Adaptive Learning and the Allocation of Time*

Lisa K. Son[†] Rajiv Sethi[‡]

Forthcoming in: *Adaptive Behavior*

*We thank John Dunlosky and an anonymous referee for helpful comments on an earlier version.

[†]Department of Psychology, Barnard College, Columbia University (lson@barnard.edu).
[‡]Department of Economics, Barnard College, Columbia University (rs328@columbia.edu).

Abstract

Time is a perishable resource that cannot be reallocated once it has elapsed. As a result, time allocation decisions must often be adjusted on the basis of information that becomes available during the process of allocation itself. We show that adaptive rules of this kind will not generally result in optimal allocations when the relationship between the optimal allocation and the time budget exhibits discontinuities, as is the case with logistic learning curves. Under such circumstances, self-guided study is likely to perform poorly relative to appropriately guided learning.

1 Introduction

The allocation of time across competing tasks differs from many other resource allocation problems in one important respect: time is irreversible and cannot be reallocated once it has elapsed. If a decision maker has both the cognitive sophistication and the information necessary to compute optimal allocations, this difference is inconsequential. When information about the underlying parameters of the problem is incomplete, however, the irreversibility of time precludes the use of certain strategies, such as allocation by trial and error. In this case, one must rely on heuristics based on information that becomes available *during the process of allocation itself*. Under what conditions might one expect such heuristics to result in optimal allocations? This is the main question addressed in this paper.

Time allocation strategies have been widely investigated in the field of cognitive psychology (Son & Kornell, 2008; Thiede & Dunlosky, 1999). A key concern in this literature is the effectiveness with which learners allocate time to the study of various items. For instance, several experiments have explored whether or not learner-controlled allocation strategies result in better performance than exogenously specified or computer-controlled strategies. The results have been mixed. Atkinson (1972) provided the earliest experimental evidence to suggest that learners use strategies that are dysfunctional or maladaptive. Learner-controlled allocation decisions were found to result in poorer performance than those controlled by a computer, leading the author to conclude that "the learner is not a particularly effective decision-maker." (Atkinson, 1972, p. 930). Similar results have been found in subsequent studies (Mazzoni, Cornoldi & Marchitelli, 1990; Metcalfe, Kornell, & Son, 2007; Nelson & Leonesio, 1988). More recent research has shown, however, that people are able to make effective allocation decisions, especially when items are presented simultaneously and the total time available is clearly displayed. Metcalfe and Kornell (2003) found that subjects were able to adjust their allocation strategies in response to the total availability of time in a manner that resulted in enhanced performance, and Kornell and Metcalfe (2006) found performance to be better under learner controlled (relative to computer controlled) conditions.

Why might learners choose effective time allocation strategies in some environments but not in others? We argue here that in the absence of complete information about the manner in which investments of time translate into increases in competence, learners are forced to rely on cues based on recent experience. For instance, learners might allocate the most time to items for which prior investments have resulted in the greatest gains. Whether or not such cues guide the learner towards optimal allocations depends on the shapes of the learning curves or uptake functions, which govern the relationship between time invested and competence attained. If all uptake functions are concave, there exist simple rules based on local cues that result in optimal allocations even with very limited information. However, when learning curves are S-shaped (as in the case of logistic uptake functions), local learning can result in suboptimal allocations. The reason is that with this latter class of learning curves, there are points of discontinuity and non-monotonicity in the relationship between optimal allocations and total available time (Son and Sethi, 2006). Since learning rules based on local cues result in terminal allocations that vary continuously in total available time, such rules cannot consistently result in optimal time allocation. We make these ideas more precise below.

2 The Learning Environment

A learning curve is a functional relationship between investments of study time and levels of attained competence. This relationship can be complex and vary substantially from one task to the next. Even for a given task, the rate of uptake can change sharply as greater levels of competence are attained. For instance, in the early stages of learning, small investments of time might result in rapid improvements, while in later stages competence may be relatively insensitive to further investments. Alternatively, learning might be difficult at the outset and become progressively easier over time, before eventually reaching plateau.

Insert Figure 1 here

The literature has focused on two types of learning curve, as shown in Figure 1. The concave curve in the left panel exhibits diminishing returns to time allocated, consistent with exponential, hyperbolic, square-root, and power functions. Such learning curves have been shown to arise in many laboratory environments (see, for instance, Anderson and Schooler, 1991). The S-shaped curve on the right, consistent with the logistic function, has generally been associated with the learning of more complex skills. Evidence for such functions have been found in a variety of experimental environments including sequence learning (Noble, 1957), language acquisition (Rice, Wexler, & Hershberger, 1998), motor learning (Newell et al., 2001), and conditioning (Frey & Sears, 1978; Klopf, 1988). Concave curves have been ubiquitous in empirical research, and many have argued that learning in general is characterized by diminishing returns (see, for instance, Ritter and Schooler, 2001). However, given the very short time frames over which most experiments are conducted, some have argued that laboratory data capture only the upper (concave) section of a learning curve that may well be S-shaped (Fischer & Pipp, 1984; Newell, Liu, & Mayer-Kress, 2001).[1] This is because, for many tasks, a substantial amount of learning has already occurred in daily life, prior to the start of the experimental session.

Individuals are typically confronted with a variety of tasks and a limited amount of time to devote to their study. Let T denote the total time to be allocated to the learning of n tasks, and

[1] Learning curves are often depicted as a downward sloping relationship between the time taken for completion of a task and the number of practice sessions previously completed (see, for instance, Ritter and Schooler, 2001). A concave uptake function (as in the left panel of Figure 1) appears as a negatively accelerated curve in this setting.

let $t_1, ..., t_n$ denote the times spent on each of the tasks, where $T = t_1 + ... + t_n$. Let $b_i = f_i(t_i)$ denote the level of competence attained in task i as a result of allocating time t_i to this task. The function f_i represents the learning curve for task i. Ideally, learners would wish to allocate time across tasks in such a manner as to attain the most preferred distribution of final competencies. For instance, they may simply wish to maximize the average competence across tasks, or if some tasks are considered to be more important than others, they may seek to maximize some weighted average of competencies. More complicated objectives are also plausible.[2]

A learner with complete information about the uptake functions and the total available time could, in principle, choose an allocation of time across tasks that perfectly achieves their goal. In many environments, however, such information is simply unavailable, and properties of the uptake function become apparent only gradually as time is allocated to study. We are particularly interested in the case where, at any given allocation, the learner behaves in an *adaptive* manner, by deciding what to study based on past experience.

3 Adaptive Learning

As an example, consider an adaptive learning rule where time is allocated to the item for which recent investments of time have resulted in the greatest increases in competence. This rule directs the learner to allocate time to the task which has the steepest current slope. It is easily seen that if all learning curves are concave, and the learner simply wishes to maximize the average competence across tasks, then this adaptive rule results in optimal allocations regardless of the total amount

[2] In general, tasks that compete for the attention of the learner over any particular time horizon will be fairly similar in nature. For instance, in allocating an hour of study time, learners might choose among different topics on a curriculum, while in allocating time over long horizons, one might choose between learning a second language or a musical instrument.

of time available. To see why, note that this rule results in a final allocation with the following properties: (i) all items that receive positive allocations of time end up at points on their respective learning curves that have the same slope, and (ii) this common slope is at least as great as that for any item that receives no allocation of time. Taken together, these conditions are necessary and sufficient for optimality when all learning curves are concave and the learner maximizes average competence. Hence we have here an example of an adaptive learning rule that ensures optimal allocations under very limited learner information.[3]

Unfortunately, the same cannot be said if learning curves are non-concave, as in the right panel of Figure 1. For S-shaped learning curves (such as the logistic) the optimal allocation varies discontinuously with total available time under very general conditions (Son and Sethi, 2006). In such cases learning based on local cues will necessarily terminate in suboptimal allocations for certain values of total available time. This can be seen most easily in the case of two items with independent learning curves, as in the following example.

Insert Figure 2 here

Suppose that each of two items has its own independent learning curve $b_i = f_i(t_i)$, as shown in Figure 2.[4] Learner's maximize the sum of competencies $b_1 + b_2$. Table 1 shows the optimal allocation and the corresponding level of performance for various values of total time availability ranging from 3 to 10 units. For each of these values of total time, the table also shows the allocation and performance level that arises under the adaptive rule. The manner in which optimal

[3] When all learning curves are concave, adaptive learning can be optimal even for more complicated learner goals. For instance, if the learner values competence in some tasks more highly than in others, and hence wishes to maximize a weighted sum of competencies, a rule which allocates time to the task with the greatest appropriately weighted slope will result in optimal behavior.

[4] The figure is based on the specification $f_i(t_i) = (10/(1 + 100e^{-s_i - t_i}) - 0.1)/9.9$,with $s_1 = 2$ and $s_2 = 3$.

and adaptive allocations change as the total available time is varied is also depicted in the top two panels of Figure 3. Note that the adaptive allocations turn out to be optimal when total available time is sufficiently small ($T \leq 4$), or sufficiently large ($T \geq 9.5$). For intermediate values of T, however, the adaptive learner fails to allocate time optimally.

Insert Table 1 here

Note that the optimal allocation in the table varies *discontinuously* with the total available time T. (This is seen most clearly in the top panel of Figure 3.) There exists a threshold value of total available time T' such that the optimal allocation requires the learner to devote all study time to the second item when $T < T'$, and requires more time to be devoted to the first item when $T > T'$. For instance, when $T = 4$, optimality requires $t_1 = 0$ and $t_2 = 4$, while when $T = 5$, optimality requires $t_1 = 3$ and $t_2 = 2$. An increase in T therefore requires a *decrease* in time allocated to t_2. This kind of change cannot be attained under the adaptive rule, which generates allocations that vary *continuously* in T. An adaptive learner faced with total available time $T = 5$ will behave for the first four time units in precisely the same manner as a learner having total available time $T = 4$. It is possible that a learner will regret having allocated too much time to the second task once it becomes apparent that this is not optimal, but this could not have been anticipated without complete prior knowledge of the uptake functions. And once allocated, time cannot then be reallocated.[5] Note that there is a second threshold T'' such that, when $T > T''$, adaptive learning again results in optimal outcomes.

[5] The data in the table are based on the implicit assumption that time units are perfectly divisible but the basic conclusions clearly also hold if time can only be allocated in discrete units. Furthermore, the arguments made here could also apply to other resources (such as calories in the context of diet composition) that cannot be reallocated if one's prior choices are discovered to be suboptimal.

Insert Figure 3 here

The last column of the table indicates the extent to which the adaptive learning rule approximates optimality. This measure of efficiency is depicted in the bottom panel of Figure 3. (Note that the range of the y axis is 80-100%.) The level of efficiency can be close to 100% even when allocations generated by the adaptive rule are quite distant from those required by optimality. The reason for this is that even though the optimal allocation varies discontinuously with total available time T, optimal *performance* is continuous in T. Since adaptive performance is also continuous, it must be approximately optimal for values of T that are close to the thresholds at which optimal and adaptive performance start to diverge.[6]

In the two task example considered here, optimal and adaptive allocations start to coincide once the total available time T becomes sufficiently large (see the threshold T'' in Figure 3). This is because the learner is able to reach the concave segments of *both* learning curves. In more realistic settings with large numbers of potential tasks, it is unlikely that the total available time will be so abundant as to allow this to occur for all tasks. For instance, when cramming for multiple exams, learners are often severely time constrained. Even for tasks pursued over a long period of time (such as language learning) there exist numerous competing tasks that might prevent the convergence of adaptive and optimal allocations.

We conjecture that the greater the number of tasks that compete for the learner's attention, the larger will be the range of values of T for which adaptive learning will be sub-optimal. Consider,

[6] We have assumed here that although learners have access only to local information about the steepness of learning curves, such information is available to them without noise. Intuition might suggest that imprecision in the estimation of learning curves would make it even more difficult for adaptive learners to reach optimal allocations, but this need not always be the case. For instance, when adaptive learning is highly suboptimal, errors in judgment might cause a learner to allocate time across tasks in a manner that better approximates optimality.

for instance, a learner faced with a third item in addition to the two described above. The optimal and adaptive allocations, and the level of efficiency attained by the adaptive rule, are depicted in Figure 4.[7] In this case there are two points of discontinuity in the relationship between optimal allocations and total available time, and two distinct ranges over which adaptive learning results in suboptimal allocations. The overall range of values of total available time over which some degree of suboptimality arises is considerably greater than in the two item case, as one might expect.

Insert Figure 4 here

How general is the phenomenon described here? When optimal allocations vary discontinuously in T, there must be *some* range of values of total available time for which *no* adaptive learning rule can result in optimal choices. This is a direct consequence of the fact that, at points of discontinuity, the amount of time optimally invested in some items jumps downwards even though the time budget has increased. If time allocation were reversible, such optima might be found by a process of trial and error. But the impossibility of such reallocations implies that there exists a range of values of the time budget for which terminal allocations must be suboptimal for all local allocation rules. Figures 3-4 illustrates this for a specific adjustment rule, but the point is considerably more general. Any algorithm that depends only on knowledge of portions of the uptake function that have already been experienced must result in allocations that vary *continuously* and *monotonically* with total available time. Whenever the optimal allocation exhibits discontinuities and non-monotonicities, therefore, there will be levels of time pressure at which such allocations can never be reached. Such discontinuities arise generically in the case of S-shaped learning curves such as the logistic (Son and Sethi, 2006).

[7] The third learning curve has the same functional specification as the other two, but with $s_3 = 4$ (see footnote 4 for details).

4 Conclusions

We have considered here the simplest possible learning environment, in which time allocated to one task does not affect competence in any other, there are no errors in the implementation of the adaptive learning rule, and items once learned are not subsequently forgotten. Even in this simple setting, adaptive rules fail to reach optimal allocations under certain conditions. An interesting extension of our work would be to incorporate forgetting and explore the extent to which adaptive rules can locate optimal spacing strategies and account for experimental findings (Pavlik and Anderson, 2005).

Time allocation is an irreversible process, which cannot be undone if it is found to be suboptimal. When learners lack the information or cognitive capacity to compute optimal allocations *ex-ante*, they are often forced to rely on local cues to adjust their decisions during the process of allocation itself. In certain cases, such as concave learning curves that are independent across items, there exists a simple adjustment rule that ensures optimal terminal allocations. Even in such cases, however, precise local knowledge of empirical learning curves is necessary in order for adaptive learning to result in optimal allocation. With certain other commonly observed learning curves such as the logistic, achieving optimality through adaptive adjustments is an even greater challenge, as learning based on local cues can result in suboptimal terminal allocations at least for some values of the time budget.

This finding has implications for educational policy. In deciding whether learners should be self-guided or instructor-guided in making time allocation decisions, it is important to first explore the general properties of the uptake functions that govern the relationship between time investments and attained competence. For complex tasks which are likely to have slow but accelerating rates of uptake at low levels of investment, self-guided learning may result in allocation levels that are suboptimally low. As most parents are intuitively aware, children often give up on tasks that are initially perceived to be too challenging, such as learning a musical instrument or a second language, even though sustained investments of effort would eventually lead to significantly accelerated rates of learning. Such investments may come to be seen as wise with the benefit of hindsight, but would not have been made by learners acting on local cues in the absence of external pressure. These are precisely the circumstances in which guided learning can result in heightened overall performance across a broad range of tasks relative to self-guided study.

Total Time	Optimal			Adaptive			Efficiency
	Task 1	Task 2	Performance	Task 1	Task 2	Performance	
3.00	0.00	3.00	0.86	0.00	3.00	0.86	100%
3.25	0.00	3.25	0.90	0.00	3.25	0.90	100%
3.50	0.00	3.50	0.93	0.00	3.50	0.93	100%
3.75	0.00	3.75	0.95	0.00	3.75	0.95	100%
4.00	0.00	4.00	0.97	0.00	4.00	0.97	100%
4.25	2.63	1.62	1.00	0.04	4.21	0.99	99%
4.50	2.75	1.75	1.06	0.29	4.21	1.01	95%
4.75	2.88	1.87	1.13	0.54	4.21	1.03	92%
5.00	3.00	2.00	1.19	0.79	4.21	1.06	89%
5.25	3.13	2.12	1.25	1.04	4.21	1.10	88%
5.50	3.25	2.25	1.30	1.29	4.21	1.13	87%
5.75	3.38	2.37	1.36	1.54	4.21	1.18	87%
6.00	3.50	2.50	1.41	1.79	4.21	1.23	87%
6.25	3.63	2.62	1.46	2.04	4.21	1.29	88%
6.50	3.75	2.75	1.51	2.29	4.21	1.35	89%
6.75	3.88	2.87	1.56	2.54	4.21	1.41	90%
7.00	4.00	3.00	1.60	2.79	4.21	1.47	92%
7.25	4.13	3.12	1.64	3.04	4.21	1.53	94%
7.50	4.25	3.25	1.67	3.29	4.21	1.59	95%
7.75	4.38	3.38	1.71	3.54	4.21	1.65	96%
8.00	4.50	3.50	1.74	3.79	4.21	1.69	98%
8.25	4.63	3.63	1.76	4.04	4.21	1.74	98%
8.50	4.75	3.75	1.79	4.29	4.21	1.77	99%
8.75	4.87	3.88	1.81	4.54	4.21	1.80	100%
9.00	5.00	4.00	1.83	4.79	4.21	1.83	100%
9.25	5.12	4.13	1.85	5.04	4.21	1.85	100%
9.50	5.25	4.25	1.87	5.25	4.25	1.87	100%
9.75	5.38	4.38	1.88	5.38	4.38	1.88	100%
10.00	5.50	4.50	1.89	5.50	4.50	1.89	100%

Table 1. Optimal and adaptive allocations for varying levels of time availability.

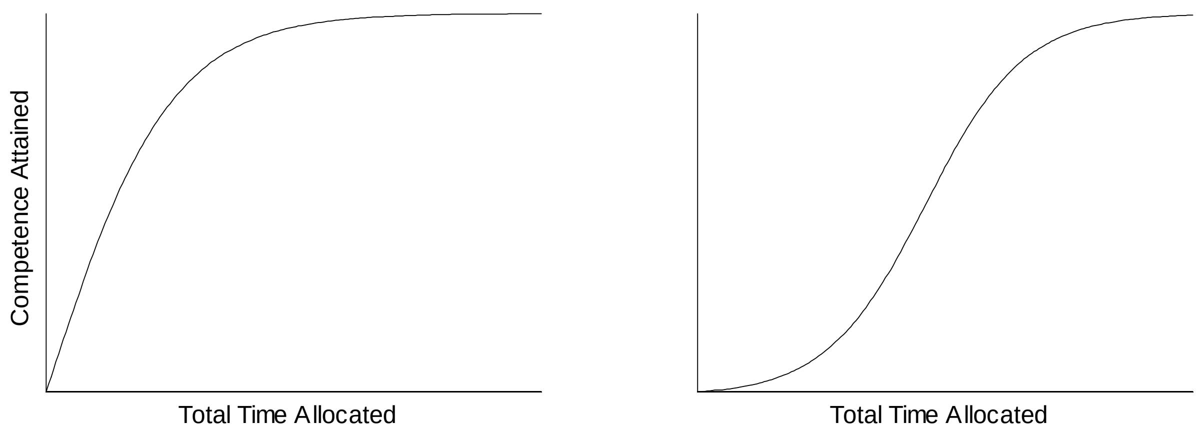

Figure 1. Two examples of learning curves: concave (left panel) and logistic.

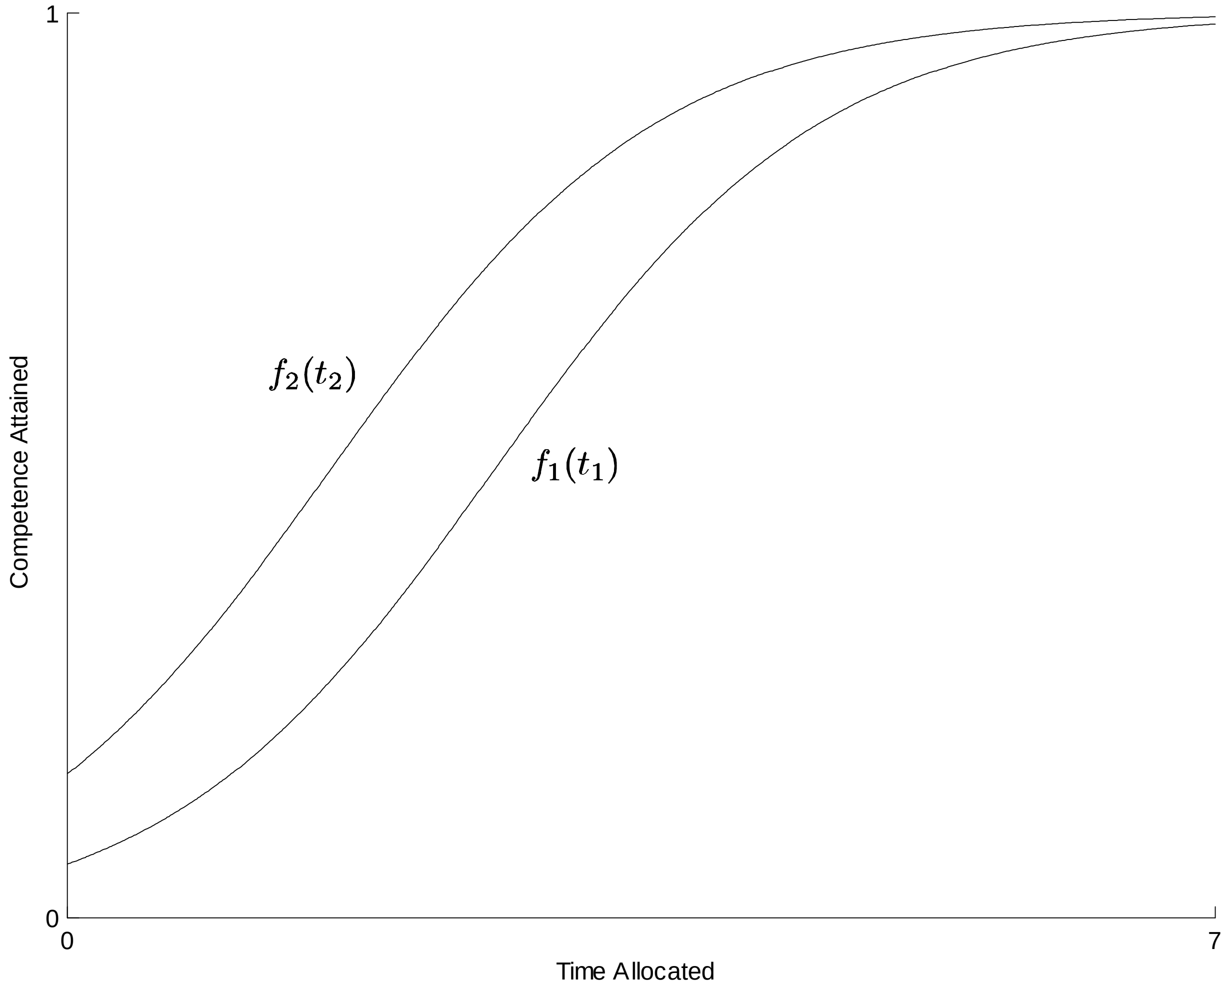

Figure 2. Learning curves for which adaptive rules can lead to suboptimal allocations.

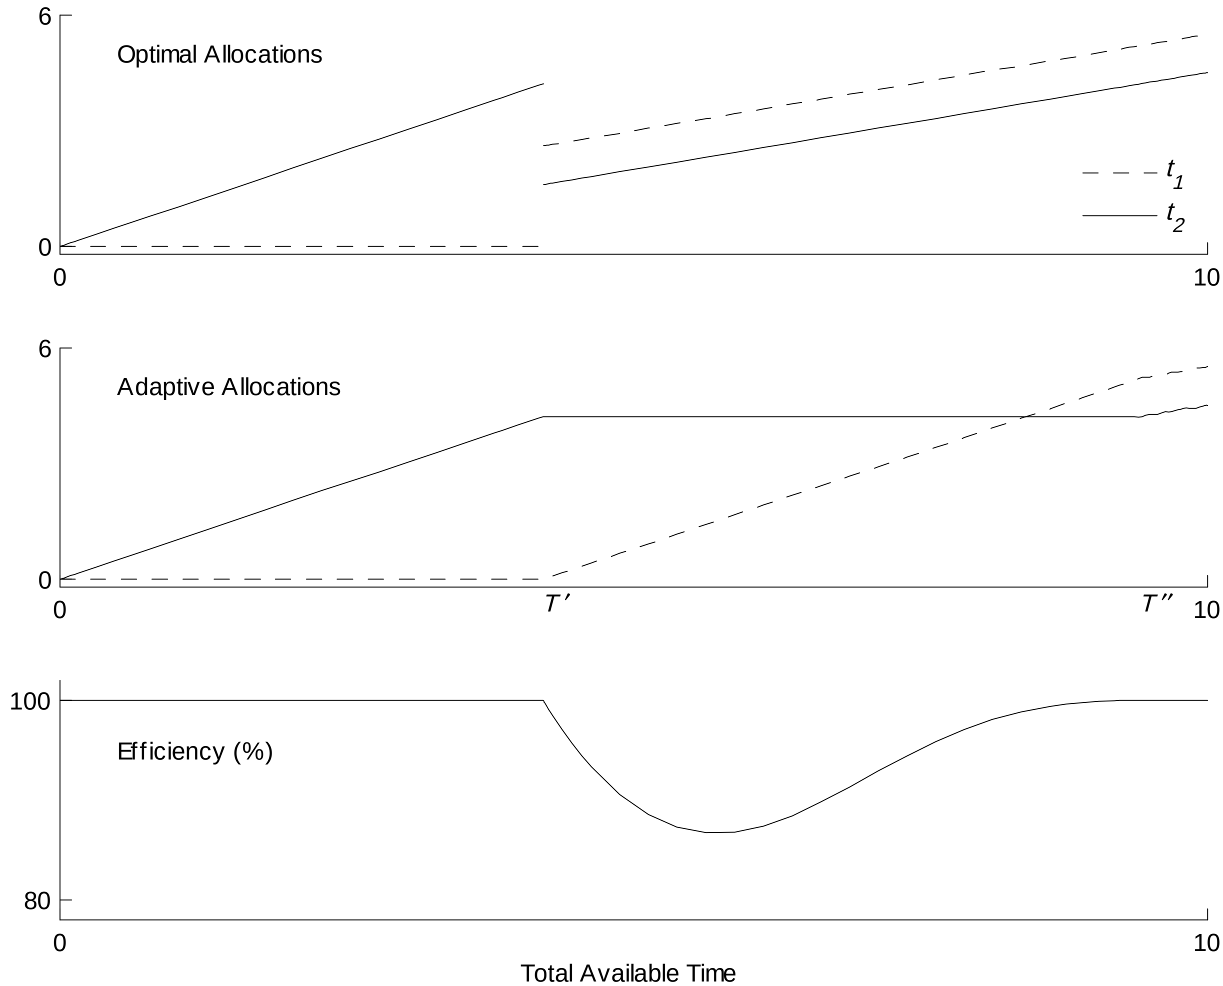

Figure 3. Optimal and adaptive allocations with two items.

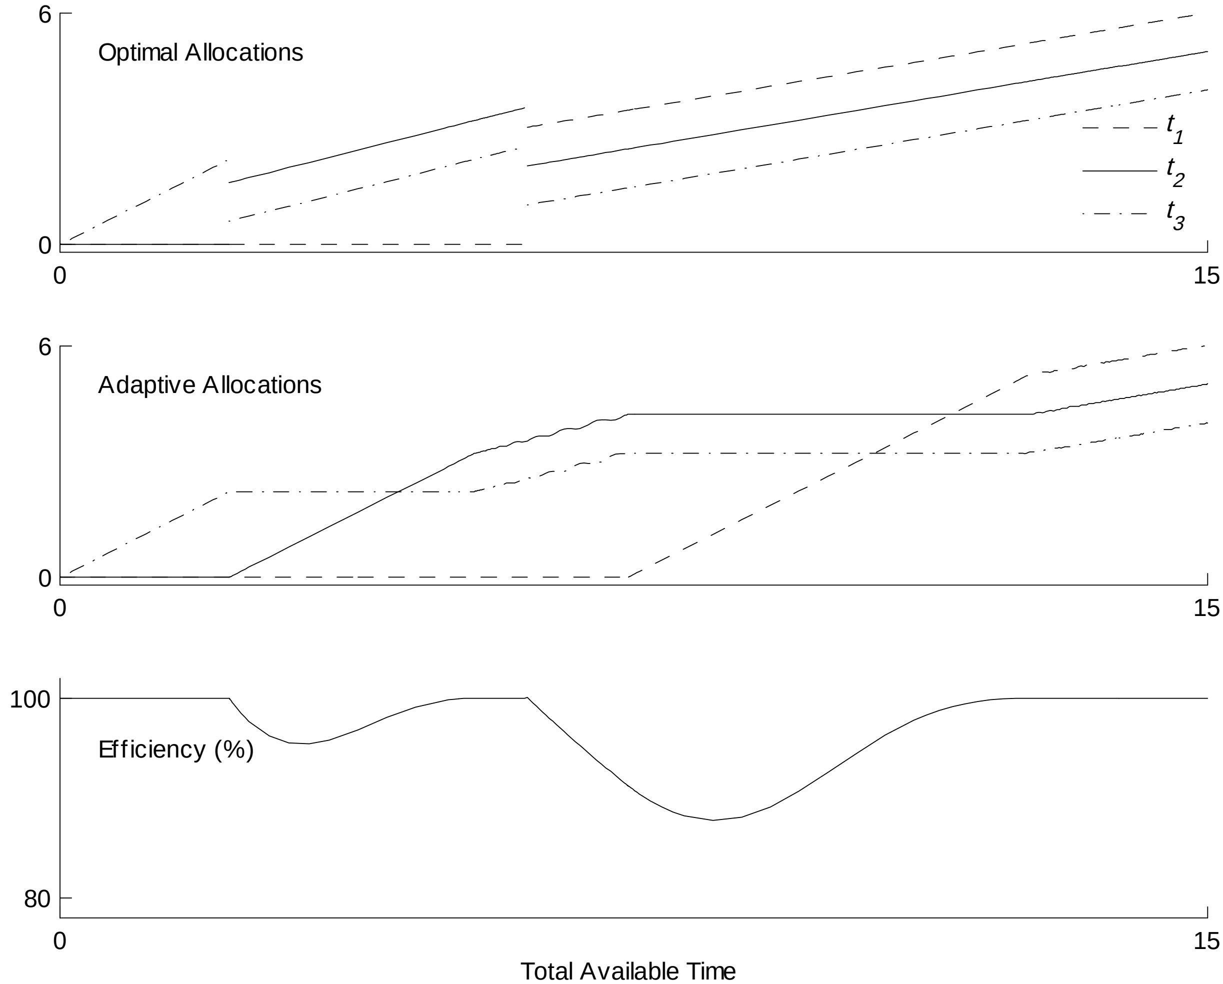

Figure 4: Optimal and adaptive allocations with three items.

References

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

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

[3] Fischer, K. W., & Pipp, S. L. (1984). Processes of cognitive development: Optimal level and skill acquisition. In R.J. Steinberg (Ed.), *Mechanisms of cognitive development* (*pp. 45–80*). San Francisco: Freeman.

[4] Frey, P. W. & Sears, R. J. (1978). Models of conditioning incorporating the Rescorla–Wagner associative axiom, a dynamic attention process, and a catastrophe rule. *Psychological Review*, *85*, 321–340.

[5] Klopf, A. H. (1988). A neuronal model of classical conditioning. *Psychobiology*, *16*, 85–125.

[6] Kornell, N., & Metcalfe, J. (2006). Study efficacy and the region of proximal learning framework. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, *32*, 609-622.

[7] Mazzoni, G., Cornoldi, C., & Marchitelli, G. (1990). Do memorability ratings affect study-time allocation? *Memory & Cognition*, *18*, 196-204.

[8] 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.

[9] Metcalfe, J., Kornell, N., & Son, L. K. (2007). A cognitive-science based program to enhance study efficacy in a high and low-risk setting. *European Journal of Cognitive Psychology*, *19*, 743-768.

[10] 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.

[11] Newell, K. M., Liu, Y. T., & Mayer-Kress, G. (2001). Time scales in motor learning and development. *Psychological Review*, *108*, 57–82.

[12] Noble, C. E. (1957). The length-difficulty relationship in compound trial-and-error learning. *Journal of Experimental Psychology, 55*, 246-252.

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

[14] Rice, M. L., Wexler, K., & Hershberger, S. (1998). Tense over time: The longitudinal course of tense acquisition in children with specific language impairment. *Journal of Speech, Language, and Hearing Research, 41*, 1412-1431.

[15] Ritter, F.E. & Schooler, L.J. (2001). The learning curve. *International Encyclopedia of the Social & Behavioral Sciences*, pp. 8602-8605.

[16] Son, L. K. & Kornell, N. (2008). Research on the allocation of study time: Key studies from 1890 to the present (and beyond). In J. Dunlosky & R. A. Bjork (Eds.), *A handbook of memory and metamemory* (pp. 333-351). Hillsdale, NJ: Psychology Press.

[17] Son, L. and Sethi, R. (2006). Metacognitive control and optimal learning. *Cognitive Science*, *30*, 759-774.

[18] Thiede, K. T., & Dunlosky, J. (1999). Toward a general model of self-regulated study: An analysis of selection of items for study and self-paced study time. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *24*, 1024-1037.