Journal of Mathematical Psychology **46**, 455–485 (2002)
doi:10.1006/jmps.2001.1400

A Dynamic Model for Rule Induction Tasks

Tom Verguts

University of Leuven

Eric Maris

University of Nijmegen

and

Paul De Boeck

University of Leuven

In this paper, a model for performance on rule induction tasks (e.g., items on intelligence tests) is developed. The model simultaneously specifies distributions for response times and response accuracies on an item-by-item basis. It is dynamic in the sense that it can be used to specify and test different ways of learning throughout a test. Three versions of the general model (i.e., with three different learning rules) are described and the fit of these versions is investigated in two datasets on solving number series. The results indicate that one of these versions (one of the learning rules) is better at accounting for the data. © 2002 Elsevier Science (USA)

INTRODUCTION

This paper develops a model to account for human problem solving in rule induction tasks. The core assumption in this model is that solving items of this type consists of sampling a sequence of possible solution rules (or solution principles). For example, possible rules for solving integrals are "substitute" and "decompose into fraction." Such rules are sampled until a useful one is found.

Each rule is assumed to have a certain level of *activation*. The probability of sampling a rule as well as the time needed to perform this sampling depend on the activation of that particular rule. Activations of individual rules can change during problem solving and this is the way that learning is incorporated in the model.

We thank William Batchelder, Siegfried Dewitte, Martin Schrepp, and Francis Tuerlinckx for their useful comments on the topic or on previous versions of this paper.

Address correspondence and reprint requests to Tom Verguts, Department of Psychology, Tiensestraat 102, 3000 Leuven, Belgium. E-mail: Tom.Verguts@psy.kuleuven.ac.be.

0022-2496/02 $35.00
© 2002 Elsevier Science (USA)
All rights reserved.

Probability distributions for accuracy and response time are derived from these assumptions.

It has been amply demonstrated that the process of finding solutions in complex cognitive tasks (such as rule induction tasks) is affected by one's history of successes and failures (i.e., learning). For example, this has been demonstrated for anagram problems (Kaplan & Schoenfeld, 1966; Lemay, 1972; White, 1988), the water jars task (Shore, Koller, & Dover, 1994), the stick building task (Lovett & Anderson, 1996; Lovett & Schunn, 1999), insight problems (Duncker, 1945; Maier & Burke, 1966), and baseball knowledge (Wiley, 1998). This paper presents a model for studying this phenomenon in more detail than is possible with traditional models.

The model is applied to number series performance. One detailed model for (letter) series performance was proposed by Simon and Kotovsky (1963). This model consists of an algorithm describing the various steps people are presumed to follow in solving such problems. Most research on this task builds on that paper (Kotovsky & Simon, 1973; Holzmann, Pellegrino, & Glaser, 1982, 1983; Schrepp, 1995). The model is a *static* model, in the sense that the probability of solving a particular item remains unchanged throughout the test.

An influential series of models that can capture learning throughout a test is the series of ACT models by Anderson and colleagues (e.g., Anderson, 1983, 1993; Anderson & Lebiere, 1998). One important difference between their models and the one we present is that ours is simpler, which makes it to easier to (a) estimate parameters, (b) investigate goodness-of-fit of the model, and (c) empirically distinguish different versions of the model with different substantive psychological interpretations.

The remainder of the paper consists of five parts. In the first section (*Number Series Task*) we briefly discuss the rule induction task that is used in our two experimental studies described later. This task is introduced in an early phase because the explanation of the model itself builds on this task. In the second section (*Model*), we derive the relevant equations. In the third section (*Study 1*) the model is applied to number series data without a time limit. In the fourth section (*Study 2*) the model is applied to number series data *with* a time limit. Finally, we close the paper with a general discussion.

NUMBER SERIES TASK

Each item of our task is a number series[1] such as "3, 5, 7, 9, 11, 13." The number series always consists of six numbers. Each such item can be described according to a certain logical rule. Four rule types are used: addition, Fibonacci, interpolation, and multiplication. An example of the *addition* type is "3, 4, 6, 9, 13, 18." In an item of this type, the increment between two consecutive numbers is itself incremented by one (so the sequence of increments is +1, +2, +3, +4, +5). The increment between two successive increments is always equal to one, but items

[1] Mathematically, the term "number series" is not completely appropriate since the task involved concerns number *sequences*, not series. However, we follow the tradition in this literature and refer to this task as a number series.

can differ in starting number and initial increment. In items of the *Fibonacci* type, each number is the sum of the previous two. The first two numbers are chosen arbitrarily, with the second number larger than the first. An example is "2, 4, 6, 10, 16, 26." In items of the *interpolation* type, two series are interpolated, each one with a constant increment rule: An example is "2, 5, 6, 8, 10, 11," which is a mix of the sequences "2, 6, 10, ..." and "5, 8, 11, ..." In items of the *multiplication* type, the increment is always multiplied by a constant value. An example is "1, 3, 7, 15, 31, 63." The multiplier is always equal to 2. Our rule induction task is constructed by creating a test of 50 such items, with items of the four types intermixed. The complete set of items is presented twice so in fact 100 items are presented to the participants. More details on test administration are presented later on.

MODEL

In the first subsection (*Response Process*) the general model is presented. The second subsection (*Activation Dynamics*) describes the three versions of the model, each corresponding to a different learning rule. In the third subsection (*Time Limits*), we discuss the changes to the model that are required when a time limit per item is imposed. In the fourth subsection (*Estimating and Testing the Model*), we discuss how to estimate the parameters of the model and how to investigate its goodness-of-fit.

Response Process

Suppose K possible rules can be used to solve all items of a certain test. In the number series example of the previous subsection, $K = 4$; this number is imposed by the experimenter. When shown an item, one of these rules is sampled. Sampling a particular rule is conceptualized as a *race* between the different rules (e.g., Marley, 1989). Specifically, each rule k is supposed to have a finishing time that is exponentially distributed with hazard rate parameter w_k. That is, if V_k denotes the finishing time for rule k, its probability density function (PDF) is given by $p(V_k) = w_k \exp(-w_k V_k)$. The first rule to finish the race is sampled and tried out by the participant. For example, if the addition rule described in the previous section wins the race, the participant tries to apply this rule to the item at hand. It is easily demonstrated (e.g., Ross, 2000, p. 295, Exercise 7) that if the race times V_k are statistically independent, then the probability that rule k is tried out is given by

$$\Pr(\text{rule } k \text{ chosen}) = \frac{w_k}{\sum_m w_m},$$

where the summation is taken over all rules.

If the rule that is sampled turns out to be inapplicable to the item, the process is started anew. That is, a new race is set up between all possible rules and the fastest rule is again tried out. Note that the sampling is *with replacement*. If the correct rule wins the race, it is applied and the solution process is terminated. However, some rules may be incorrect but still lead to a termination of the iterated sampling process. For example, suppose a participant incorrectly concludes that the series

"1, 3, 5, 7, 9" is governed by the rule "each number is the previous one multiplied by three." In this case, the sampling process will stop, although the correct rule is not found. Another possible "rule" which may stop the solution process is the answer "I don't know." In this case, the participant decides she cannot solve the item and goes on to the next one. We aggregate all such misleading rules (misleading in the sense that they incorrectly lead to a termination of the solution process) into an additional response category $(K+1)$, with corresponding activation w_{K+1}. Hence, the race between rules is iterated until either the correct rule is found or one of the misleading ones wins the race.

Let us denote the activation of the correct rule by w_C (so $C \in \{1, \ldots, K\}$). Then, this iterated sampling process can be conceptualized as a three-state Markov chain, with the three states corresponding to (a) having sampled the correct rule (denoted state C), (b) having sampled the misleading rule (denoted state M), and (c) having sampled one of the remaining rules (a rule with an index $k \in \{1, \ldots, C-1, C+1, \ldots, K\}$), denoted state R. The transition probability matrix is given by

	C	M	R
C	1	0	0
M	0	1	0
R	π_C	π_M	$1-\pi_C-\pi_M$

(1)

where $\pi_C = w_C / \sum_{k=1}^{K+1} w_k$ and $\pi_M = w_{K+1} / \sum_{k=1}^{K+1} w_k$. The process starts in state R.

It is now possible to derive the probability of success and the response-time distribution on an item-by-item basis. Let us introduce the item index i $(= 1, \ldots, I)$. This index is attached to the variables introduced earlier. For example, w_{ki} now denotes the activation of rule k at the presentation of item i. Further, let $C(i)$ denote the correct rule number for item i. For example, if rule 2 is the correct one for item $i = 1$, then $C(1) = 2$. Then, the probability of sampling the correct rule at item i is equal to $\pi_{C(i)\,i}$.

An item is solved correctly if the chain eventually ends up in state C (rather than state M). Let X_i (with realizations x_i) denote the accuracy variable. That is, $X_i = 1$ if item i is solved correctly and $X_i = 0$ otherwise. It then follows from elementary stochastic process theory (e.g., Ross, 1996, Proposition 4.4.2) that

$$\Pr(X_i = 1) = \frac{\pi_{C(i)\,i}}{\pi_{C(i)\,i} + \pi_{(K+1)\,i}} = \frac{w_{C(i)\,i}}{w_{C(i)\,i} + w_{(K+1)\,i}}. \tag{2}$$

Let T_i (with realizations t_i) be the variable denoting the decision time. This is the time needed to reach either one of the two states C or M. It is now possible to show that

$$p(T_i = t_i \mid X_i = x_i) = p(T_i = t_i) = (w_{C(i)\,i} + w_{(K+1)\,i}) \exp(-(w_{C(i)\,i} + w_{(K+1)\,i})\, t_i), \tag{3}$$

so T_i follows an exponential distribution. This is shown in Appendix A (Result 3).

Combining Eqs. (2) and (3), we get the following PDF for the responses (accuracies and decision times) to all I items jointly:

$$p(\text{data}) = \prod_{i=1}^{I} \Pr(X_i = x_i)\, p(T_i = t_i \mid X_i = x_i)$$

$$= \prod_{i=1}^{I} (w_{C(i)\,i})^{x_i}\, (w_{(K+1)\,i})^{1-x_i} \exp[-(w_{C(i)\,i} + w_{(K+1)\,i})\, t_i]. \tag{4}$$

The probabilities in these equations are conditional on the parameters w_{ki}. Thus, *given* these parameters, the different X_i and T_i are all statistically independent.

The observed response time is assumed to be composed of two parts, a decision time and a motor response time. Until now, we have only focused on the decision time. The motor response time is the time needed to transform a decision into a motor response (pressing a button, for instance). For simplicity, the motor response time is assumed to be a person-specific constant t_{re}. The observed response time can hence be denoted $T_{i,re} = T_i + t_{re}$. In our analysis, we first estimate the person-specific constants t_{re} and then calculate $T_i = T_{i,re} - t_{re}$. These T_i are then treated as observed decision times that are assumed to be governed by Eq. (4).

Although the assumption of a constant motor-response time is an oversimplification, it should be noted that for the tasks modeled the motor response time is relatively small as compared to the decision time. Also, it is not our intention to describe the response time distribution in great detail. Rather, we will try to model the change in mean response time (the mean taken over blocks of items) to find out what this change tells us about the process of solving number series.

Activation Dynamics

Activations w_{ki} may change over items as a function of experience and feedback. In general, activations change according to the formula

$$w_{ki} = f(w_{k(i-1)}). \tag{5}$$

That is, the current activation is a function of the activation on the previous item. It may be that w_{ki} also depends on a few parameters, but this is implicit in Eq. (5). We will consider three special cases of f.

1. The activations are constant throughout; that is, $w_{ki} = w_{k(i-1)}$. The resulting model will be called the *constant model* in the remainder.
2. The activations of the responses w_{ki} change according to the formula

$$w_{ki} = \begin{cases} \alpha w_{k(i-1)} + (1-\alpha)\,\theta & \text{if item } i-1 \text{ is of type } k \text{ and } X_{i-1} = 1, \\ \alpha w_{k(i-1)} & \text{otherwise.} \end{cases} \tag{6}$$

This model will be called the *incremental model*, which is shorthand for "incremental learning model." In Eq. (6), $0 \leqslant \alpha \leqslant 1$ and $\theta \geqslant 0$. This equation states that if the

previous item $(i-1)$ is of type k *and* that item is solved correctly, the activation of rule k increases. Otherwise, the activation decreases. For example, if an item is not solved correctly, then all activation values will decrease. Equation (6) is similar to formulas that were used in earlier mathematical models of learning (e.g., Bush & Mosteller, 1955; Luce, 1959).

If a rule is not useful or cannot be found for some time, its activation will decrease at an exponential rate. In particular, the activation of the $(K+1)$th response, which is sometimes used but is never useful, can be written as $w_{(K+i)\,i} = \alpha^{i-1} w_{(K+1)\,1}$. This assumption may be unrealistic because the tendency to give up (which is one of the responses in the $(K+1)$th category) may increase rather than decrease over time. However, if the giving-up option is not often chosen, the model may very well be appropriate (see the Results and Discussion sections of Studies 1 and 2).

The parameters w_{k1} refer to the initial biases of a participant. If w_{k1} is high, the participant is likely to sample rule k right from the start when responding to the first item. The parameter θ partly determines how strongly rules are activated by a previous successful application. Indeed, if θ is high, the activation w_{ki} will increase strongly after successfully solving an item (cf. Eq. (6)). As for the parameter α, the higher it is (the closer to 1), the slower the activations change, as can be seen from Eq. (6), and hence the slower the response distributions change. The parameter α may therefore be called a *stability* parameter.

3. The third model is the *all-or-none model*, in which the activations w_{ki} can be in only two possible states. We formalize this learning rule as follows:

$$w_{ki} = \begin{cases} \theta & \text{if rule } k \text{ is used successfully earlier,} \\ w_{k1} & \text{otherwise.} \end{cases} \tag{7}$$

When a rule is activated, it jumps to the point θ and never leaves that position. This explains the name "all-or-none model." Note that, in this model, the activation of the $(K+1)$th response is constant throughout the test ($w_{(K+1)\,i} = w_{(K+1)\,1}$ for all items i).

The model in Eq. (7) does not correspond in detail to the all-or-none models presented earlier in the literature (e.g., Restle, 1962, 1965). However, the old all-or-none models and the present one share the philosophy that a learner can be in one of only two states, corresponding to unlearned (or weakly active) and learned (or strongly active), respectively. A more detailed comparison between the current and the older all-or-none models is provided in the General Discussion. Figure 1 presents examples of how activities develop in all three models (constant, incremental, and all-or-none).

Time Limits

Suppose only a limited time, say 15 seconds, is allowed per item. Then, every participant has only $15 - t_{re}$ s (15 minus the motor response time) to think. This threshold $(15 - t_{re})$ will be denoted by τ.

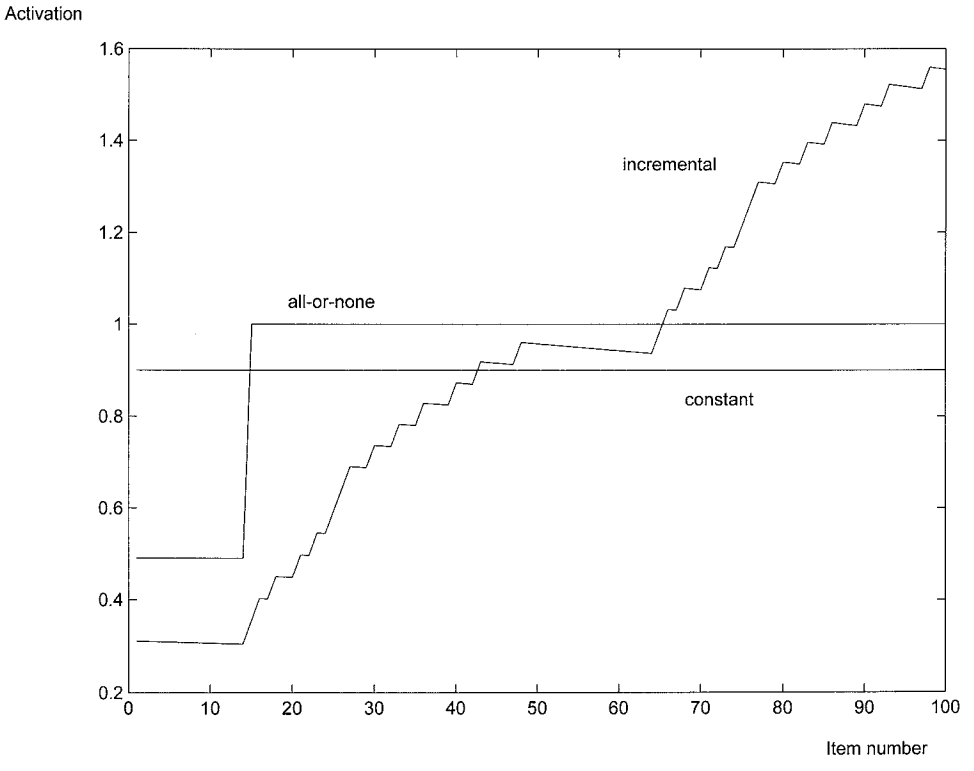


FIG. 1. Activation development in the constant, incremental, and all-or-none models.

In the time-limited case, the variables X_i and T_i are defined as before, but they are not always observed. In particular, X_i and T_i are not observed when $T_i \geqslant \tau$. We therefore define a new variable to denote accuracy. This new variable, Y_i, equals 1 if both $X_i = 1$ and $T_i < \tau$ and equals 0 otherwise. Then,

$$\Pr(Y_i = 1) = \frac{w_{C(i)\,i}}{w_{C(i)\,i} + w_{(K+1)\,i}} \{1 - \exp(-(w_{C(i)\,i} + w_{(K+1)\,i})\,\tau)\}. \tag{8}$$

This is derived in Appendix A (Result 4). In Appendix A, we also derive the PDF of accuracies and response times to all I items jointly, for the case of time-limited item presentation (Result 5).

Estimating and Testing the Model

We first consider the estimation of the motor response time t_{re}. This parameter is estimated by $\min_i\{T_{i,re}\}$. This procedure can be justified as follows. According to the model, $T_{i,\mathrm{re}} = T_i + t_{\mathrm{re}}$, in which t_{re} is constant and T_i has an exponential distribution. Assume that T_i has a constant hazard rate w and that I observations $T_{i,\mathrm{re}}$ are given. Then,

$$E[\min_i \{T_{i,re}\}] = t_{\mathrm{re}} + \frac{1}{Iw}, \tag{9}$$

since the minimum has itself an exponential distribution with parameter Iw. This result shows that the estimator $\min_i\{T_{i,re}\}$ is biased, but also that the bias is small if the sample size (I) is sufficiently large. It might be suspected that this procedure is susceptible to outliers at the left-hand side of the response-time distribution, but we will show that there are no such outliers in our data.

After subtracting the estimated motor response times from the observed response times, these *adjusted* response times are entered in the likelihood function, which is then maximized. Thus, maximum likelihood estimates are computed on accuracies and adjusted response times. The maximization is performed by a *protected steepest ascent* algorithm which is described in Appendix B. The estimation can be performed for every participant separately or for all participants together. The second procedure will be used when parameters are restricted to be equal across participants to increase parsimony of the model.

For the constant and the all-or-none models, parameters will be estimated for each participant separately. For the incremental model, parameters are estimated for all participants together. This is because the parameter α is restricted to be equal across participants. Restricting α to be equal across participants is motivated by the fact that we found this parameter to vary little over participants. Different w_{k1} (initial activation) and θ parameters are assumed for every participant. The θ parameter is assumed to be constant across rules. Hence, the resulting number of parameters of the incremental model is $(K+2)\,P+1$, where P denotes the number of participants. After this restriction, the incremental model has only one parameter more than the all-or-none model. Specifically, the incremental and all-or-none models now have $(K+2)\,P+1$ and $(K+2)\,P$ parameters, respectively.

The model will be tested using the *parametric bootstrap* procedure (Davison & Hinkley, 1997). This procedure is as follows:

1. Using the estimated parameter vector $\hat{\xi} = (\hat{\theta}, \hat{\alpha}, \ldots, \hat{w}_{(K+1)1})$ for a given person, a new data matrix $\mathbf{D}^{\text{rep},n}$ is generated N times ($n = 1, \ldots, N$).

2. Consider a statistic S (different possibilities for S will be discussed later). Let S^{obs} denote the statistic S as evaluated in the observed data. For each data set n, we check whether $S^{\text{rep},n} \geqslant S^{\text{obs}}$, where $S^{\text{rep},n}$ denotes the statistic evaluated at the nth replicated data set $\mathbf{D}^{\text{rep},n}$.

3. The proportion of times that $S^{\text{rep},n} \geqslant S^{\text{obs}}$ is our bootstrapped p-value. Small values (e.g., below .05) indicate that the model is inadequate.

The idea of this procedure is to approximate the true sampling distribution of the statistic S by replacing the true parameter vector (ξ) with its estimate ($\hat{\xi}$) and to use the resulting distribution as the reference distribution for the observed value of S. Thus, the bootstrapped p-value is an approximation of the actual p-value.

Recording the number of times that $S^{\text{rep},n} \geqslant S^{\text{obs}}$ (see Step 2) is useful for goodness-of-fit statistics in which a large S^{obs} indicates a bad fit and a small S^{obs} indicates a good fit (e.g., a Pearson chi-square statistic). This will be called a *one-sided* bootstrap test in the following. However, it is not appropriate for statistics in which both very small and very large values of S^{obs} indicate a bad fit. In these cases, the

p-value will be calculated *two-sided* instead of one-sided. Specifically, in the two-sided case we will check whether

$$|S^{\text{rep},n} - \overline{S^{\text{rep}}}| \geqslant |S^{\text{obs}} - \overline{S^{\text{rep}}}|,$$

for all $n = 1, \ldots, N$. In this case, if the p-value is low, this indicates that S^{obs} is far from the mean of S^{rep} (which is denoted by $\overline{S^{\text{rep}}}$). In the following, when a test statistic is presented, it will also be indicated whether a one- or two-sided test is performed. The value of N, the number of bootstrap samples, is set to 100 (per person and per statistic).

STUDY 1

In this section, we apply the model to a data set on number series performance without a time limit. At this point, a note concerning the uniqueness of the rules is in order. Korossy (1998) discusses the fact that even in a very restricted class of items there are often many rules that are possible for a single item. He advises researchers to tell participants the class of rules under study, so participants can pick one of the intended rules. However, this is neither necessary nor advisable in the present study, for the following reasons. First, it is not advisable since the ecological validity of the task would suffer: Telling the intended rules would make the test much easier and the test could hardly be called a rule discovery test in this case. Second, it is not necessary because we ask participants to *state the rule* they are using, rather than to complete the item series. Hence, if people indeed use alternative rules, we would notice this. However, this did not happen in the studies we report on here. People either use the rules we used to construct the item series or rules that are incorrect or inaccurate.

Method

Material. The number series task was described earlier. There are four rule types: addition, Fibonacci, interpolation, and multiplication. In items of the interpolation type, the series can in principle be nonincreasing, but we constructed the items so that this did not occur (i.e., the series is always increasing). For the multiplication type, the multiplier used is always equal to 2. The reason for these two conventions will be explained, later.

A test of 50 items is constructed and presented twice in identical order. Hence, $I = 100$ items are presented. The different rule types are introduced gradually. In items 1 to 10, the addition and Fibonacci types are used, each rule being used 5 times out of 10. ln items 11 to 20, the interpolation type is added to these two: Addition, Fibonacci, and interpolation rules are used 2, 4, and 4 times, respectively, in this set of 10 items. In items 21 to 30, the addition and interpolation types are used, each rule used 5 times out of 10. In items 31 to 40, the multiplication type is added to these two, with addition, interpolation, and multiplication being used 3, 3, and 4 times, respectively. Finally, in items 41 to 50 all 4 types are used, and addition, Fibonacci, interpolation, and multiplication are used 2, 3, 2, and 3 times respectively. The exact same sequence is repeated for items 51–100. The same sequence of items is presented to each participant.

Procedure. Items are presented on a computer screen. The participant is asked to find the rule involved in the item. If the rule is found, the participant has to press the spacebar, which causes a box to cover the item. In this way, we try to enforce that the participant can no longer think about the rule after pressing the space bar. She should type the rule in the box that appears after pressing the space bar. It should be stressed that the participant does not have to complete the number series, but simply has to describe the rule she thinks is the correct one. Hence, the process of finding the rule is separated from the process of applying it. The participant has unlimited time to type in the rule.

In the instructions, the procedure that is required of the participant (typing a rule) is illustrated with two examples of number series involving a constant increment (e.g., 2, 5, 8, 11, 14, 17), a rule that is not used in the actual test. Also, it is mentioned that two series may be intermixed: In this way, a hint is given about the interpolation series. This hint is given because in a pilot experiment it had turned out that the interpolation rule is very difficult to discover without a hint.

After the participant has typed his/her rule, the same series appears again with two extra numbers added. For example, if the series is "1, 3, 5, 7, 9, 11," the extended series would be "1, 3, 5, 7, 9, 11, 13, 15." This way, the participant receives feedback about the rule he/she typed because he/she can check whether the rule he/she used is in accord with the extended series. The time available to look at the extended series is unlimited.

Each participant is tested separately. Instructions are presented on the computer screen, and if everything is clear the participant starts solving the items. The experimenter does not intervene but always stays in the room. Participants are told that they are performing in a response time experiment and that they should work as fast as possible without compromising accuracy.

Data coding. After data collection, two types of data are registered per item, response time and accuracy, the latter indicating whether or not the rule described is the correct one. The coding is such that the participant is credited with a 1 if the correct principle is described even if some details in the description are incorrect. To make this more concrete, we now give some details on the scoring of the responses to the different rule types. For the addition type, a rule description is considered correct if it refers to increasing increments. For example, if the participant describes as a rule "+1, +2, +3, +4, ...," this is considered correct, but also if the participant writes "+2, +3, +4, +5, ..." for that item. For a Fibonacci item, a response is considered correct if there is some reference to the two previous numbers. For example, correct descriptions would be "add two numbers to obtain the next one," "the previous number and the last one give the following one," and so on. In the last description, adding is considered to be the default operation, and this is why this description is considered correct too, even though there is no explicit reference to the operation that should be performed. If a participant explicitly mentions an incorrect operation performed on the last two numbers, this would be considered an incorrect response. For an interpolation item, a response is considered correct if there is some reference to the mixing of two series. For example, in the series "1, 4, 5, 7, 9, 10, 13," if a participant's rule description is "for the odd numbers, +4, and

for the even numbers, also +4," this is considered correct because there is a clear reference to the mixing of series, even though the exact rule is not correct (for the even numbers, the relevant rule is +3 rather than +4). For a multiplication item, a response is considered correct if there is some reference to the multiplication of increments. In the item "3, 4, 6, 10, 18, 34," correct rule descriptions would be "+1, +2, +4, +8, ...," "multiply the differences by two," "differences are increasing powers of two," and so on. Some people also used the mathematically equivalent rule, "Each number is the previous one times two plus a constant." This was also considered correct. Finally, the additional response category $(K+1)$ was scored if the participant gave an incorrect rule or typed a give-up response such as "I don't know."

Participants. Eight persons participated for a small monetary reward. They were students of the University of Leuven who had responded to an advertisement.

Trends. Trends in the data will be examined. The two learning models (incremental and all-or-none) predict that accuracy (probability of success) will increase and response time will decrease over time. This is not predicted by the constant model. To investigate this, all instances of each item type are divided into two blocks as follows: The addition items that were presented in the first half of the test are assigned to Block 1, and the addition items that were presented in the second half are assigned to Block 2. The same procedure is followed for the Fibonacci, interpolation, and multiplication items.

We look at the mean accuracy and mean response times over these two blocks, for every (person, item type) pair separately, and report the number of times that the difference is in the direction expected by the two learning models. That is, according to the two learning models, accuracy should be higher in the second block than in the first and response time should be shorter in the second block than in the first.

Note that investigating the trend in the accuracies is exactly analogous to investigating the frequency of occurrence of the $(K+1)$th category, since, in this study, choosing the $(K+1)$th category implies $X_i = 0$, and vice versa. The same will no longer be true in the next study, where a time limit is imposed.

Goodness-of-fit-statistics. Here, we describe the goodness-of-fit statistics that will be used in the parametric bootstrap procedure to test the model. For the first statistic to be developed, smaller blocks are constructed than those described in the previous subsection in order to perform more fine-grained analyses. For item-type addition, blocks are constructed as follows: Addition items in the first quarter of the test are assigned to Block 1, those in the second quarter to Block 2, those in the third quarter to Block 3, and those in the fourth quarter to Block 4. A similar procedure is followed for Fibonacci and interpolation items. For type multiplication, the two-block structure of the previous paragraph is maintained since there are no multiplication items in the first quarter of the test. Hence, for addition, Fibonacci, and interpolation, four blocks are constructed, and for multiplication items only two.

Let the item type (addition, Fibonacci, interpolation, or multiplication, in that order) be indexed by $k = 1, \ldots, 4$. The first statistic to be used is defined as

$$Q_{\text{acc}}(k) = \sum_{b=1}^{B} \frac{[\bar{X}_{bk} - E(\bar{X}_{bk})]^2}{E(\bar{X}_{bk})}. \tag{10}$$

Here, $\bar{X}_{bk}$ denotes the mean accuracy of all items of type k within block b. The variable $E(\bar{X}_{bk})$ is its expectation, which is calculated as

$$E(\bar{X}_{bk}) = \frac{1}{n_{bk}} \sum_{\substack{i \text{ in block } b \\ \text{and of type } k}} E(X_i), \tag{11}$$

where n_{bk} denotes the number of items of type k in block b and $E(X_i)$ is used as a shorthand notation for $E(X_i \mid X_1, \ldots, X_{i-1})$. The dependence is needed because successes on previous items influence current success through the activation values w_{ki} (see Eqs. (6) and (7)). The summation in Eq. (11) is taken over all items belonging to block b and of type k (e.g., items 1, ..., 25 for $b = 1$ and item type addition). This statistic will be called the *Pearson accuracy statistic* because it is modeled after the familiar Pearson chi-square statistic. We will perform a one-sided bootstrap test for this statistic.

A similar discrepancy measure is calculated for response times,

$$Q_{\text{time}}(k) = \sum_{b=1}^{B} \frac{[\bar{T}_{bk} - E(\bar{T}_{bk})]^2}{E(\bar{T}_{bk})}. \tag{12}$$

In Eq. (12),

$$E(\bar{T}_{bk}) = \frac{1}{n_{bk}} \sum_{\substack{i \text{ in block } b \\ \text{and of type } k}} E(T_i),$$

where $E(T_i)$ is used as a shorthand notation for $E(T_i \mid X_1, \ldots, X_{i-1})$. Again, the dependence on $X_1, \ldots, X_{i-1}$ is needed because successes on previous items partially determine the w_{ki} values which are needed to calculate $E(T_i)$. The statistic Q_{time} will be called the *Pearson response time statistic*.

The third and fourth types of statistic are inspired by a statistical test proposed by Suppes and Ginsberg (1963), which was intended to check a critical aspect of an all-or-none learning model (see our General Discussion for a more elaborate description of that model). Their work was concerned with the paired-associate learning paradigm, and they noted that, prior to learning the (all-or-none) connection between a word and its associated word, successes on the stimulus word should be binomially distributed. This implies that the correct response probabilities are *stationary*: Prior to learning the association between words, the probability of successfully answering an item remains constant. Suppes and Ginsberg developed a test for this aspect of the data. Similar statistics can be constructed in our case, with the difference that we will check participants' behavior *after* a rule has been learned, rather than before. These statistics will be called *stationary* statistics. One will be constructed for the accuracy data and one for the response time data.

To construct the stationarity statistic for item type k, define item i_k to be the first item of type k on which a participant gives a correct answer (i.e., $x_{i_k} = 1$ and $x_i = 0$ for all items $i < i_k$ of type k). Divide all items $i > i_k$ of type k into blocks of five: The first five items i of type k for which it holds that $i > i_k$ constitute Block 1, the next five items constitute Block 2, and so on. Suppose there are B such blocks, and let $\bar{X}_{bk}$ be the mean accuracy of type-k items in a block b. (Note that this definition of $\bar{X}_{bk}$ is different from that used in Eqs. (10)–(12).) Finally, let $\bar{X}_k$ be the mean accuracy of all items $i > i_k$ of type k. Then, the *stationarity accuracy statistic* is defined as

$$G_{\text{acc}}(k) = \frac{1}{B} \sum_{b=1}^{B} (\bar{X}_{bk} - \bar{X}_k)^2. \tag{13}$$

The motivation for this statistic is as follows. Under the constant model, all accuracies X_i involved in this statistic are obtained from the same distribution, namely, a Bernoulli distribution with probability $w_{C(i)1}/(w_{C(i)1} + w_{(K+1)1})$. Note that the item index $i = 1$ can be chosen in this case since the activation values are the same for all items of the same type under the constant model. Under the all-or-none model, all accuracies are from a Bernoulli distribution with probability $\theta/(\theta + w_{(K+1)1})$. The important thing to note is that in both cases (constant and all-or-none models), these probabilities are the same for all items of rule type k. Hence, under the constant and all-or-none model, the probability of success does not change for items $i > i_k$ of type k.

On the other hand, for the incremental model, let the correct rule for item i be rule k (so $C(i) = k$). Then, item accuracies arise from Bernoulli distributions with the probability of a correct response being

$$\frac{w_{ki}}{w_{ki} + w_{(K+1)i}} = \frac{\alpha^{i-1} w_{k1} + \theta(1-\alpha) \sum_{j=0}^{i-2} \alpha^j I_{\{k\}}(i-1-j)\, X_{i-1-j}}{(\alpha^{i-1} w_{k1} + \theta(1-\alpha) \sum_{j=0}^{i-2} \alpha^j I_{\{k\}}(i-1-j)\, X_{i-1-j}) + \alpha^{i-1} w_{(K+1)1}}, \tag{14}$$

where $I_{\{k\}}(i)$ is an indicator function which is equal to 1 if item i is of type k and is zero otherwise. The equality in Eq. (14) can be derived by applying mathematical induction on Eq. (6). Contrary to the probabilities for the constant and the all-or-none models, the probabilities for the incremental model change after every new item.

The motivation for the statistic in Eq. (13) is now that there will be less variability among the X_i if they all arise from the same distribution (as for the constant and all-or-none models) than if they all arise from different distributions (as for the incremental model). The statistic in Eq. (13) is intended as a measure of the variability of the X_i's. Under the constant and the all-or-none models, this statistic will be smaller than under the incremental model. Although this is difficult to derive analytically, it is easily demonstrated by simulations. Hence, this statistic can be expected to be diagnostic in distinguishing between the constant and all-or-none models, on the one hand, and, on the other hand, the incremental model.

Strictly speaking, we could also use the data before the first success to distinguish between the incremental and the all-or-none models. This is not done here for two reasons. First, the activations w_{ki} of the incremental model change only very weakly

during this period, so there is not a big difference with the all-or-none model in this phase. Second, we are especially interested in the type of learning after a rule is used correctly for the first time; namely, are activations fixed as in the all-or-none model or variable as in the incremental model? Therefore, the data before the first success do not bear much on our main point of interest.

In a similar manner, we introduce the *stationarity response-time statistic*,

$$G_{\text{time}}(k) = \frac{1}{B} \sum_{b=1}^{B} (\bar{T}_{bk} - \bar{T}_k)^2, \tag{15}$$

in which $\bar{T}_{bk}$ and $\bar{T}_k$ are defined in an analogous way as $\bar{X}_{bk}$ and $\bar{X}_k$ are in Eq. (13). Again, this statistic is expected to be smaller under the constant and the all-or-none models than under the incremental model for a similar reason to that outlined in the previous paragraphs.

Results and Discussion

Two participants are excluded from the analysis because they gave up more and more often and faster and faster in the course of test taking. This pattern of data cannot be captured by any of the models, so we decided not to analyze these data. Hence, only the 6 participants who did not give up (or not often) were considered in the following analyses.

Descriptive statistics. Mean accuracies (over items) range (over persons) from .620 to .950. The overall mean accuracy is .887; its standard deviation (SD) over persons equals 0.136. Mean response times (over items) range (over persons) from 4.259 to 25.437 s. The overall mean response time is 15.163 s, and its SD over persons equals 7.138. Mean accuracies per item type are 1, .924, .779, and .821 for addition, Fibonacci, interpolation, and multiplication, respectively. Addition and Fibonacci items seem to be quite easy. In the following analyses, we will focus on interpolation and multiplication items because they allow for learning to show up in both accuracy and response time data.

All minimum response times are within one standard deviation of the mean response time, calculated for every person separately. This illustrates our statement in the context of Eq. (9) that there are no outliers at the low end of the response time scale that can harm the estimation of t_{re}.

Trends. For the accuracy data, we performed $6 \times 4 = 24$ (number of persons $\times$ number of item types) comparisons between performances in the first and in the second block of items. Of these 24 comparisons, seven are in the direction predicted by the learning models (accuracy higher in the second block). Three out of these seven comparisons are statistically significant at the level .05. In 16 comparisons, the accuracies in Blocks 1 and 2 are equal. This is mainly due to people who have null or perfect scores throughout the test for a certain item type (which occurred 13 times out of 16). For one person–item type combination, there was higher accuracy in the first block than in the second.

For the response time data, all 24 comparisons are in the direction predicted by the two learning models (response time is lower in Block 2). Of these, 11 are significant at the .05 level. Hence, these trends in the data, especially in the response time data, seem to speak in favor of the two learning models. The accuracy data may not be appropriate for distinguishing between the two learning models and the constant model because of a very high overall accuracy.

Pearson statistics. Table 1 shows the one-sided p-values for the Pearson (Q) statistics for each person separately, for the interpolation and multiplication items, and for all three models. The last row (denoted "All") reports the p-value if the statistics are added over all persons. The results for the two other item types (addition and Fibonacci) are not shown because they did not distinguish between the three models. Here, as in the following tables, p-values significant at the level .05 are printed in boldface.

First, it should be noted that the data of participant 5 cannot be simulated with the constant model. This is because, in this model, $w_{k1} = 0$ for items of type interpolation and $w_{(K+1)1} = 0$ such that sampling rules for items of this type continue indefinitely. Also, for none of the models could the accuracy statistic be calculated for the interpolation data of this participant since $w_{k1} = 0$ for all three models and $x_i = 0$ for all interpolation items. Hence $w_{ki} = 0$ for all items and $E(\bar{X}_{bk}) = 0$ for all blocks b. This implies that calculating the statistic of Eq. (10) would lead to a division by zero.

We now discuss the responses to the items of type interpolation and multiplication. It is clear that items of type *interpolation* cannot be fitted by the constant model. This holds for both accuracy and response time: People do become faster and more accurate while solving these items. The two learning models (incremental and all-or-none) achieve a reasonable fit for this item type.

The pattern in the data is shown in Fig. 2. In this figure, means are shown for items of interpolation type, separately for each of the four blocks. The expected values are computed according to the incremental model. Each double panel represents the data of one participant (numbered 1, ..., 6). The top and bottom panels depict accuracy and response time data (in seconds), respectively. One can see that one participant (number 2) becomes more accurate and faster during test taking. Two other participants (4 and 6) are at maximal accuracy from the start, but become faster during the test. Participants 1, 3, and 5 do not become more accurate or faster during test taking. Participants 1 and 3 remain at maximal accuracy and Participant 5 at minimal accuracy. These data patterns can be well captured by the incremental or the all-or-none model, but not by the constant model. The plots for the all-or-none model (not shown) look quite similar to those of the incremental model. One exception is that the response time data for Participant 4 could not be fitted well by the all-or-none model. This is also clear from Table 1.

For items of the type *multiplication*, the constant model again gives a bad fit to the accuracy data. The incremental and all-or-none models do fine.

For the response times on the multiplication items, the incremental and constant models give a good fit, but the all-or-none model does not. To investigate the reason for this discrepancy between the three models, we examine the data of the

TABLE 1

Pearson Statistic *p*-Values, Study 1

	Constant model	Incremental model	All-or-none model
Interpolation			
Accuracy			
Participant			
1	1.00	.81	1.00
2	**.00**	.08	.11
3	.47	.50	.78
4	1	1	1.00
5	/	/	/
6	.85	.95	1.00
All	**.00**	.44	.57
Response Time			
Participant			
1	.47	.72	.39
2	**.02**	.16	.10
3	**.01**	.76	.52
4	**.01**	.81	.05
5	/	.89	.91
6	**.02**	.99	.66
All	**.00**	.74	.37
Multiplication			
Accuracy			
Participant			
1	**.00**	.14	.06
2	1.00	1	1.00
3	.69	.49	**.02**
4	**.00**	.14	.18
5	/	.81	.62
6	1.00	1	1.00
All	**.02**	.68	.22
Response Time			
Participant			
1	.55	.44	.10
2	.14	.79	.05
3	.63	.75	**.01**
4	.20	.97	**.00**
5	/	.95	.27
6	**.03**	.15	.15
All	.31	.95	**.01**

worst-fitting participant, person 4. His response times on the multiplication items are plotted in Fig. 3. (Recall from the *Method* section that $B = 2$ for multiplication items.) Plots are shown for the constant, incremental, and all-or-none models. This

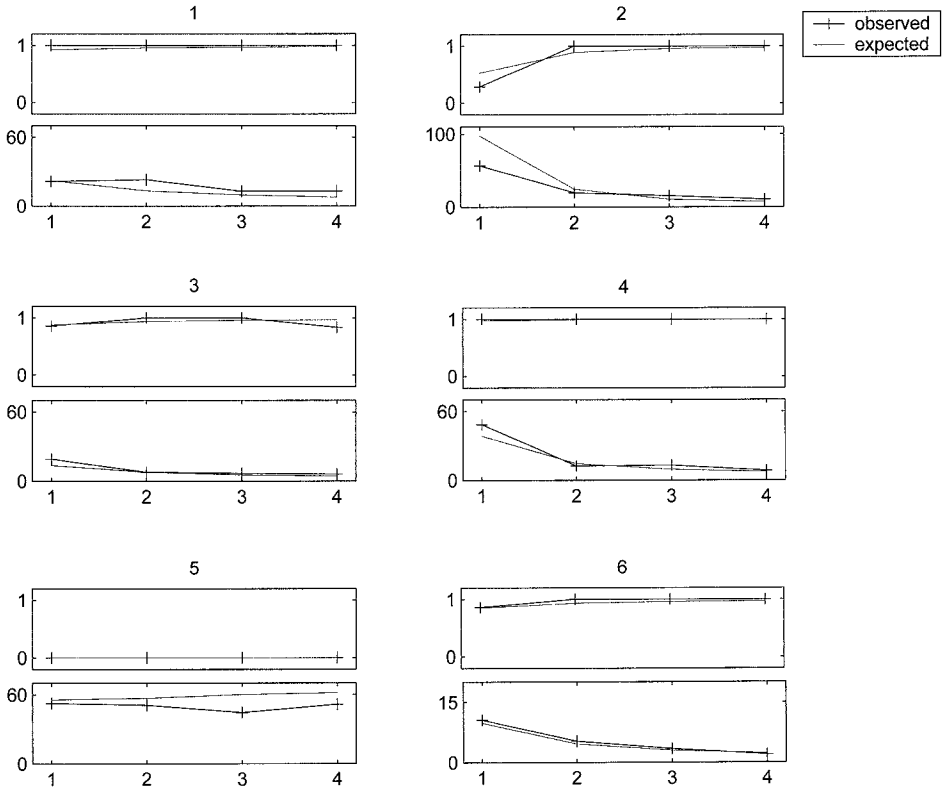


FIG. 2. Observed and expected accuracy measured by proportion correct and response time in seconds for items of type interpolation in Study 1. Each pair of panels corresponds to one participant. Expectations are according to the incremental model. Item blocks appear on the abscissa. Accuracy appears on the ordinate of the top panel of each of the six plots and response time on the ordinate of the bottom panel.

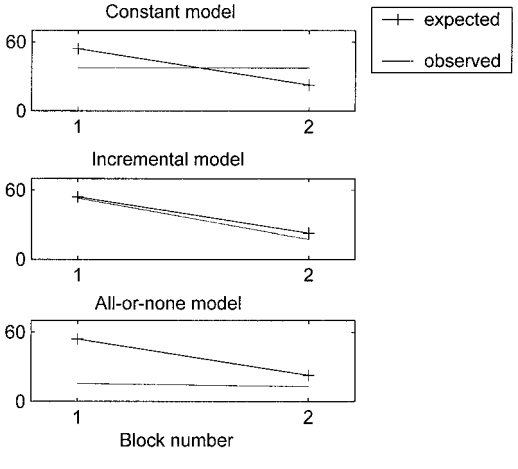


FIG. 3. Observed and expected response times in seconds for multiplication items for one participant in Study 1. Item blocks appear on the abscissa, response times on the ordinate.

participant finds the correct rule very quickly in Block 1, so the all-or-none model predicts that there is no speedup from Block 1 to Block 2. Of course, the constant

TABLE 2

Stationarity Statistic *p*-values, Study 1

	Constant model	Incremental model	All-or-none model
Interpolation			
Accuracy			
Participant			
1	1.00	.43	.38
2	.13	.23	.48
3	.74	.72	.75
4	1.00	1.00	.98
5	/	/	/
6	.34	.22	.34
All	.06	.11	.10
Response Time			
Participant			
1	.83	.71	.36
2	.15	.17	.07
3	.18	.15	.17
4	**.01**	.75	**.00**
5	/	/	/
6	.60	.42	.05
All	.07	.21	**.00**
Multiplication			
Accuracy			
Participant			
1	**.00**	.10	**.02**
2	.58	.88	.90
3	.57	.58	.83
4	**.00**	.29	.11
5	/	.23	.46
6	.59	.59	.55
All	**.03**	.55	.34
Response Time			
Participant			
1	.90	.78	**.01**
2	.93	.53	.49
3	.46	.31	.11
4	.46	.24	.14
5	/	.56	.56
6	.07	.09	.10
All	.41	.16	**.02**

model also predicts no speedup. Only the incremental model can capture the fact that this participant becomes faster after the first correct application of a rule. Another question is why the all-or-none model puts its expected response time so low (see the third plot in Fig. 3), rather than placing it in between the observed response times, just like in the constant model. The reason is that type-multiplication items have to share the same value θ with other item types. Therefore, the value θ, which determines the expected response time, is a compromise between "fast" (addition, Fibonacci) and "slow" (interpolation, multiplication) item types.

Until now, we have found clear evidence against the constant model but the evidence favoring either the incremental or the all-or-none model was less clear. In general, the incremental model was supported better (see Table 1), but not always.

Stationarity statistics. We now consider the stationarity statistics, which were constructed with the aim of distinguishing between the incremental and the all-or-none models. Again, we restrict our discussion to the interpolation and multiplication items. Table 2 reports the *p*-values for the three models. Participant 5 solved no items of interpolation type, so calculating this statistic is not meaningful here.

For the interpolation items, the accuracy data do not distinguish between the three models. However, from the response time data, it is clear that the incremental model does a better job than either the constant or the all-or-none model. This suggests that people still learn *after* the first use of a solution principle.

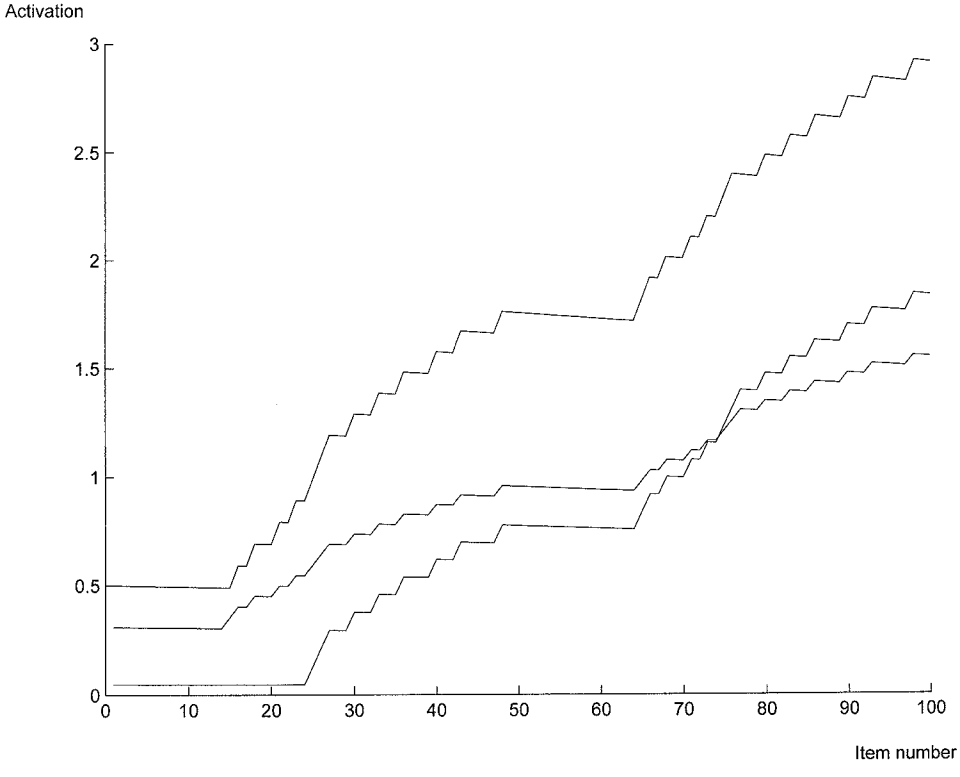


FIG. 4. Activation development in the incremental model of item type interpolation for the first three participants in Study 1.

For the accuracy data of the multiplication items, the constant model again turns out to fit worse than the other two. For the response time data, the all-or-none model fits worst (see the row "All" for multiplication items in Table 2).

Our tentative conclusion from these data is that the incremental model performs better than either of the other two. However, this could be ascribed to two sources of differences between this model and the other two. It may be the case that the incremental model fits better because rule activation develops gradually over time (as we suggest). An alternative explanation, however, is that it is better because it is the only model that allows for forgetting or gradually diminishing activation. However, in the present data, this second explanation is implausible because the estimated parameter α is quite high ($\hat{\alpha} = .998$). Hence, if a certain rule is not used for a while, its activation will diminish very slowly. Figure 4 shows the estimated activation profiles for the first three participants for the rule type interpolation (the profiles for the other three participants are very similar). One can see that indeed if rule activation diminishes, it does so very slowly. It should be noted that Eq. (6) seems to imply that, with a high α value, rule activation also increases very slowly after successfully using a rule. However, this is not the case (as is clear from Fig. 4), because the parameter θ can compensate for a high α value. For the incremental model, the mean estimated θ parameter is 61.844, with a minimum value of 31.684. Hence, our incremental activation account of the data seems warranted.

STUDY 2

Method

The material and data coding are the same as in Study 1. The only difference in the procedure is that only 15 s are allowed to solve an item. If that limit is reached the extended number series is shown as in Study 1, and the participant can proceed to the next item by pressing the space bar.

Participants. Eight persons participated for a small monetary reward. All persons were different from those in Study 1. They were students that had responded to an advertisement.

Trends. Trends in the data are investigated as in the previous study.

Goodness-of-fit statistics. The Pearson Q statistic for the new accuracy variable Y_i is defined in the same way as the statistic for X_i was defined earlier. Analogous statistics for response times are not used in this study. This is because preliminary analyses indicated that the time limit of 15 s is quite severe, and the response-time differences are not diagnostic for distinguishing between the three models. A stationarity statistic is also not used here because it turned out not to be diagnostic in preliminary analyses of these data.

Results and Discussion

Descriptive statistics. Mean accuracies (over items) range (over persons) from .500 to .670. The overall mean accuracy is .568, its standard deviation (SD) over persons equals 0.055. As expected, the mean accuracy is much lower than in Study 1, due to the time limit.

The mean response time is calculated over all response times smaller than 15 s. Response-time means (over items) range (over persons) from 5.240 to 9.240 s. The overall mean response time is equal to 7.246 s; its SD equals 1.286. As expected, both the mean and the standard deviation are lower than in Study 1. The proportions of success are .831, .672, .259, and .473 for, respectively, addition-, Fibonacci-, interpolation-, and multiplication-type items.

All minimum response times are within two standard deviations of the mean response time, calculated for every person separately. Again, this illustrates our statement in the context of Eq. (9), that there are no outliers to disturb the estimation of t_{re}.

Trends. For the accuracy data, $8 \times 4 = 32$ comparisons can be performed. Of these, 21 are in the direction predicted by the learning models (higher accuracy in the second block than in the first). Eight of these 21 comparisons are significant at the level .05. In six of the comparisons, the mean accuracy is equal across blocks, and in five the mean accuracy is higher in the first block than in the second. None of these differences is statistically significant.

For the response time data, the truncated data are included to perform the following calculations. Twenty-seven comparisons out of 32 result in faster response times in the second block than in the first. Of these 27, 12 are statistically significant. Four comparisons are in the inverse direction, and for one comparison the response times are equal in both blocks (equality of mean response time is possible because of the truncation).

We now look at the trend in the proportion of incorrect responses given before the deadline (i.e., the $(K+1)$th response category). In three out of eight comparisons this proportion decreases over blocks. In three comparisons the number is equal, and in two the number increases from the first to the second block. None of these differences is statistically significant. Hence, the proportion of incorrect responses does not decrease strongly.

Pearson statistics. Table 3 shows p-values for the Pearson statistic applied to accuracy data for the three models and for the interpolation and multiplication items. If one looks at the last row (denoted "All") of Table 3 for the accuracy and response-time data, it can be seen that the incremental model has a better fit than the all-or-none model has. This test of the all-or-none model has lower power than the test of the incremental model has (due to two fewer participants (5 and 7) being used in the computation of this statistic for the all-or-none model, for a similar reason as that excluding participant 5 in Study 1). Despite this, Table 3 shows that the incremental model does better than the all-or-none model does. Also, the incremental model is better than the constant model.

As in the previous study, one might wonder whether the better fit of the incremental model is due to its ability to explain forgetting rather than to gradually increasing activation. Again, it turns out that the latter explanation is probably correct, since the estimated α is again quite high ($\alpha = 0.9995$). In case the first part of the learning rule in Eq. (6) applies (for increasing activation), this is compensated by relatively high θ-values (mean estimated $\theta = 188.016$, with a minimum

TABLE 3

Pearson Statistic *p*-values, Study 2

	Constant model	Incremental model	All-or-none model
Accuracy			
Interpolation			
Participant			
1	**.00**	.06	.11
2	**.01**	.06	.26
3	.98	.54	**.02**
4	.11	.33	.34
5	**.00**	.06	/
6	**.01**	**.00**	**.00**
7	1.00	1.00	/
8	**.01**	.27	.33
All	**.00**	**.02**	**.00**
Multiplication			
Participant			
1	.94	.40	**.00**
2	.92	.69	**.03**
3	.71	.78	.05
4	.18	.27	.27
5	1.00	.80	/
6	.70	.34	.64
7	.92	.86	/
8	.83	.84	**.02**
All	.95	.90	**.00**

of 112.073). The activation profile for the different participants is similar to that shown in Fig. 4. There is incremental learning but little if any forgetting. So why do we include the α parameter at all, rather than use a model where there is only incremental learning and no forgetting? The main reason is the interpretability of the θ parameter in the present formulation of the incremental model. Specifically, suppose that beyond a certain item i, all items are solved correctly. In that case, it holds that $\lim_{i \to \infty} \sum_k w_{ki} = \theta$. Hence, θ can then be interpreted as the maximum activation that is possible. No such interpretation is possible for the θ parameter if the α parameter is taken out of the model.

To summarize our findings from both studies, we note that a constant model is inadequate for these data: People do learn while solving this test. Also, the all-or-none learning rule does not seem to account for the data as well as the incremental model does.

GENERAL DISCUSSION

We have developed a general model to account for speed and accuracy data in rule-induction tasks and have shown that this model gives an adequate fit to several

aspects of our data. We have formulated, fitted, and tested three specific models in two contexts (with and without a time limit). The constant model was clearly inadequate, and of the two other models, the incremental model fitted better than the all-or-none model. To end this paper, we discuss two limitations of our model and relate our work to some other research traditions.

Stimulus Information

We have assumed that only the activations of rules that have been built up in the past determine which response is chosen: There is no direct information from the stimulus itself. A more general model might be one in which certain stimulus aspects cooperate in determining a response. Although this is certainly a relevant direction to extend the model, we did not incorporate it here because the stimuli themselves usually lend few hints about the correct rule involved. For example, consider the two series

4 7 11 16 22 29

and

5 6 9 11 13 16.

The first is an example of the addition rule and the second an example of the interpolation rule. The most obvious way to solve these items is to simply try out the different possible rules rather than to look for specific cues that may differentiate between different item types. A point relevant to this argument is that we took care to avoid any cues that might distinguish between different item types. For example, only items of the interpolation type can in principle show a nonincreasing item series (e.g., 4, 7, 5, 10, 6, 13). However, as mentioned earlier, the item series we constructed for these studies were all increasing. Hence, this could not help participants in making a choice between rules. A second example is that the numbers used in multiplication-type items can increase very quickly if large multipliers are used, but we used only the multiplier 2, the smallest one possible.

Other Learning Mechanisms

A second limitation concerns our localization of learning in the activation of solution rules. Other ways to become more proficient in number series performance would be faster encoding of the number series or faster and more accurate rule checking (a process that is ignored in the present model). Our model cannot deal with these kinds of learning. However, in our opinion, in the number series task these processes can probably be performed relatively quickly and accurately in comparison with the problem of finding the correct rule.

Dynamics of Test Behavior

Earlier, we noted the fact that many cognitive task analyses assume that no change takes place during test performance. The untenability of this assumption in cognitive testing is most dramatically shown in participants' behavior on the water jars task (e.g., Dover & Shore, 1991; Luchins, 1942; Shore *et al.*, 1994). The same point was also clearly demonstrated in the present paper by the bad fit of the constant model. The misfit was more subtle with number series than with the water jars

task, so that we had to resort to statistical methods to show this. On the other hand, our test represents a more realistic situation in the sense that it more closely resembles a typical intelligence test.

Existing Learning Models

Some learning models have been proposed in the experimental literature, and we now make a comparison between our work and a few of these.

The incremental vs all-or-none learning controversy. In the fifties and sixties, there was a debate on whether learning is incremental or all-or-none. We have also chosen to call two of our models incremental and all-or-none. However, these models, especially the all-or-none model, differ from the ones presented earlier (e.g., Restle, 1962, 1965). We will discuss the similarities and differences between our models and two well-known earlier learning models, the *linear operator model* and the *one-element stimulus sampling model.*

Consider a paired-associate learning experiment. A set of pairs of words is to be learned, and the first words of the set are presented repeatedly. At each presentation of a word, the participant is required to recall the other word of the pair. After a response, the correct answer is provided by the experimenter. The linear operator model assumes that at each presentation the probability P_n of providing the correct missing word, at the nth presentation of that word, is equal to

$$P_n = (1-\theta)\, P_{n-1} + \theta.$$

This model is similar to our incremental model in that the probability of success increases gradually. However, there are also important differences. First, there is no dependence on previous performance in the linear operator model (i.e., X_{i-1}). Second, the linear operator model operates directly on probabilities, whereas we assume that the probability of success depends on underlying activations that change over time. And third, no response time predictions are derived from the linear operator model, in contrast to ours. Despite the differences, we think it is nevertheless useful to place our incremental model in this tradition, since it shares the basic idea of gradually increasing probabilities of success.

On the other hand, the one-element model (or the all-or-none model) is a model allowing only two states; the "unlearned" and the "learned" states. In the unlearned state, people guess the correct word and have a probability g of guessing the correct one, where g is usually the reciprocal of the number of response options. After every item processed in the unlearned state, there is a fixed probability of jumping to the learned state. Once in the learned state, the participant never leaves that state and solves the item correctly with probability 1. This model differs from our all-or-none model in important ways. First, in our model, the initial state (i.e., if the activation of the rule is w_{k1}) is not a completely unlearned state because there is still some probability $w_{k1}/(w_{(K+1)1}+w_{k1})$ of finding the correct rule (for rule type k). On the other hand, in the one-element model, nothing at all is known in the unlearned state, and successful item performance can occur solely through correct guessing. Second, in our model, in the learned state (i.e., if the activation of a rule is equal to θ) the probability of success is $\theta/(w_{(K+1)1}+\theta)$, rather than 1. And third, in

the one-element model, the event of jumping to the learned state is independent of the performance on the previous item. Despite these differences, we have chosen to call our model an all-or-none model as well, because the two models (our all-or-none model and the one-element model) share the same basic idea that learning is a "discontinuous" phenomenon; that is, a chunk of knowledge is either known or unknown.

Modern learning models. Recently, interest in learning has revived, and many new learning models have been proposed. Many of these new models are of a connectionist type (e.g., Dragoi & Staddon, 1999; Gluck & Bower, 1988; Kruschke, 1992). A related but nonconnectionist model is the dynamic version of the fuzzy logic model (Friedman, Massaro, Kitzis, & Cohen, 1995). Friedman *et al.* (1995) and Kitzis, Kelley, Berg, Massaro, and Friedman (1998) present theoretical and empirical comparisons between several of these models. None of these recent models adhere to all-or-none learning rules. Instead, they all use some type of incremental learning rule. For example, one often-used rule is the *backpropagation* learning rule in neural networks (e.g., Ballard, 1995). In backpropagation learning, the learner is given the correct answer, and weights in the network are optimally adapted such as to minimize the difference between the correct and the stated answer. However, the empirical validity of these learning rules is seldom tested explicitly. Methods proposed in this paper may prove to be useful for making detailed comparisons between different learning rules.

APPENDIX A

Derivations

In this appendix, we derive some observations stated in the main text. The chronology of the text is followed. The item index i will be dropped whenever this does not lead to confusion.

Suppose $K+1$ responses perform a race, in which each response races with a time that is independent and exponentially distributed with parameter (hazard rate) w_k. Then, it is an elementary exercise in probability theory to show that the race time and the choice are statistically independent in this case (see Ross, 2000, p. 295, Exercise 7). Further, it also follows from that exercise that the probability of the kth rule winning the race equals $w_k/[\sum_{k=1}^{K+1} w_k]$.

Using the definition $w_{\text{tot}} = \sum_{k=1}^{K+1} w_k$, we recall that $\pi_C = w_C/w_{\text{tot}}$, as in the main text, and we introduce the notation $\pi_{K+1} = w_{K+1}/w_{\text{tot}}$ and $\pi_{\text{stop}} = \pi_C + \pi_{K+1} = (w_C + w_{K+1})/w_{\text{tot}}$. The variable w_{tot} denotes the total activation, π_C is the probability of sampling the correct rule, π_{K+1} is the probability of sampling the misleading rule, and π_{stop} is the probability that the sampling process stops (i.e., the probability of sampling the correct or the misleading rule).

We first discuss the time-unlimited case. In the main text, it is already noted that the probability of success equals

$$\Pr(X=1) = \frac{w_C}{w_C + w_{K+1}}, \tag{A1}$$

where C denotes the number of the correct rule. Our first result concerns the distribution of the number of sampled rules in a given item, denoted J (with realizations j). It turns out that J has a geometric distribution with parameter π_{stop}.

RESULT 1. $\Pr(J = j \mid W = x) = \Pr(J = j) = \pi_{\text{stop}}(1-\pi_{\text{stop}})^{j-1}$.

Proof.

$$\begin{aligned}\Pr(J = j \mid X = x) &= \frac{\Pr(J = j, X = x)}{\Pr(X = x)} \\ &= \frac{(1-\pi_{\text{stop}})^{j-1}\,\pi_C^x \pi_{K+1}^{1-x}}{\dfrac{w_C^x w_{K+1}^{1-x}}{w_C + w_{K+1}}} \\ &= (1-\pi_{\text{stop}})^{j-1}\,\pi_{\text{stop}}. \quad \blacksquare\end{aligned}$$

The second equality of the proof follows from applying Eq. (A1) to $\Pr(X = x)$.

The next result states that the total decision time t for a given item, conditional on the number of rules that are sampled (j), is Gamma distributed with parameters j and w_{tot}.

RESULT 2. $p(T = t \mid J = j, X = x) = p(T = t \mid J = j) = \text{Gamma}(j, w_{\text{tot}})$.

Proof. The variable T is the sum of j component times. Every component time is the winning time of a race between $K+1$ independent exponential random variables. From Ross (2000, p. 295, Exercise 7), it follows that the distribution of each component time is independent of x. Therefore, T (the sum of the component times) is also independent of x. This proves the first equality. Now, since every component time is distributed as an exponential with parameter w_{tot}, it follows that T is distributed as a Gamma with parameters j and w_{tot}. This proves the second equality. $\blacksquare$

The next result shows that the decision time follows an exponential distribution with hazard rate $(w_C + w_{K+1})$. We introduce the moment-generating function of T, denoted by $M_T(s)$. The idea of the proof is due to McGill (1963).

RESULT 3. $p(T = t \mid X = x) = p(T = t) = (w_C + w_{K+1}) \exp[-(w_C + w_{K-1})\, t]$.

Proof.

$$\begin{aligned}M_{T \mid X = x}(s) &= \sum_{j=1}^{\infty} \Pr(J = j)\, M_{T \mid X = x, J = j}(s) \\ &= \sum_{j=1}^{\infty} (1-\pi_{\text{stop}})^{j-1}\,\pi_{\text{stop}} \left(\frac{w_{\text{tot}}}{w_{\text{tot}} - s}\right)^j \\ &= \frac{w_C + w_{K+1}}{w_C + w_{K+1} - s}.\end{aligned}$$

The second equality of the proof follows from Result 2. Hence, $p(T \mid X = x)$ is an exponential distribution with parameter $(w_C + w_{K+1})$. $\blacksquare$

The following two results are referred to in the subsection "Time Limits" of the main text.

RESULT 4. $\Pr(Y=1) = (w_C/(w_C+w_{K+1}))\,\{1-\exp(-(w_C+w_{K+1})\,\tau)\}$.

Proof.

$$\begin{aligned}\Pr(Y=1) &= \Pr(X=1 \,\&\, T<\tau)\\ &= \Pr(X=1 \mid T<\tau)\,\Pr(T<\tau)\\ &= \frac{w_C}{w_C+w_{K+1}}\,\{1-\exp(-((w_C+w_{K+1})\,\tau)\}. \quad \blacksquare\end{aligned}$$

We now derive the PDF of accuracies and response times to all I items jointly for the case of time-limited item presentation. We first define, for item i,

$$u_i = \begin{cases} t_i & \text{if} \quad t_i < \tau \\ \tau & \text{otherwise.}\end{cases}$$

Then, we can state and prove the following result.

RESULT 5. $p(\text{data}) = \prod_i [w_{C(i)\,i}^{x_i} w_{(K+1)\,i}^{1-x_i}]^{I_{(0,\tau)}(u_i)} \exp[-(w_{C(i)\,i}+w_{(K+1)\,i})\,u_i]$.

The symbol I denotes an indicator function, where $I_{(0,\tau)}(u)=1$ if $0<u<\tau$ and $I_{(0,\tau)}(u)=0$ otherwise.

Proof. With time-limited presentation, the variables T_i and X_i are sometimes not observed, namely, when $T_i \geqslant \tau$. The variables T_i with $T_i \geqslant \tau$ are said to be truncated. We define the variable O_i which indicates whether T_i and X_i are observed ($O_i=1$ if $T_i<\tau$) or not ($O_i=0$ if $T_i \geqslant \tau$). We denote the set of items with $t_i<\tau$ as D_1 and the items with $t_i \geqslant \tau$ as D_0. The probability density of all observations, including the observations where $T_i \geqslant \tau$, can be written as

$$\begin{aligned}p(\text{data}) &= \left[\prod_{i=1}^{I} \Pr(O_i=o_i)\right] \times \left[\prod_{i\in D_1} p(X_i=x_i, T_i=t_i \mid O_i=1)\right]\\ &= \left[\prod_{i\in D_0} [1-F_i(\tau)]\right] \times \left[\prod_{i\in D_1} F_i(\tau)\right] \times \left[\prod_{i\in D_1} \frac{p(X_i=x_i, T_i=t_i, O_i=1)}{\Pr(O_i=1)}\right],\end{aligned} \tag{A2}$$

where the symbol $F_i(\,)$ denotes the cumulative distribution of the response time for item i. We can rewrite (A2) as

$$\begin{aligned}p(\text{data}) &= \prod_{i\in D_0} [1-F_i(\tau)] \prod_{i\in D_1} p(X_i=x_i, T_i=t_i)\\ &= \prod_{i\in D_0} \exp[-(w_{C(i)\,i}+w_{(K+1)\,i})\,\tau]\\ &\quad \times \prod_{i\in D_1} (w_{C(i)\,i})^{x_i}\,(w_{(K+1)\,i})^{1-x_i} \exp[-(w_{C(i)\,i}+w_{(K+1)\,i})\,t_i]\\ &= \prod_i [(w_{C(i)\,i})^{x_i}\,(w_{(K+1)\,i})^{1-x_i}]^{I_{(0,\tau)}(u_i)} \exp[-(w_{C(i)\,i}+w_{(K+1)\,i})\,u_i]. \quad \blacksquare\end{aligned}$$

Note that x_i is not observed for items with $t_i \geqslant \tau$. However, in this case x_i does not appear in the PDF because $I_{(0,\tau)}(u_i)=0$ for these items.

APPENDIX B

Protected Steepest Ascent: A Flexible Algorithm for Optimization

The method combines three aspects that are found in the numerical optimization literature: steepest ascent, protection, and finite differences. It requires that the objective function has finite first derivatives.

Steepest Ascent

We start by describing the method of *steepest ascent*, which is the basis of the present method, and some notation will be introduced simultaneously. Suppose a function L is to be optimized as a function of an H-dimensional parameter ξ. Denote the gradient vector with elements $\partial L/\partial \xi_h (h = 1, \ldots, H)$ by ∇. Then, in step n of the algorithm, the parameter vector $\xi^{(n)}$ is obtained by

$$\xi^{(n)} = \xi^{(n-1)} + a\nabla,$$

where a is a scalar that determines the size of the step taken in direction ∇ (Gill, Murray, & Wright, 1995).

Protected Steepest Ascent

The parameter a is chosen in some suitable way. We choose $a = 1$ and then check whether the function L is increased in the new point $(\xi + \nabla)$. If so, this new point is chosen as the new starting point; if not, a is divided by 2 and the function value is checked in the resulting point (i.e., $\xi + \nabla/2$). If the function is still not increased in comparison with $L(\xi)$, a is again divided by 2, and so on, until the function $L(\xi + \nabla/2^m)$ is increased for some value of m. Since the direction ∇ is a direction of increase, such a value m will always be found. The resulting algorithm will be called the *protected* steepest ascent, since the parameter jumps are protected from overshooting. The chief difficulty in the previous method is that the gradient vector ∇ has to be calculated. This vector is approximated with the method of *finite differences* (Gill *et al.* 1995). One should note that this yields a method that requires only calculation of the likelihood function L.

Boundary Restrictions

One advantage of the method is that it can easily handle boundary restrictions. We distinguish three cases.

1. The parameter ξ is restricted to lie within the (finite) interval (A, B). In this case, whenever the parameter ξ appears in the likelihood function, exchange it with the expression

$$(B - A) \times \frac{\exp(\xi^*)}{1 + \exp(\xi^*)} + A \tag{B1}$$

and maximize it with respect to ξ^*, which is an unrestricted parameter. The algorithm results in an estimate of ξ^*. By applying expression (B1) the corresponding estimate for ξ is obtained. In the context of maximum likelihood (ML) estimation, the ML invariance property (Mood, Graybill, & Boes, 1974, p. 284) guarantees that this results in the correct estimate for ξ.

2. The parameter ξ is restricted to lie in the interval $(-\infty, B)$. Exchange the parameter ξ in L with B-exp$(-\xi^*)$ and proceed as in Case 1.

3. The parameter ξ is restricted to lie in the interval $(A, +\infty)$. Exchange the parameter ξ in L with $A+\exp(\xi^*)$ and proceed as in Case 1.

We now turn to an example of the protected steepest ascent algorithm which is relevant in the present paper.

EXAMPLE 1. The likelihood function we will consider is of the form

$$L = \prod_{i=1}^{I} (w_{C(i)\,i})^{x_i} (w_{(K+1)\,i})^{1-x_i} \exp\left[-((w_{C(i)\,i} + w_{(K+1)\,i})\, t_i\right], \tag{B2}$$

which is explained in the main body of the text. Under the incremental model, all activations w_{ki} change according to the equation

$$w_{ki} = \begin{cases} \alpha w_{k(i-1)} + (1-\alpha)\,\theta & \text{if item } i-1 \text{ is of type } k \text{ and } X_{i-1} = 1, \\ \alpha w_{k(i-1)} & \text{otherwise,} \end{cases} \tag{B3}$$

which is also explained in the main body of the text. Our objective is to maximize (B2) under the constraint (B3) on the activations w_{ki}. If K responses and P persons are involved, there are $(K+2)\,P+1$ parameters: $(K+1)\,P$ person-specific initial activations $w_{k1}(k = 1, \ldots, K+1)$, P person-specific parameters θ, and one common parameter α.

To study the performance of our method of parameter estimation, we consider data from six simulated participants and $K = 4$ (as in Study 1 in the main text). This results in 37 parameters to be estimated. We also used the following restrictions: $w_{k1} > 0$ (for all k), $\theta > 0$, and $0 < \alpha < 1$. The resulting function L is quite complicated, but can easily be maximized with the protected steepest ascent method.

A total of 100 items is presented to these pseudo-participants, and parameters are estimated based on their response time and accuracy data. The convergence criterion is that all gradient elements have to be smaller than $\varepsilon = 0.01$ in absolute value. Starting the estimation from different starting points always results in the same maximum $\hat{\xi}$, which suggests that the maximum is unique.

Ten datasets are generated. In Table B1, we show the true parameter values (ξ), the mean estimated parameter values ($\bar{\hat{\xi}}$), and the standard error of this parameter (SE($\hat{\xi}$)). We show these results separately for the best and worst fitting simulated participants. (*Best* and *worst* are defined in terms of the average absolute difference between true and estimated parameter values.) Table B1 shows that recovery is quite good and that the standard error is low, especially for the parameters α and w_{k1}.

TABLE B1

Parameter Estimates for $I = 100$

	Best-fitting participant			Worst-fitting participant		
Parameter	ξ	$\bar{\hat{\xi}}$	SE($\hat{\xi}$)	ξ	$\bar{\hat{\xi}}$	SE($\hat{\xi}$)
α	0.990	0.991	0.002	0.990	0.991	0.002
θ	1.000	0.908	0.379	5.000	5.513	0.968
$w_{1,1}$	0.100	0.124	0.056	0.100	0.130	0.078
$w_{2,1}$	0.050	0.063	0.025	0.050	0.060	0.046
$w_{3,1}$	0.060	0.080	0.045	0.060	0.094	0.059
$w_{4,1}$	0.020	0.019	0.016	0.020	0.059	0.073
w_{K+1}	0.100	0.089	0.021	0.100	0.102	0.030

REFERENCES

Anderson, J. R. (1983). *The architecture of cognition.* Hillsdale, NJ: Erlbaum.

Anderson, J. R. (1993). *Rules of the mind.* Hillsdale, NJ: Erlbaum.

Anderson, J. R. (1996). ACT: A simple theory of complex cognition. *American Psychologist*, **51**, 355–365.

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

Ballard, D. H. (1997). *An introduction to natural computation.* Cambridge, MA: MIT Press.

Bush, R. R., & Mosteller, F. (1955). *Stochastic models for learning*. New York: Wiley.

Davison, A. C., & Hinkley, D. V. (1997). *Bootstrap methods and their applications.* Cambridge, MA: Cambridge University Press.

Dover, A., & Shore, B. M. (1991). Giftedness and flexibility on a mathematical set-breaking task. *Gifted Child Quarterly*, **35**, 99–124.

Dragoi, V., & Staddon, J. E. R. (1999). The dynamics of operant conditioning. *Psychological Review*, **106**, 20–61.

Duncker, K. (1945). On problem solving. *Psychological Monographs*, **58**, No. 5.

Friedman, D., Massaro, D. W., Kitzis, S. N., & Cohen, M. M. (1995). A comparison of learning models. *Journal of Mathematical Psychology*, **39**, 164–178.

Gill, P. E., Murray, W., & Wright, M. H. (1995). *Practical optimization.* London: Academic Press.

Gluck, M. A., & Bower, G. H. (1988). From conditioning to category learning: An adaptive network model. *Journal of Experimental Psychology: General*, **117**, 227–247.

Holzman, T. G., Pellegrino, J. W., & Glaser, R. (1982). Cognitive dimensions of numerical rule induction. *Journal of Educational Psychology*, **74**, 360–373.

Holzman, T. G., Pellegrino, J. W., & Glaser, R. (1983). Cognitive variables in series completion. *Journal of Educational Psychology*, **75**, 603–618.

Kaplan, I. T., & Schoenfeld, W. N. (1966). Oculomotor patterns during the solution of visually displayed anagrams. *Journal of Experimental Psychology*, **72**, 447–451.

Kitzis, S. N., Kelley, H., Berg, E., Massaro, D. W., & Friedman, D. (1998). Broadening the tests of learning models. *Journal of Mathematical Psychology*, **42**, 327–355.

Korossy, K. (1998). Solvability and uniqueness of linear-recursive number sequence tasks. *Methods of Psychological Research*, **3** [www.mpr-online.de].

Kotovsky, K., & Simon, H. A. (1973). Empirical tests of a theory of human acquisition of concepts for sequential patterns. *Cognitive Psychology*, **4**, 399–424.

Kruschke, J. K. (1992). ALCOVE: An exemplar-based connectionist model of category learning. *Psychological Review*, **99**, 22–44.

Lemay, E. H. (1972). Anagram solutions as a function of task variables and solution word models. *Journal of Experimental Psychology*, **92**, 65–68.

Lovett, M. C., & Anderson, J. R. (1996). History of success and current context in problem solving. *Cognitive Psychology*, **31**, 168–217.

Lovett, M. C., & Shunn, C. D. (1999). *Journal of Experimental Psychology: General*, **128**, 107–130.

Luce, R. D. (1959). *Individual choice behavior*. New York: Wiley.

Luchins, A. S. (1942). Mechanization in problem solving—The effect of Einstellung. *Psychological Monographs*, **54**.

Maier, N. R., & Burke, R. J. (1966). Test of the concept of "availability of functions" in problem solving. *Psychological Reports*, **19**, 119–125.

Marley, A. A. J. (1989). A random utility family that includes many of the classical models and has closed form choice probabilities and choice reaction times. *British Journal of Mathematical and Statistical Psychology*, **42**, 13–36.

McGill, W. J. (1963). Stochastic latency mechanisms. In R. D. Luce, R. R. Bush, and E. Galanter (Eds.), *Handbook of mathematical psychology* (Vol. 1, pp. 309–360). New York: Wiley.

Mood, A. M., Graybill, F. A., and Boes, D. C. (1974). *Introduction to the theory of statistics*. Singapore: McGraw–Hill.

Restle, F. (1962). The selection of strategies in cue learning. *Psychological Review*, **69**, 329–343.

Restle, F. (1965). Significance of all-or-none learning. *Psychological Bulletin*, **64**, 313–325.

Ross, S. M. (1996). *Stochastic processes*. New York: Wiley.

Ross, S. M. (2000). *Introduction to probability models*. San Diego: Academic Press.

Schrepp, M. (1995). Modeling interindividual differences in solving letter series completion problems. *Zeitschrift für Psychologie*, **203**, 173–188.

Shore, B. M., Koller, M., and Dover, A. (1994). More from the water jars: A reanalysis of problem-solving performance among gifted and nongifted children. *Gifted Child Quarterly*, **38**, 179–183.

Simon, H. A., and Kotovsky, K. (1963). Human acquisition of concepts for sequential patterns. *Psychological Review*, **70**, 534–546.

Suppes, P., and Ginsberg, R. (1963). A fundamental property of all-or-none models, binomial distribution of responses prior to conditioning, with application to concept formation in children. *Psychological Review*, **70**, 139–161.

White, H. (1988). Semantic priming of anagram solutions. *American Journal of Psychology*, **101**, 383–399.

Wiley, J. (1998). Expertise as mental set: The effects of domain knowledge in creative problem solving. *Memory and Cognition*, **26**, 716–730.

Received: May 23, 2000; published online: January 10, 2002