3, *c* is the decay scale parameter, and *a* is the intercept of the decay function. For the first practice of any sequence, $d_1 = a$ because m_0 is equal to negative infinity. These equations are recursive because to calculate any particular m_n at the time of trial *n*, one must have previously calculated the d_k needed for each of the prior *n-1* trials, which requires knowing the prior m_ns.

These equations result in a steady decrease in the long-run retention for more presentations in a sequence where presentations are closely spaced. As spacing gets wider in such a sequence, activation has time to decrease between presentations, decay is therefore lower for new presentations and long-run retention effects do not decrease as much.

Individual and Item Differences and β

The use of β to represent constant noise has been implicit in the system since at least Anderson and Lebiere (1998). In that documentation of the declarative memory system two sources of the noise are proposed. These two types of noise are referred to as permanent activation noise and temporary activation noise. Permanent activation noise is exactly equivalent to the sum of the three β components. In contrast, random trial-to-trial noise is captured in the *s* parameter in Equation 4 below, which represents the probability of recall given a random logistic distribution of activation (with *s* controlling the variance).

(*Appendix continues*)

Recall Equation

In ACT-R, an item will be retrieved if its activation is above a threshold. Because activation is noisy, an item with activation m as given by Equation 4 has only a certain probability of recall. ACT-R assumes a logistic distribution of activation noise, in which case the probability of recall is:

$$p(m) = \frac{1}{1 + e^{\frac{\tau - m}{s}}} \qquad (4)$$

In Equation 4, τ is the threshold parameter. As discussed, the s parameter controls the noise in activation and it describes the sensitivity of recall to changes in activation.

Retrieval Time Equation

The time to retrieve in ACT-R is shown in Equation 5. In Equation 5, F is the parameter that scales the effect of latency on activation. Fixed time cost refers to the fixed time cost of perceptual motor encoding and response.

$$l(m) = Fe^{-m} + fixed\ time\ cost \qquad (5)$$

Intersession Forgetting in the Model

Anderson, Fincham, and Douglass (1997) found that although Equation 2 could account for practice and forgetting during an experiment, it could not fit retention data over long intervals. Because of this, they supposed that between sessions, intervening events erode memories more slowly than during an experimental session. This slower forgetting was modeled by scaling time as if it was slower outside the experiment. Forgetting is therefore dependent on the "psychological time" between presentations rather than the real time. This psychological time factor is implemented by multiplying the portion of time that occurs between sessions by the h parameter when calculating recall. This is done by subtracting $h*$total intersession time from each age (t_k) in Equation 1. Because of this mechanism, time in the model is essentially a measure of destructive interfering events. The decay rate, therefore, is a measure of fragility of memories to the corrosive effect of these other events.

Study Duration Model

Because it was shown in Pavlik Jr. (2006) that ACT-Rs assumptions about study practice and test practice having equal memorial consequences did not fit the data, and that a new model of study duration resulted in interesting predictions about efficiency, this new, better fitting model was used for the current study. In this model strength of a study practice, $b = u$ $(1-e^{\ (duration-study\ trial\ fixed\ cost)\ \cdot\ (-v)/number\ of\ items\ in\ the\ stimulus})$, where u is the maximum benefit of study and v describes the rate of approach to the maximum. This simple model captures the fact that study practice has diminishing marginal returns and appears to reach an asymptotic level of encoding given a long enough encoding interval (Metcalfe & Kornell, 2003; Nelson & Leonesio, 1988). The study value is used as a factor for the term representing the respective study learning events in the activa-

tion equation (Equation 6). In the case of test practices b is fixed at 1.

$$m_n(t_{1..n}) = \beta_s + \beta_i + \beta_{si} + In\left(\sum_{k=1}^{n} b_k\, t_k^{-d_k}\right) \qquad (6)$$

Further, this study effect model had an additional component that was found to be useful in capturing the fact that study following a failed retrieval appeared to be more effective than study practice alone. To capture this effect of study after failure, the v parameter was divided by the number of terms in the stimulus. This component of the model says that during study trials participants deploy an attentional resource (typically in a strategic fashion, but also through rote processes) to encode the stimulus being studied. Because this resource is limited, it must be divided among the components of the stimulus (this is done in the model by dividing the encoding rate by the stimulus size.) This limited-resource mechanism implies that the advantage of a study after a failed test comes from the opportunity to pre-encode the cue. Because of this pre-encoding of the prompt during the failed test study opportunity, the encoding of the single response term proceeds twice as quickly compared to a study-only trial.

Parameter Determination

The details of how parameters for these models were determined, using data from Pavlik (2006) and Pavlik (2007), have been omitted, but are available on request from the author. Table 3 shows the final parameters used by the model in the schedule optimization or simulation. The right hand side of Table 3 details that parameters were used to simulate the experiment (simulation was completed before running the experiment), which parameters were used to determine the decision criteria the schedule optimization algorithm would use (again, before the experiment), and which parameters were used to predict memory strengths for the application of the decision criteria (during the experiment).

Determining Optimality Conditions

Given the ACT-R model, it was possible to compute the average activation gain after a 9-day retention interval (the average for the experiment being optimized) for both recall-or-restudy trials and study-only trials as a function of different levels of activation. [One might have considered calculating this value for each day of the experiment separately, using 11, 9, and 7 days for each of the three learning sessions, but the predictions vary little for these relative small differences in retention so the average was used.] Given a set of parameter values, average activation gain for a practice depends on the activation (m) at the time of practice and the duration of the retention interval (r). Similarly, one can compute the estimated time cost for both study-only trials and recall-or-restudy trials, including fixed costs of intertrial times. Again, the cost is a function of activation. These values were used to compute the functions describing the activation gain per second for selecting a recall-or-restudy trial or selecting a study-only trial.

Equation 7 describes the activation gain per second of practice for recall-or-restudy trials.

gain test$_n$(r, m$_n$)

$$= \frac{p\,(m_n)r^{-(ce^{m.}+a)}\ +\ (1\ -\ p(m_n))\cdot 95\cdot r^{-(ce^{m.}+a)}}{p(m_n)\,(F_e^{-m}\ +\ fixed\ success\ cost)\ +\ (1\ -\ p(m_n))\,(fixed\ failure\ cost)} \quad (7)$$

In this equation, p is the probability of recall, r is the retention interval, and the other parameters are defined in Equation 2 and Equation 5. For simplicity, Equation 7 ignores the influence of the natural logarithm in the activation equation (Equation 2). This was found to have an insignificant effect on predictions. Equation 7 uses a 0.95 (b parameter) weight to capture the average effect of the 3-s feedback study after failed tests (see Equation 6).

Equation 8 describes the activation gain per second of practice for study-only trials.

$$gain\ study_n(r, m_n)\ =\ \frac{.79\cdot r^{-(ce^{m.}\ +\ a)}}{fixed\ study\ cost} \quad (8)$$

Equation 8 uses a 0.79 weight to capture the effect of a 4-s study delivered alone (see Equation 6).

Correcting for Overall Difficulty

As the data showed here, there is also an effect of the context of difficulty. Because this effect was not captured in the model, but also appeared to be present in the data from Pavlik (2007) a correction was made to the model to capture this variance. For this correction, the model assumed an additional β correction component that was a linear function of correctness during learning. This was implemented by a linear model describing how overall difficulty during learning mapped to β during learning. This line was the equation β correction = average activation during learning * 1.45 + 0.916. This correction was meant merely to capture variance, and it seems that more data and theoretical work are needed to explain the phenomenon as discussed in the introduction. The correction was applied when determining the values −0.66 and −0.33 from Equations 6 and 7 and was used when computing Figure 1. Further, this correction was applied to computing activations during learning. During learning, the current correction was estimated after each trial for each participant by using the inverse of Equation 3 and the overall average prior probability of recall for that participant to compute the average activation during learning.

Like Pavlik (2007), the algorithm was implemented with a requirement for a spacing of two trials between repetitions because it seemed risky to assume the model would be effective at very short intervals. It seemed possible that some sort of study-phase retrieval spacing effect mechanism would severely block encoding at very short lags (a mechanism discussed by Thios & D'Agostino, 1976). The algorithm had an interesting adaptation to this stricture because the average pair dropped below −0.63 activation after two intervening trials. For this reason, the schedule optimization began by introducing most new pairs with two study-only trials, spaced depending on the activations of other pairs. After the second study average pairs were within the testing range (greater than −0.63 activation) and received recall-or-restudy trials. Interestingly, because the program updated an estimate of the participant β value every 300 trials, the number of introductory study-only trials varied and could go as high as 4 or 5 for very slow learning participants and as low as 1 for able participants. Another interesting consequence of the algorithm was that if a pair was responded to incorrectly several times in a row then its estimated β would fall low enough that its activation would dip below −0.63. At that point, the pair will be scheduled for several study-only trials before being scheduled for further recall-or-restudy practice.

Each recall-or-restudy trial was also used to update the estimate of participant/item β for the practiced word pair. This estimation began by assuming a prior distribution $F(x)$ of βs with a prior participant/item mean β equal to 0. The participant/item β was then updated based on the success and failure for each test given this prior. This Bayesian procedure provided current participant/item β estimates for each pair that reflected the history of success or failure in recalling the pair. Practically this procedure involved a successive shifting and narrowing of the participant/item estimate of β as data were gathered.