Journal of Mathematical Psychology **42**, 327–355 (1998)
Article No. MP981215

Broadening the Tests of Learning Models

Stephen N. Kitzis

Fort Hays State University,

and

Hugh Kelley, Eric Berg, Dominic W. Massaro, and Daniel Friedman*

University of California, Santa Cruz

For many years psychological studies of the learning process have used a simulated medical diagnosis task in which symptom configurations are probabilistically related to diseases. Participants are given a set of symptoms and asked to indicate which disease is present, and feedback is given on each trial. We enrich this standard laboratory task in four different ways. First, the symptoms have four possible values (low, medium low, medium high, and high) rather than just two. Second, symptom configurations are generated from an expanded factorial design rather than a simple factorial design. Third, subjects are asked to make a continuous judgment indicating their confidence in the diagnosis, rather than simply a binary judgment. Fourth, cumulated performance scores, payoffs, and the availability of a historical summary of the outcomes are varied in order to assess how these treatments modulate performance. These enrichments provide a broader data set and more challenging tests of the models.

Using 123 subjects each in 480 trials, we compare five existing learning models plus several variants, including the well-known Bayesian, fuzzy logic, connectionist, exemplar, and ALCOVE models. We find that the subjects do learn to distinguish the symptom configurations, that subjects are quite heterogeneous in their response to the task, and that only a small part of the variation across subjects arises from the differences in treatments. The most striking finding is that the model that best predicts subjects' behavior is a simple Bayesian model with a single fitted parameter for prior precision to capture individual differences. We use

The research was supported in part by a grant from the National Science Foundation (SBR-9310347) to Friedman and Massaro. We have benefitted from the comments of special issue editor Gary Bolton and two anonymous referees.

* Send correspondence to Dan Friedman, Department of Economics, University of California, Santa Cruz, CA 95064; email: dan@cash.ucsc.edu.

0022-2496/98 $25.00

rolling regression techniques to elucidate the behavior of this model over time and find some evidence of over-response to current stimuli.

INTRODUCTION

How humans learn persists as one of the most challenging puzzles facing social scientists. Some eminent scientists such as Chomsky and Fodor argue that the laws underlying learning are highly domain specific. How we learn language may not resemble how we learn mathematics. Altogether different processes may govern learning environmental contingencies such as the likelihood of morning fog in Santa Cruz given morning fog the previous day, or learning strategic interactions such as an entrepeneurs' decisions on how to position their products. But recently the domain specificity even of language learning has been empirically challenged (e.g., Pullum, 1996).

We conjecture that domain independent laws will be found for learning processes. Evidence for this conjecture comes from domain independent laws in other fields of inquiry. Important examples include Weber's law, Gestalt laws of organization in both visual and auditory modalities, and Shepard's (1987) law of generalization (Massaro, in press). Within the field of learning, regularities across domains can be described by the law of effect (or time on task), massed versus spaced practice, and the power law of learning. These successes encourage attempts to uncover trial-to-trial learning and decision rules that will apply across a broad range of information domains.

Our conjecture might appear overly optimistic and naive without making a crucial distinction between information and information processing. Information refers to what the stimulus input means to the perceiver. Different domains will necessarily involve different stimuli and these different stimuli make available different content and experience to different individual perceivers. Information processing, on the other hand, refers to the general learning and decision operations postulated by a model, and it could be identical across domains. It is, of course, an empirical question whether domains differ with respect to just information or also differ with respect to information processing.

In this paper we offer some empirical evidence. We collect five of the most prominent existing models of the learning process, adapt them to a common task, and compare their ability to predict the choices over 480 trials of each of our 123 subjects. The models include Bayesian statistical decision (the normative or optimal model) and several descriptive models including the fuzzy logic model of perception (FLMP), a neural network or connectionist model (CMP), an exemplar model, and a simple version of the ALCOVE model. We also consider several variants including a simple reinforcement model in the style of Roth and Erev (1995).

The task, an expanded version of medical diagnosis, is chosen with three conflicting goals in mind. First, to give all the models a fair chance, we want the task to resemble the simple tasks for which the models were developed and that have been used in previous empirical literature on learning. Second, we want the task to challenge subjects' learning capacities so we can observe the process over many

periods. Third, we want a varied task that resembles interesting domains outside the lab so we can begin to address the issue of domain specificity.

Our medical diagnosis task, like its predecessors in the literature, tests isolated individual subjects in many trials. In each trial, the subject sees a set of symptoms, is asked to indicate which disease is present, and then is told the actual disease. The object of learning is the stochastic relation between symptoms and diseases. Thus the task resembles important and nontrivial domains outside the laboratory. We enrich the standard laboratory task in four different ways. First, the symptoms have four possible values (low, medium low, medium high, and high) rather than just two. Second, we present each symptom level separately as well as each combination of symptom levels. Third, subjects are asked to make a continuous judgment indicating their confidence in the diagnosis, rather than simply a binary judgment. Fourth, cumulated performance scores, payoffs, and access to an historical summary of the outcomes are varied in order to assess how these treatments modulate performance. We believe that the enrichments provide subjects with a suitable learning challenge and provide researchers with a broader and sharper data set for evaluating the competing models.

Continuous judgements are the focus of our data analysis because they allow us to observe trial-by-trial learning for individual subjects. The individual trial-to-trial data in most previous research consists of binary choices which must be averaged across subjects to become comparable to the probabilistic predictions of learning models. Such data cannot account for individual differences in information (as defined above) and thus cannot address the key question of whether individual subjects also differ in information processing. Thus our data allow us to investigate whether some subjects are better described by one model (say, CMP) and other subjects by another learning model (say, exemplar).

Another methodological point worth mentioning is that our data analysis relies mainly on least squares fits to the data and compares squared forecast error across competing models. Moreover we use the quadratic scoring rule in some treatments to rate subjects' continuous judgements. Selten (1996) offers an axiomatic justification of these conventions and points out some shortcomings of the currently more popular maximum likelihood techniques in the present context of probabilistic prediction.

We find that subjects do learn to distinguish the symptom configurations, that subjects are quite heterogeneous in their response to the task, and that only a small part of the variation across subjects arises from the differences in access to history and score. The most striking finding is that the model that best predicts subjects' behavior is a simple Bayesian model with a single fitted parameter for prior precision to capture individual differences. We use recent rolling regression techniques to elucidate the behavior of this model over time and find some evidence of overresponse to the current symptoms.

The work presented here builds on three related papers. Friedman, Massaro, Kitzis and Cohen (1995, henceforth denoted FMKC95) tests competing learning models in a simple medical diagnosis task with binary choice only. The concluding discussion explains the need for enhanced experiments with continuous choices and fits of models to individual subjects. Friedman and Massaro (1997, henceforth

denoted FM97) describes the current data set in more detail and tests competing models of decision stage, but does not examine learning behavior or the other stages of information processing. Kelley and Friedman (1998, henceforth KF98) studies a related learning task (called orange juice futures prediction) and introduces the rolling regression techniques we use in the last part of the results section. We also build on many published papers by other investigators which we cite when we present the medical diagnosis task and the models.

Since the models are widely known but the task has novel elements, we begin in the next section by describing the experiment. Then we adapt the competing learning models to our task and present the results.

EXPERIMENT

Gluck and Bower (1988) and several other authors surveyed in FMKC95 investigate a medical diagnosis task in which each subject is presented in each trial with a "medical chart" $s = s_1\, s_2 \cdots s_n$ of up to $n = 4$ binary symptoms. For example, $s_1 = 1$ might indicate a sore throat and $s_3 = 0$ might indicate the absence of dizzy spells. The subject responds to each chart (or symptom configuration) by stating which of two diseases a patient has. Then the subject is told the actual disease.

The medical diagnosis task is very simple in that all stimuli s_i and all responses d are binary variables, coded here as 0 or 1. But with one or more symptoms presented the task is nontrivial for two reasons. First, all trials are training trials in that the subject always receives feedback (the actual value of d) and the subject has no initial knowledge of the relation between d and the given s. Second, the relationship between the symptoms and the disease is stochastic, so the subject sometimes will be misled. For example, the subject may pick the less likely disease and turn out to be correct on that trial, or pick the more likely disease and turn out to be incorrect.

The current research extends the previous experimental task in four different ways. First, our two symptoms can assume any of four values (low, medium low, medium high, and high), rather than just two values. Second, stimulus combinations are generated by an expanded factorial design rather than a simple factorial design. As will be explained below, we have 24 symptom configurations, rather than 16. Third, subjects are asked to make a continuous judgment indicating their confidence in the diagnosis, rather than simply a binary (disease A or disease B) judgment. Fourth, we alter the environment by varying the Score and History treatments described below. These elaborations of the previous procedure provide a richer data set as well as much more challenging tests of the models.

Subjects

A total of 123 undergraduates from the University of California at Santa Cruz participated in this experiment as an option to fulfill a class requirement. One third of these subjects also received pay for their participation—two of the six treatment

cells in the experiment involved the distribution of pay based on individual performance. Individual earnings during the experiment ranged from 5 to 17 dollars, with a mean of $13.63. The entire experimental procedure was approximately two hours in length. Written instructions are available on request.

Apparatus

Subjects were placed in sound dampened isolated testing rooms and responded to a computer program written in C++ run on Power Macintosh 7500/100 computers with full color monitors. Subjects viewed stimuli on the monitor and responded by clicking the mouse on various icons in the display. Figure 1 provides examples of the monitor displays.

Procedure

Participants were told that they would be diagnosing a series of fictitious patients based on medical charts, which consisted of values on the symptoms of temperature and blood pressure (sometimes with one symptom missing). They were told that this was a learning experiment in which the goal was to learn the associations between symptoms and diseases. It was explained that the experiment reflected the real-life fact that the relationship between particular symptoms and a specific disease is sometimes weak or strong but never entirely certain. Subjects were informed that the experiment was designed to be difficult (especially in the beginning), that they could expect to make mistakes, and that the presentation order of the diseases was random.

Stimuli

Each trial began with the presentation of the symptom values on the left side of the monitor display. As can be seen in Fig. 1A, the symptom values were displayed using two thermometer icons labeled temperature and blood pressure. Each thermometer was partially filled in red to indicate the symptom level: low, medium low, medium high, or high. For example, a 1/4 filled thermometer represented low, and a 3/4 filled thermometer represented medium-high temperature or blood pressure. Also present on the monitor were icons representing the two possible diseases: Autochus and Burlosis. In some conditions, subjects had access at this point to previous case histories, as explained below. The subject then made a binary choice of which disease he/she believed that the current patient possessed by clicking the mouse on the appropriate icon. A continuous response was then collected in the form of a confidence rating entered by using the mouse to move a slide bar as in Fig. 1C. After collecting both responses the program revealed which of the two diseases the patient turned out to have by highlighting the appropriate icon, e.g., Burlosis in Fig. 1C. In some conditions a score (explained below) was presented at this point. When ready, the subject would advance to the next trial via another mouse click.

Expanded Factorial Design

Table 1 presents the expanded design used in our experiment. As in a single factor design, each of the symptoms is presented unimodally, for a total of $4 + 4 = 8$ symptoms. As in the factorial design, each of the four temperature symptoms is

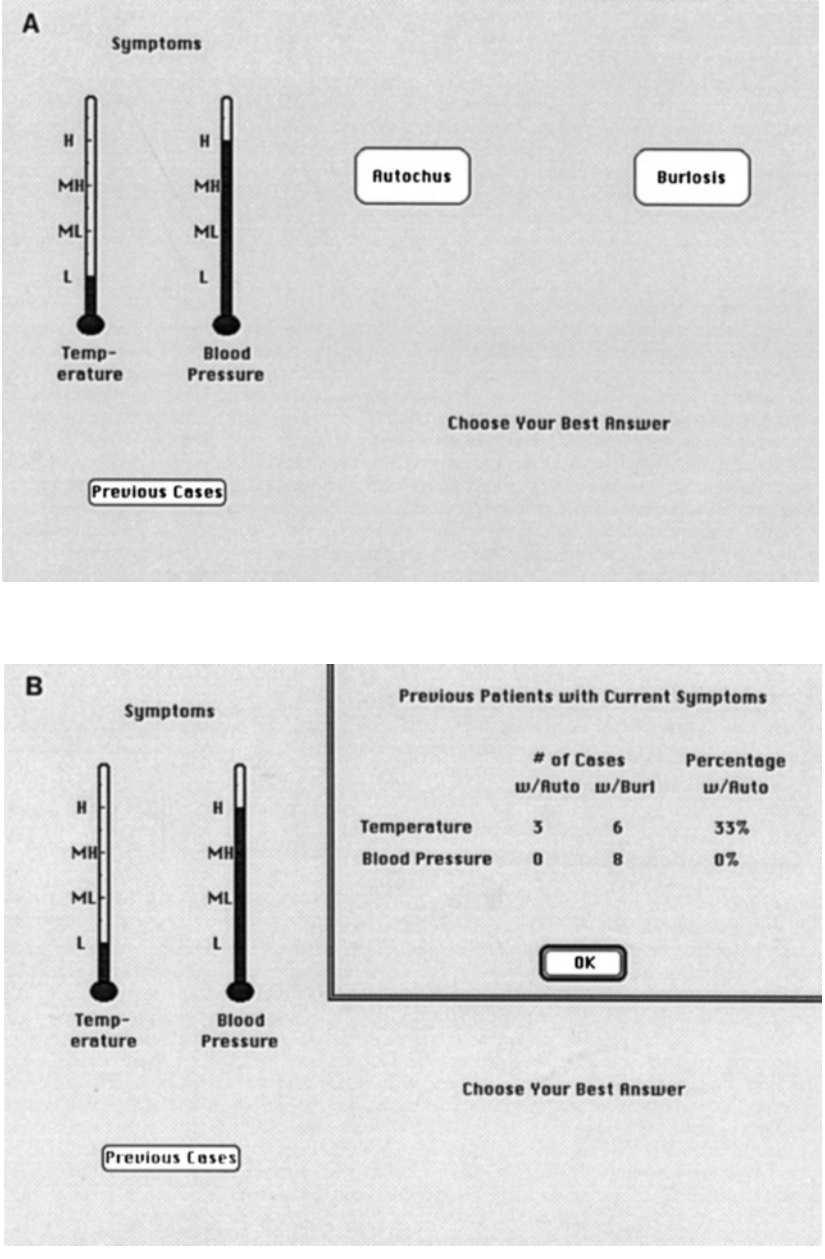

FIG. 1. Medical diagnosis screens: (A) basic screen; (B) history screen; (C) score screen.

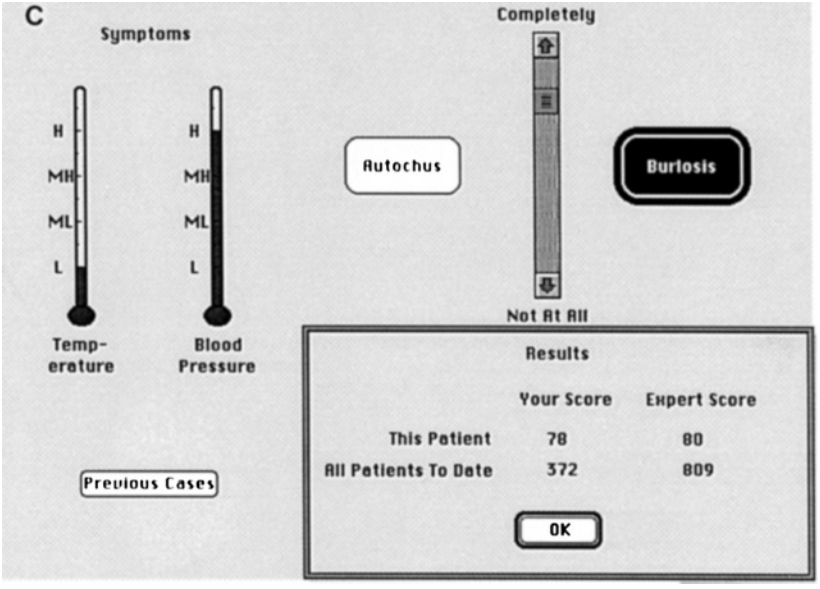

FIG. 1—*Continued*

combined with each of the four blood pressure symptoms for another 16 symptom configurations. Thus there are a total of 24 configurations.

The experimental session for each subject consisted of 480 trials. Table 1 shows that the number of observations for the different symptom configurations ranged from 11 to 33, with 18 of the 24 having between 16 and 26 observations. These frequencies were determined by the conditionally independent symptom likelihoods listed in Table 2. The same randomized presentation order of stimuli was used for

TABLE 1

Expanded Factorial Design: Instances of Disease (*A*, *B*) by Symptom Configuration

	Blood Pressure				
Temperature	High	Medium high	Medium low	Low	None
High	3,8	10,5	19,2	32,1	16,4
Medium high	2,16	7,10	14,5	24,2	12,8
Medium low	2,24	5,14	10,7	16,2	8,12
Low	1,32	2,19	5,10	8,3	4,16
None	2,20	6,12	12,6	20,2	

Note. The first entry in the top row, for example, means that there were 3 cases of Autochus (disease *A*) and 8 cases of Burlosis (disease *B*) in medical charts showing High Temperature and High Blood Pressure; the last entry in the top row indicates 20 charts (16*A* and 4*B*) with no blood pressure reported and High Temperature.

TABLE 2

Objective Likelihoods

	p(Temperature | Disease)			p(Blood pressure | Disease)		
Symptom level	A	B	Log odds	A	B	Log odds
High	0.4	0.1	1.39	0.05	0.5	−2.30
Medium high	0.3	0.2	0.41	0.15	0.3	−0.69
Medium low	0.2	0.3	−0.41	0.3	0.15	0.69
Low	0.1	0.4	−1.39	0.5	0.05	2.30

Note. The underlying likelihoods used to generate the medical charts. For example, when a chart for a patient with disease A shows temperature, it reports a High level with probability 0.4. With disease B the corresponding probability is 0.1, and the log-odds are $\ln(0.4/0.1) = 1.39$.

all subjects. The experiment was participant-paced. Subjects were told this fact and also that all previous subjects had finished in less than the two allotted hours. (This information helped limit subjects' tendency to hurry through the experiment, especially towards the end.) The experiment was broken into three blocks of 160 trials and subjects were permitted five-minute breaks between blocks.

Treatments

A 2 by 3 factorial between subjects treatment scheme was used, with 20 subjects in each cell. (Attendance fluctuations gave us 2 extra subjects in one cell and 1 extra in another.) The evidence treatment had two conditions, History and No History, and the payoff treatment had three conditions, Pay + Score, Score, and None, as explained below.

The History condition gave subjects the option on each trial, before making their response, to view a chart of relevant cases previously encountered. If selected by clicking the "previous cases" icon, the chart stayed on the screen until the subject finished viewing it and clicked an "OK" icon. The chart displayed the number of previous patients with each disease that had possessed the symptom levels present in the current patient, as in Fig. 1B. Subjects in the No History condition had no access to such a chart.

Some treatments involved the computation of a score calculated in each trial from the continuous response $c \in [0, 1]$ and the actual disease $d = 1$ (Autochus) or $d = 0$ (Burlosis) using the quadratic scoring rule $S(c, d) = A - B(c - d)^2$ with $A = 80$ and $B = 280$. The maximum possible score on a trial (correct binary response with complete confidence, so $c = d = 0$ or 1) was $A = 80$ points. The lowest possible score (incorrect binary response, with complete confidence so $|c - d| = 1$) was $A - B =$ negative 200 points. A "Not at All" confident answer, coded $c = 0.5$, always resulted in $A - 0.25B = 10$ points. For comparative purposes, subjects were also told the "expert" score, which would be earned by an ideal statistician familiar with

Table 2, as discussed in the Results section below. MF97 has a more complete discussion of the quadratic scoring rule.

In the Score and Pay + Score conditions, following each trial the screen presented the subject's score on that trial and the cumulative score to that point of the experiment as in Fig. 1C. Subjects in the Pay + Score condition were also paid 1 dollar per 1000 points at the end of the experiment. The payment procedures were explained at the beginning of the experiment. In the None condition, no scores were presented, and no pay was given or discussed.

MODELS

It will be useful to present the competing models of learning within an analytic framework that can clearly distinguish information from information processing. The framework, developed more fully in Massaro and Friedman (1990), is illustrated in Fig. 2. It assumes that learning modifies the first stage of information processing, evaluation of evidence, and that evaluated information from various sources (e.g., separate symptoms) is then integrated and passed to the final stage, decision. Thus evaluation in our medical diagnosis task must be described for the eight possible single symptom values, denoted $k = 1, 2, 3, 4$ respectively for low, medium low, medium high, and high temperature and denoted $k = 5, 6, 7, 8$ respectively for low, medium low, medium high, and high blood pressure. Then for each competing model we describe integration, decision and learning in response to feedback each trial.

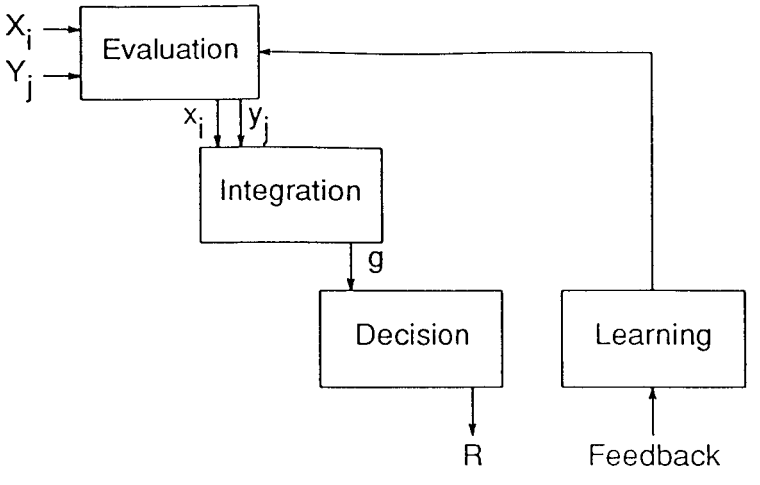

FIG. 2. Schematic representation of the stages of information processing. The evaluation stage transforms the symptom values (indicated by upper-case letters) into psychological values (indicated by lower-case letters). The integration stage takes the output of the evaluation stage and combines or integrates the psychological values to give an overall value for each of the relevant alternatives. The decision stage maps these values into some response, such as a discrete decision or rating. Learning from feedback modifies the evaluation stage.

Bayesian Model

The Bayesian model is normative (or optimal or unboundedly rational) in that it assumes no loss of information or biases introduced by irrelevant information. The basic version of the model, explained more fully in FMKC95, assumes that subjects process information as if they were Bayesian statisticians with initially diffuse priors. An extension allows individual subjects to hold initial beliefs about the relative disease frequencies.

Feature evaluation in the Bayesian model responds to the observed symptom and disease frequencies. Let N_{kA} and N_{kB} be the frequencies of diseases A and B respectively when symptom k is present. The evaluation of feature k is simply the likelihood ratio $r_k = N_{kA}/N_{kB}$ for $k = 1, ..., 8$. The base ratio $r_o = N_A/N_B$ is also relevant; it is the overall frequency N_A of disease A relative to the overall frequency N_B of disease B, irrespective of symptoms.

Learning is implicit in the Bayesian model of feature evaluation. If the feedback on the current trial is disease A, then for each of the current symptoms k the numerator of r_k (and the numerator of r_o) increments by 1. The other numerators and all denominators remain unchanged on that trial. On the other hand, if the current feedback is disease B, then the denominators of r_k (and r_o) each increments by 1 for the current symptoms k, the other denominators and all numerators remain unchanged.

Information integration is described by the well-known Bayes formula. It can be conveniently expressed in ratio (or odds) form as

$$\frac{P[d=A \mid s]}{P[d=B \mid s]} = \frac{P[s \mid d=A]}{P[s \mid d=B]} \frac{P[d=A]}{P[d=B]} \tag{1}$$

The verbal statement of (1) is that the posterior odds are equal to the likelihood ratio times the prior odds. But the prior odds $(P[d=A])/(P[d=B])$ are simply r_o, and (since symptoms are conditionally independent) the likelihood ratio $(P[s \mid d=A])/(P[s \mid d=B])$ simply the product of the r_k for the symptoms k present at that trial.

The formula can be written compactly using the indicator functions $I_k(s_t) = 1$ if symptom k is present on trial t, and $I_k(s_t) = 0$ otherwise. Recall that $k = 1$ to 4 codes temperature and $k = 5$ to 8 codes blood pressure. Hence in each two symptom trial of the current experiment we have $I_k = 1$ for exactly two values of k, one less than or equal to 4 and the other greater than 4. In single symptom trials, we have only one indicator value at 1 and seven at 0. (By contrast, in previous medical diagnosis experiments with up to four binary symptoms, we have up to four values of k with $I_k = 1$.) The notation is convenient because ${r_k}^{I_k}$ is just r_k when symptom k is present and is 1 otherwise. Thus, writing the posterior odds given symptoms s as $Y(s)$, Eq. (1) becomes

$$Y(s) = r_o \prod_{k=1}^{8} r_k^{I_k(s)}. \tag{2}$$

Taking logs and writing the log-odds ratio as $y(s) = \ln(Y(s))$, we obtain a linear version of Bayes formula

$$y(s) = \ln r_o + \sum_{k=1}^{8} I_k(s) \ln r_k. \tag{3}$$

Probabilities can be recovered in the usual fashion from the odds or log-odds ratio. For example, the posterior probability given symptom configuration s is $p(s) = S(y(s))$, where S is the standard sigmoid (or logistic) function $S(y) = (1 + e^{-y})^{-1}$.

Decision stage can be modelled very simply. The continuous response (rating) can be assumed to be the posterior probability, and the binary choice can be assumed to match the posteriori probability or to be simply the more likely disease. More complex models of noisy decision can also be considered, as in FM97; we will touch on them briefly in the concluding discussion.

Implementation of the Bayesian model is quite direct. A version with no free (fitted) parameters, called Bayes.0, begins in the first trial by setting all frequencies to 1, to avoid zero-divide problems. Thus at trial 1 we have $N_{kA} = N_{kB} = N_A = N_B = 1$ for $k = 1, ..., 8$. The model then calculates the likelihood ratios r_k (all equal to 1 in the first trial), plugs these values into Eq. (3) to produce y, and computes the associated posterior probabilities by applying the sigmoid function to y. The associated decisions, the predictions of model Bayes.0, are then recorded in a file. If the symptoms are $k = 1$ (low temperature) and $k = 6$ (medium low blood pressure) and the observed feedback is disease A then learning increments N_A, N_{1A} and N_{6A} to 2 while the other Ns remain at 1. The same procedure is applied in all later trials. In each trial t, for $t = 1, ..., 480$, the previous history provides ratios r_{kt} and r_{ot}, and the current symptom combination s_t is observed. The predicted responses are computed from the logistic transform applied to Eq. (3) and the new values of the rs are computed for trial $t + 1$ by incrementing the appropriate Ns.

We also consider a Bayesian model with one free parameter to account for the possibility that individual subjects hold initial beliefs more or less strongly. In this version, called Bayes.1, we suppose that the initial counts N_A and N_B have the same value n at the beginning of the first trial, and that the initial counts $N_{kA} = N_{kB} = n/2$. For each fixed value of n, we update the Ns after every trial exactly as in Bayes.0 and record the predictions for the entire sequence of trials. Then, using the usual least squares criterion, we find the value of n that best fits a given subject's actual continuous responses. The fitted parameter n measures the strength of a symmetric prior belief that each likelihood ratio is 1.0. In Bayesian jargon, n is a precision parameter for the symmetric prior. Note that larger n implies that beliefs respond less sensitively to trial by trial feedback so learning is slower.

Fuzzy Logical Model of Perception

The fuzzy logical model of perception (FLMP) incorporates developments in fuzzy logic (Zadeh, 1965), pattern recognition (Selfridge, 1959), and choice theory (Luce, 1959) to provide a systematic account of perceptual judgment, decision making, and learning. More flexible in several respects than the Bayesian model, the FLMP

has been tested extensively in a wide variety of domains and has considerable explanatory power in situations with multiple sources of information (Massaro, 1987, 1989, 1994, in press; Massaro & Friedman, 1990). A strong prediction of the FLMP is that the contribution to performance of one source of information increases when other available sources of information become more ambiguous.

The currency of the FLMP model is a truth value that can range from 0 or completely false to 1 or completely true. For example, a 0.3 truth value for the proposition "a whale is a fish" means that the proposition is true to degree 0.3; i.e., there is a moderate degree of similarity between whales and fish. It is psychologically different from a probability. An objective 0.3 probability that "a whale is a fish" would mean that 3 out of 10 whales are fish. A subjective 0.3 probability would mean that a bet (that a whale is a fish) with 7 to 3 odds is fair. Neither of these probabilistic interpretations works because a whale is never a fish.

Feature evaluation in FLMP for the medical diagnosis task produces a truth value (or "feature value") denoted f_k for the statement "the disease is A" when symptom k is present. The truth value of the complementary statement "the disease is B" is denoted g_k. Independently of the symptoms the statement "$d = A$" has some truth value denoted b, and the complementary statement has value b'. Thus there are 8 f_k and 8 g_k for a total of 16 individual feature values.

The integration stage in FLMP combines the feature values to produce integrated truth values, using the same multiplicative formula as in Bayes. That is, even though the feature values are psychologically different from elementary conditional probabilities, the integration stage treats them in an analogous fashion. Specifically, the result of FLMP integration is a truth value for $d = A$ of

$$f(s) = \frac{b[\prod_{k=1}^{8} f_k^{I_k}]}{b[\prod_{k=1}^{8} f_k^{I_k}] + b'[\prod_{k=1}^{8} g_k^{I_k}]}. \tag{4}$$

The truth value for the complementary proposition $d = B$ is $g(s) = 1 - f(s)$. Thus the truth odds take the same form as Bayesian posterior odds:

$$Y(s) = b_0[u_1^{I_1} u_2^{I_2} \cdots u_8^{I_8}], \tag{5}$$

where $Y(s) = f(s)/g(s)$, $b_0 = b/b'$, and $u_k = f_k/g_k$.

The decision stage for continuous choice again is the same as in the Bayesian models: subjects simply report the final truth value as computed in Eq. (4). For binary choice we impose the relative goodness rule (RGR) that the odds with which subjects choose the diseases are given by Eq. (5); see FMKC95 for further discussion of the RGR.

Since Eq. (2) is equivalent to (4) and (5), it might at first seem that, despite their differing psychological roots, in practice the FLMP model is the same as the Bayes model. But the models differ in the way they implement learning, the updating of the feature values $v_{kt} = f_k$, g_k and b_0 from trial $(t-1)$ to trial (t). Learning in the FLMP is described by the general rule

$$v_{kt} = v_{kt-1} + \lambda e_t I_k(s_t). \tag{6}$$

Here λ denotes the constant learning rate and $e_t = d_t - f(s_t)$ denotes the perceived error, given the current feedback d_t and the truth value $f(s_t)$ of that disease assessed using the current feature values for the current symptom configuration s_t.

The learning rule (6) has a long history and extensive literature, as documented in FMKC95. Three points are worth emphasizing at this juncture. First, learning is error-driven in that it occurs only to the extent that there is a nonzero perceived error. By contrast, Bayesian learning takes place on every trial. Second, all features (and the base odds) are learned at the same rate λ. This is restrictive since in practice some features (or the base rate) might be learned faster than other features. Third, the learning rate λ is constant. By contrast, although Bayesian updating can be expressed in a manner analagous to Eq. (6), the Bayesian learning rate λ is not constant over time but rather declines as $1/t$; see Appendix of FMKC95 for a derivation and discussion.

The basic implementation of FLMP analyzed below is called FLMP17.1 since it has 17 internal parameters—the base odds b_0 and the feature values f_k and g_k for $k = 1, ..., 8$—that are updated every trial via Eq. (6) using one free parameter. That free parameter, the learning rate λ, is estimated via ordinary least squares.

To check robustness, at one point we also estimate FLMP17.2, which has a second free parameter, a learning decay rate called γ. This version replaces λ in Eq. (6) by $\lambda_t = \lambda t^{-\gamma}$. Thus $\gamma = 0$ corresponds to the constant learning rate in the basic FLMP, $\gamma = 1$ corresponds roughly to the Bayesian decay rate, and intermediate values of γ indicate slower decay in the learning rate.

CMP

Connectionist models, also known as neural network models, have a central place in the psychological and computer science literature on learning (e.g., Rosenblatt, 1958; Minsky & Papert, 1969, 1988; Rumelhart & McClelland, 1986). We confine our attention to simple two layer feed-forward networks in which each symptom $k = 1, ..., 8$ defines an input node connected to the two output nodes corresponding to diseases *A* and *B*. We also include a constant node (indexed as 0) to account for possible bias towards disease *A* or *B* in the absence of evidence. Thus there are 18 connection weights to be learned, but standard normalizations (e.g., that weights sum to 0 over connections leaving any given input node) reduce the number of internal parameters to 9. The model can be summarized in terms of log-odds $y(s)$ given symptom configuration s by

$$y(s) = w_0 + \sum_{k=1}^{8} w_k I_k(s) \tag{7}$$

with the learning algorithm in (6) applied to the w_ks instead of the v_ks. The decision rule is the same as in FLMP, applied to the output of (6) transformed (from log odds to values in the range $[0, 1]$) by the standard sigmoid function $S(y) = 1/(1 + e^{-y})$. These conventions are quite standard (Rescorla & Wagner (1972), Gluck & Bower (1988), Estes *et al.* (1989), & FMKC95).

We refer below to this implementation as CMP9.1 since learning (specified by a single free parameter, λ) updates nine internal parameters, the connectionist weights $w_0, ..., w_8$. Some variants are worth brief mention. One could use un-normalized weights for connections to the symptom nodes (CMP17). By the same token, one could reduce the 17 internal FLMP parameters to 9 by applying the learning scheme (6) to the ratio $r_k = f_k/g_k$ instead of applying it separately to numerator and denominator (FLMP9). As shown in FMKC95, one can absorb the base odds (or zero node) parameter into the other 8 parameters in CMP9 or FLMP9, producing models denoted CMP8 and FLMP8. Despite their very different psychological underpinnings, the CMPx learning model is mathematically equivalent to FLMPx ($x = 8, 9, 17$) except that the learning algorithm (6) is applied to internal parameters w that correspond to logarithms of the internal FLMP parameters v. Thus the CMP updating rule is multiplicative rather than additive. See Kivinen and Warmuth (1995) for a theoretical justification for using multiplicative updates.

Other Reinforcement Learning Models

The models considered so far involve rather sophisticated integration of information across symptoms. Perhaps subjects do not integrate the information but just learn the disease association for each symptom configuration separately. The simplest implementation of such reinforcement learning is to say for each of the 24 symptom configurations s a subject has propensities $q_{At}(s)$ and $q_{Bt}(s)$ for decisions A and B at time t; that her continuous response (or probability of choosing A in binary mode) is the relative propensity $q_{At}(s)/(q_{At}(s) + q_{Bt}(s))$, and that she increments by the rule $q_{Dt+1}(s) = q_{Dt}(s) + aI_{Dst}$ for $D = A, B$, where a is a positive constant and I_{Dst} is the indicator function which $= 1$ if the symptom configuration on trial t was D and the symptom configuration was s and which $= 0$ otherwise. Thus learning consists of incrementing by a the propensity for the observed disease/symptom configuration, and the other 47 propensities are left unchanged. Roth and Erev (1995) are influential recent advocates of this sort of reinforcement learning, which embodies the classic "law of effect" and "power law of practice."

In our implementation, we set $a = 1$ without loss of generality, so the propensities are simply counts of disease occurrence by symptom configuration. (By contrast, the Bayesian scheme involves counts of disease occurrence by individual symptom levels, and integrates across the symptoms in a configuration.) Initial counts at $t = 1$ are all set to 1 in the implementation RL.0 or are all set to a free parameter q_o in RL.1. Thus RL.0 and RL.1 parallel Bayes.0 and Bayes.1, but assume less sophisticated integration of information across symptoms.

Exemplar models are versions of reinforcement learning that have previously been applied to the medical diagnosis task (Estes *et al.*, 1989; Medin, Altom, Edelson, and Freko, 1982). Exemplar models store the complete symptom configuration and the associated disease. At each trial, the exemplars in memory are used to categorize the input. In our implementation of the Exemplar model, each of the symptom configurations (here indexed $j = 1, ..., 24$) has a weight W_j that reflects the fraction of stored exemplars associated with disease A. Thus evaluation stage is essentially the same here as in the RL models. Again like the RL models,

the integration stage is bypassed because each symptom configuration is evaluated separately. But decision and learning in our implementation of Exemplar are like CMP, not like RL (or Bayes). The weight W_j assigned to the current symptom configuration is transformed at decision stage using the sigmoid function. Thus the predicted continuous response (and the predicted probability that the binary response is disease A) is $S(W_j) = 1/(1 + e^{-W_j})$.

Exemplar learning applies Eq. (6) to the weight W_j corresponding to the symptom configuration observed in the current trial and leaves the other 23 weights unchanged. The learning rate parameter λ is fitted to the data in the usual fashion. We refer to this model as Exemplar24.1 because a single learning parameter operates on the 24 internal parameters.

A final model we consider is a hybrid of the exemplar and connectionist models known as ALCOVE (Kruschke, 1992). Noted for its flexibility and good fits to a variety of data, the model can be regarded as a three layer feed-forward network with 24 connection weights. Our present implementation is essentially the same as in FMKC95, which explains this relatively complex model. The conventions on all stages of information processing (evaluation, integration and decision) and learning are exactly the same as in the CMP model. Again we rely on ordinary least squares to estimate a single free parameter λ.

RESULTS

We present the results in three subsections. The first presents several summary descriptions of learning by individual subject and by treatment condition. The second compares the explanatory power of competing learning models. The third develops some extensions of the Bayesian learning model.

Summary Statistics

The first question is whether subjects were able to distinguish among the symptom configurations. Figure 3 presents some favorable evidence. For each of the 24 symptom configurations, Fig. 3A takes the average continuous response over the twenty subjects in the History/Score treatment cell and the first block of 240 trials, and plots that average against the true posterior probability. The entries all lie reasonably close to the diagonal; the correlation coefficient of average continuous response and true probability is 0.944. The same subjects do even better in the second block (trials 241–480); panel B shows that the correlation rises to 0.971.

Binary responses can be analyzed in a parallel manner since we are averaging across trials within a block and across subjects within a treatment cell. Figure 3C shows that the subjects' average binary choice in the first block lies further from the diagonal than their average continuous choice, but (by the same token) they reliably choose the more likely disease for almost all the symptom configurations. An exception is symptom configuration *J* (low temperature and medium low blood pressure): The actual choice frequency for disease *A* is almost 0.5, while the true (Bayesian posterior) probability is 0.33. Figure 3D shows that the binary choices

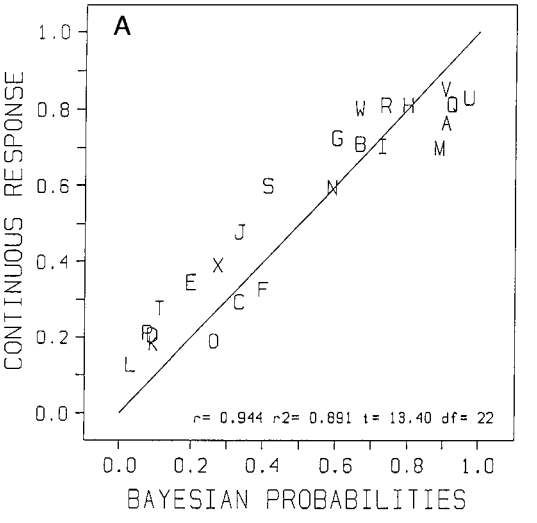
A
CONTINUOUS RESPONSE
BAYESIAN PROBABILITIES
r= 0.944 r2= 0.891 t= 13.40 df= 22

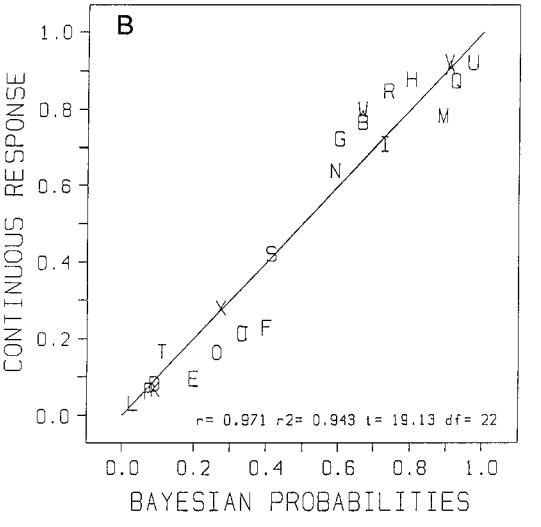
B
CONTINUOUS RESPONSE
BAYESIAN PROBABILITIES
r= 0.971 r2= 0.943 t= 19.13 df= 22

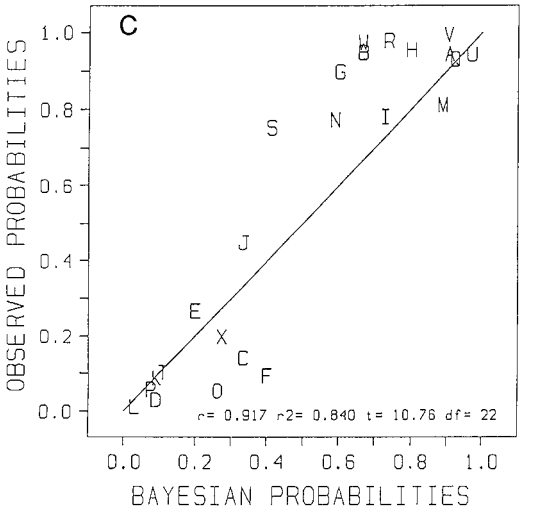
C
OBSERVED PROBABILITIES
BAYESIAN PROBABILITIES
r= 0.917 r2= 0.840 t= 10.76 df= 22

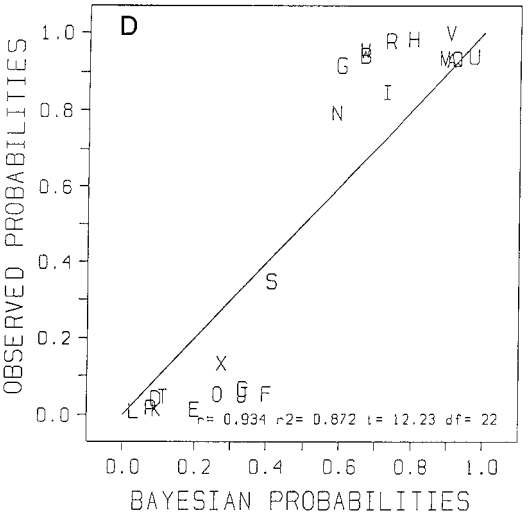

FIG. 3. Mean response by symptom configuration in score/history condition: (A) continuous response - block 1; (B) continuous response - block 2; (C) binary response - block 1; (D) binary response - block 2. *Note.* The list of symptom configurations appears below with the labels (A–X) used in the scatter plots, the values (0 = absent, 1 = low, 2 = med. low, etc.) of temperature (Temp) and blood pressure (bp), the number of trials (n) and the true posterior Bayesian probabilities (Bay) over all trials.

Source Data

Label	Temp	bp	n	Bay
A	0	1	22	0.90909
B	0	2	18	0.66667
C	0	3	18	0.33333
D	0	4	22	0.09091
E	1	0	20	0.20000
F	2	0	20	0.40000
G	3	0	20	0.60000
H	4	0	20	0.80000
I	1	1	11	0.72727
J	1	2	15	0.33333
K	1	3	21	0.09524
L	1	4	33	0.03030
M	2	1	18	0.88889
N	2	2	17	0.58824
O	2	3	19	0.26316
P	2	4	26	0.07692
Q	3	1	26	0.92308
R	3	2	19	0.73684
S	3	3	17	0.41176
T	3	4	18	0.11111
U	4	1	33	0.96970
V	4	2	21	0.90476
W	4	3	15	0.66667
X	4	4	11	0.27273

are even more reliable by the second block, e.g., disease A is rarely chosen for symptom configuration J.

The same general picture emerges for other subsets of subjects; the scatterplots are suppressed to conserve space. Even in the first block, the average responses indicate that subjects have already learned to interpret properly the more informative symptom configurations, and performance improves slightly in the second block. Correlation coefficients ranged from 0.916 for binary response in the first block of the History/Score + Pay cell to 0.981 for continuous response in the second block of the History/Score + Pay cell.

Another summary description of learning success is the score earned by subjects. Recall that we use the quadratic scoring rule $S(c, d) = 80 - 280(c - d)^2$, where c is the subject's continuous response and $d = 0$ (for B) or 1 (for A) is the actual disease on that trial. This score is reported to subjects in four of the six treatment cells, and can be computed ex post for subjects in the other two cells. A subject who reported $c =$ the true Bayesian posterior probability of disease A each trial would earn a total score of 19599 in the 480 trial session. But this benchmark (sometimes referred to as the "expert score") is unrealistic for subjects who have to learn the relation between symptoms and diseases. An ideal learner who faithfully used the Bayes.0 model would earn a total score of 18330. By contrast, a subject who never learned anything but chose the best constant response $c = 0.5$ every trial (referred to below as "optimal ignorance") would earn a score of 4800, and an omniscient but perverse subject who always reported the wrong disease with complete confidence ($c = 1 - d$) every trial would earn the theoretical minimum total score of $-96,000$.

Figure 4 shows that the scores of our 123 subjects varied widely. The final score was highly negative for one subject, but the majority of subjects received scores in the 10,000 to 18,000 range. (The one aberrant subject is omitted from subsequent analysis.) We observe some systematic differences by treatment conditions. The mean final score across all of the subjects in the History conditions was 13325, while the mean score for those subjects in the No History conditions was 12474. The difference is significant according to an ANOVA test: $F(2, 121) = 4.71$, $p < 0.05$. Mean final scores (collapsed across the the two levels of History) are 13088 in the No Score condition, 12917 in the Score condition and 13978 in the Score + Pay condition; but the effect of the score treatment is not significant: $F(2, 116) = 1.70$, $p = 0.188$. Also, the interaction between Score and History was not significant: $F(2, 116) = 0.92$, $p = 0.40$. We conclude that subjects generally did a fairly respectable job of learning the relationship between symptoms and diseases, and learned faster with the benefit of the History treatment.

Learning Model Fits

The main competing models we consider are Bayes (with 0 and with 1 fitted parameter, n) and FLMP17, CMP9, Exemplar24, and Alcove24 (each with 1 fitted parameter, λ). We will also look at several variants of the models. The fits of all models are trial-by-trial: Given the stimuli and diseases observed so far, the model predicts the continuous response on the next trial. We choose free parameters

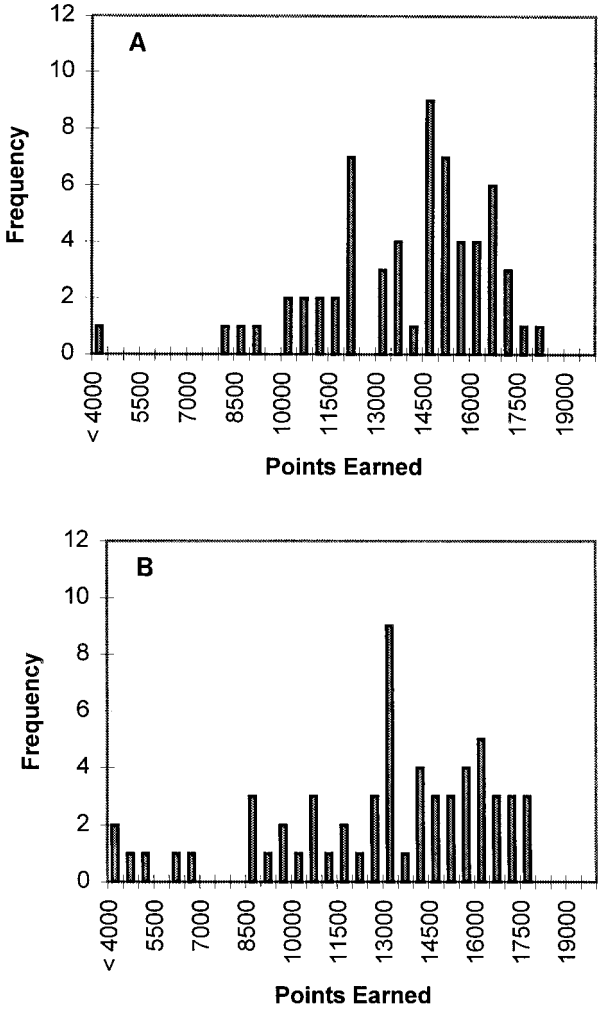

FIG. 4. Distribution of subjects scores: (A) all history conditions; (B) all no history conditions.

separately for each subject to minimize the total squared prediction error or, equivalently, the root mean squared deviation (RMSD).

Figure 5A shows that subjects varied in their resemblance to the ideal statistician described in Bayes.0. One subject's RMSD was about 0.12, indicating that her continuous response typically was within 12% of this ideal. The modal RMSD was about 0.16 and the vast majority of subjects had RMSDs less than 0.30 for Bayes.0. By comparison, for the optimal ignorance benchmark, the average RMSD across all subjects was 0.327 in block 1 and 0.358 in block 2. Figure 5B indicates relatively good fits for the Bayes.1 model, with almost all subjects' RMSDs in the 0.10 to 0.30 range and median near 0.20. The fits of the other models in Figs. 5C–5F seem not quite as good.

With the exception of Bayes.0 (which has no fitted parameters), the learning models all show considerable dispersion in estimated learning rates across subjects. Recall that the prior precision parameter n in Bayes.1 (and RL.1) is inversely

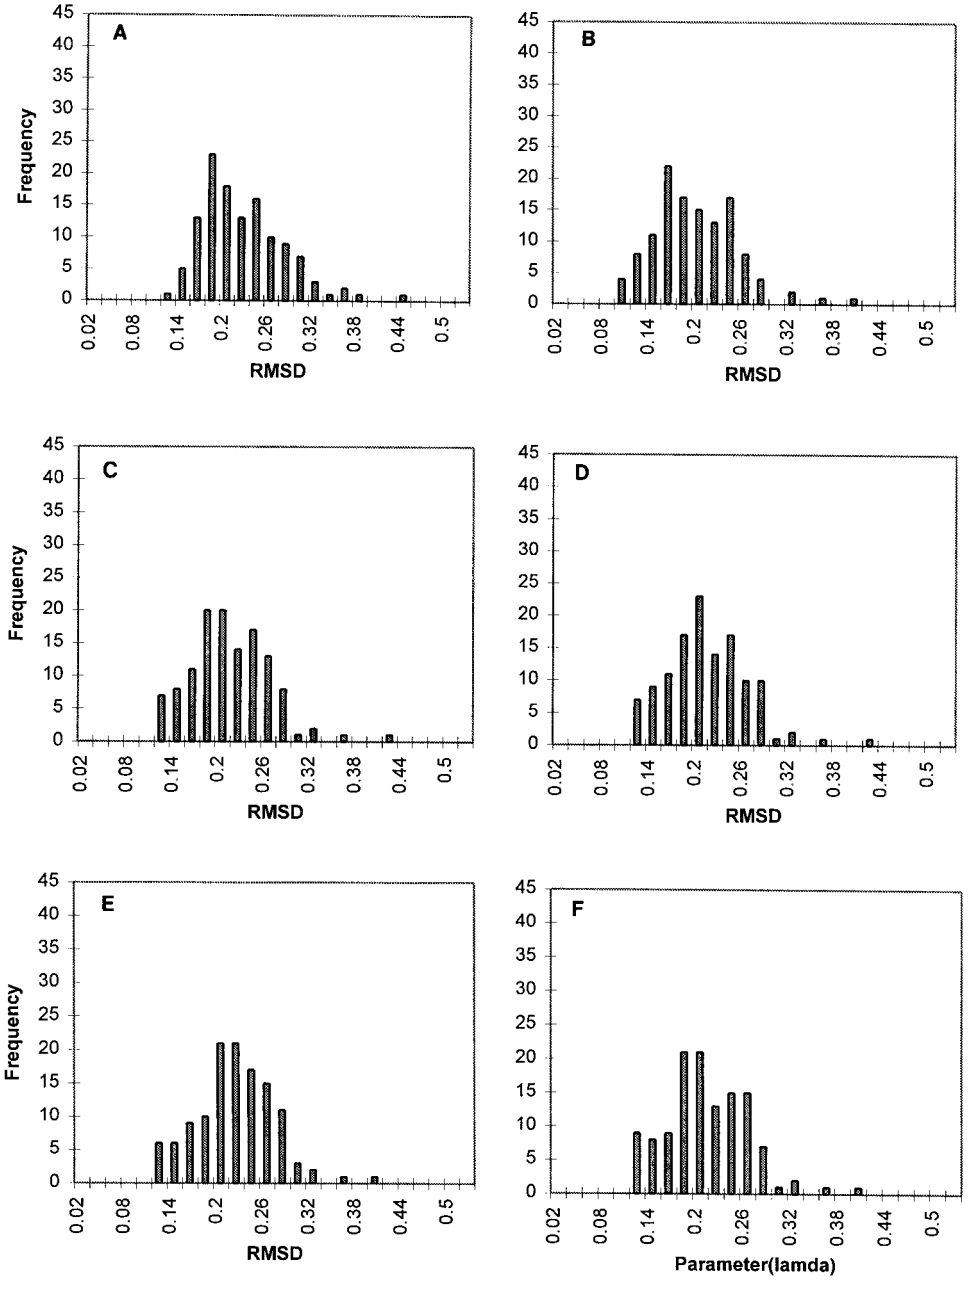

FIG. 5. Distribution of RMSDs: (A) Bayes.0 model fits; (B) Bayes.1 model fits; (C) FLMP17.1 model fits; (D) CMP9.1 model fits; (E) Exemplar24.1 model fits; (F) Alcove24.1.

related to the learning rate. Figure 6B shows that the modal value is about 5, but a scattering of subjects have fitted values up to 70 and about 8 subjects have values exceeding 125, indicating very slow learning. In the other models, the fitted value of λ varies by an order of magnitude, typically from 0.01 to 0.1 in FLMP17.1, from 0.02 to 0.4 in CMP9.1, from 0.03 to 0.3 in Exemplar24.1, and from 0.1 to 1.0 in

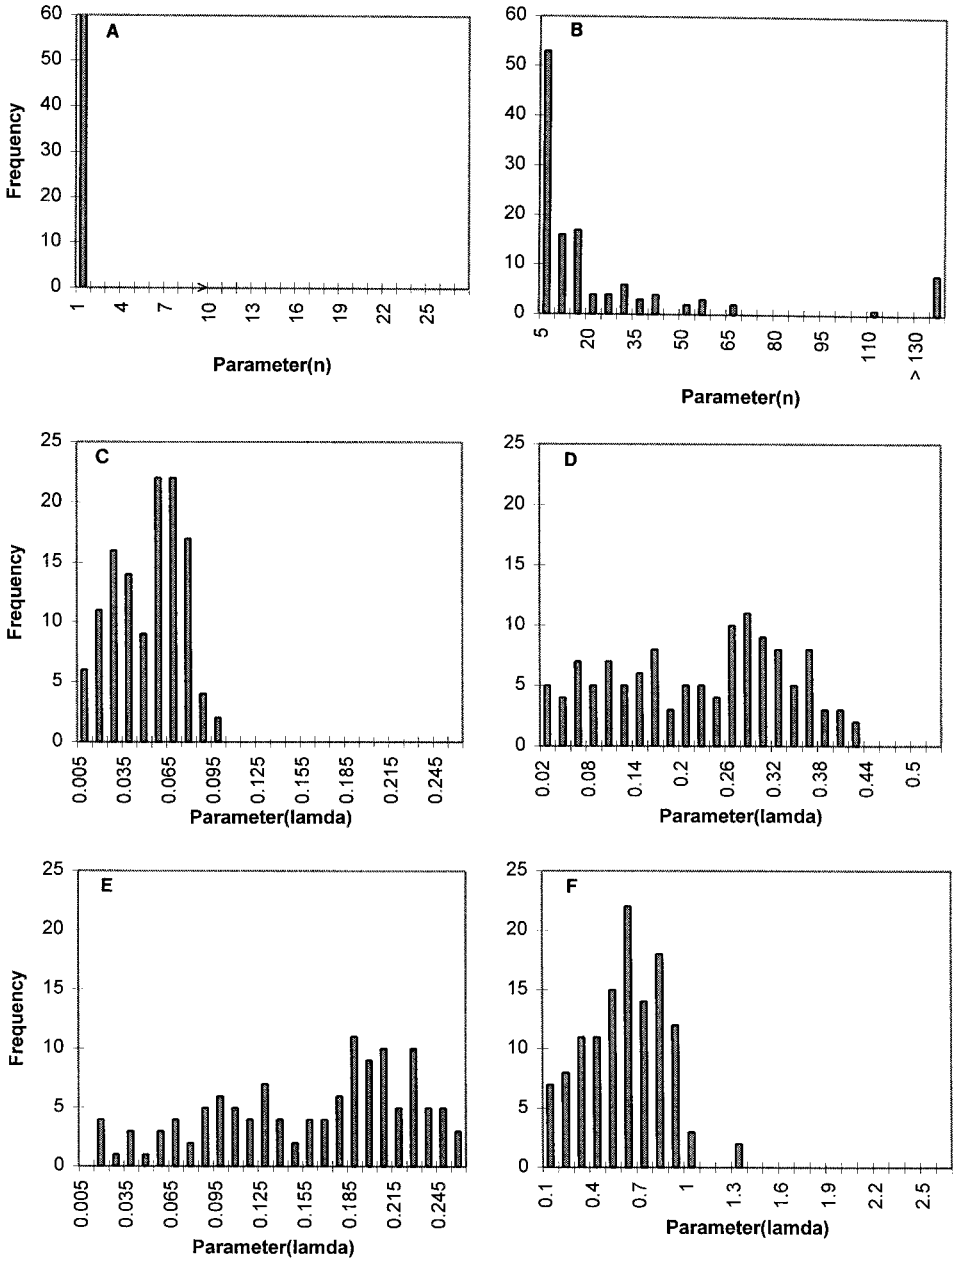

FIG. 6. Distribution of learning parameters: (A) Bayes.0; (B) Bayes.1; (C) FLMP17.1; (D) CMP9.1; (E) Exemplar24.1; (F) Alcove24.1.

Alcove24.1. Larger estimated values of λ are to be expected in the models that bypass integration (Exemplar and Alcove) because they update their internal parameters in fewer trials. The dispersion of parameter estimates across subjects, however, is even greater than we expected.

Tables 3 and 4 report the main results of the learning model competition. Despite the absence of free parameters, the normative Bayes.0 model has respectable fits

TABLE 3

Continuous Response RMSD by Treatment

		No history			History			
Model	NP	No score	Score	Score + pay	No score	Score	Score + pay	All Subjects
1. Bayes	0	.2302	.2276	.2335	.2231	.1775	.1803	.2124
2. Bayes	1	.1791	.2049	.1952	.2016	.1680	.1580	.1847
3. FLM17	1	.1891	.2177	.2040	.2180	.1919	.1774	.1999
4. CMP9	1	.1886	.2182	.2034	.2183	.1926	.1762	.1998
5. Ex24	1	.1936	.2255	.2106	.2267	.2075	.1906	.2092
6. Alcove24	1	.1892	.2180	.2030	.2137	.1876	.1724	.1975
7. Bayes w/l rates	18	.1627	.1814	.1763	.1845	.1567	.1446	.1679
8. Ex24 w/l rates	25	.1543	.1736	.1692	.1790	.1547	.1432	.1626
9. RL	0	.2150	.2289	.2267	.2305	.2014	.1912	.2159
10. RL	1	.1899	.2206	.2071	.2190	.1959	.1812	.2025

Note. The models are described in the text. NP is the number of fitted (or "free") parameters fitted for each subject to minimize squared prediction error over all 480 trials. The twenty or more subjects were used in each of the six treatment conditions (History, Score, etc.); there were 122 subjects in all. Table entries are (minimized) root mean squared prediction error (RMSD) for each model in subjects' trial-by trial continuous response.

with RMSDs ranging from about 23% in the No History conditions down to about 18% in the History conditions with Score. The main competitors did better, but typically reduced the forecast error by only a few percentage points. Over all subjects and trials, the average prediction error was about 21% in Bayes.0 and was a bit under 20% given a single fitted learning parameter for each subject in the CMP9, FLMP17, and Alcove24 models. The forecast error exceeded 20% for Exemplar24 even with a free parameter. Bayes.1 was the top performer, with RMSD averages as low as 16% in the History/Score + Pay cell.

We were surprised that the competing models offered so little improvement over Bayes.0 and no improvement over Bayes.1, so we ran a number of robustness checks. First we confirmed that changing the number of internal parameters (as described in the previous section) made little difference in the CMP and FLMP fits. We confirmed that adding a second learning parameter (the decay rate γ in the models using Eq. (6) or a prior asymmetry parameter in Bayes.1 and RL.1) improved fits only very slightly, typically reducing RMSD by 0.001–0.005, and the fitted values indicated that very few subjects exhibited rapid decay of the learning rate or asymmetric priors. We considered logistic transforms of the choice data using estimates reported in FM97, but found no consistent improvement in model fits. At one stage of our research we also examined model fits using maximum likelihood (ML) rather than least squares (LS) techniques, and found no notable differences.

The last four rows of Table 3 report some of the more interesting robustness checks. Perhaps different internal parameters are learned at different rates. Row 7

TABLE 4

RMSD Comparisons all Subjects

Model	NP	Continuous response					Binary response
		First 60	First 240	Last 240	Last 60	All trials	All trials
1. Bayes	0	.2625	.2269	.1948	.1949	.2124	.3184
2. Bayes	1	.1753	.1873	.1755	.1759	.1847	.3162
3. FLMP17	1	.1915	.2033	.1872	.1840	.1999	.3387
4. CMP9	1	.1917	.2036	.1882	.1844	.1998	.3400
5. Ex24	1	.2138	.2177	.1911	.1846	.2092	.3519
6. Alcove24	1	.1914	.1995	.1858	.1821	.1975	.3378

Note. The models are described in the text. NP is the number of fitted (or "free") parameters fitted for each subject to minimize squared prediction error over the indicated set of trials. Table entries are (minimized) root mean squared prediction error (RMSD) for each model in subjects' trial-by-trial continuous and binary response.

allows each of the 16 likelihoods (of each level of each symptom given each disease) in Bayes as well as the unconditional disease probability to have its own prior precision. Row 8 allows each symptom configuration in Exemplar to have its own learning rate and also allows for asymmetric prior beliefs. These rows indeed have the lowest forecast errors, below 15% in the most favorable conditions. But given the excessive number of free parameters, one should regard the entries as lower bounds on forecast error in our data and not as the performance of serious candidates to explain human learning. The reinforcement learning model RL.0 is potentially a serious candidate but it only narrowly outperformed its counterpart Bayes.0 in the noisier treatments and elsewhere had larger prediction error. Performance improved with a fitted prior precision parameter, but RL.1 still was bested in every condition by its Bayesian counterpart Bayes.1.

The averages in Table 3 might disguise diversity across subjects. We know from Fig. 6 that subjects are quite diverse in their fitted parameters for a given model. More importantly, one might ask whether different models (not just different parameters) are needed to explain the learning processes of different subjects. To explore this issue, for each of the five main contenders, we counted the number of subjects for which it had the smallest RMSD. It turns out that Bayes.1 was the best fit for 111 of 122 subjects; Exemplar24.1 for 5 subjects, Alcove24.1 for 3, FLMP17.1 for 2 and CMP9.1 for 1. The scattering of the 11 non-Bayesian winners seems unrelated to the Score and History treatments.

The averages might also disguise diversity over time. Perhaps some models perform better in early periods when learning is rapid and not as well in later periods. Table 4 reports the average performance over time of the main contenders. Indeed we do see more separation between the contenders in earlier periods than in later periods, but the main finding is reinforced: Bayes.1 provides the best explanation of continuous choice in each time period as well as the best overall. The last column shows that the prediction errors for all candidate models are substantially larger for

average binary choice than for continuous choice, but that the relative performance of Bayes.1 (and even Bayes.0) is quite good.

Bayesian Learning Curves

The Bayesian model's predictive success encouraged us to analyze more carefully its trial-by-trial performance. Recall from Eq. (3) that the Bayesian relationship between symptoms and disease takes a linear form in terms of log odds:

$$y(s) = \ln r_o + \sum_{k=1}^{8} I_k(s) \ln r_k.$$

We can observe the subject's learning process through a Bayesian lens by replacing the left hand side $y(s)$ of (3) by the log odds of the actual response, and estimating the symptom coefficients over various time periods. Thus we construct the dependent variable $L(c) = \ln[(c + 0.01)/(1.01 - c)]$, shifting the continuous choice c by 0.01 away from 0 and 1 to avoid taking the log of zero, and we take the symptom indicators as explanatory variables. We estimate the rolling regression

$$L(c_t) = \beta_{oT} + \sum_{k=1}^{8} \beta_{kT} I_k(s_t) + \varepsilon_t \tag{8}$$

over a moving window of 160 consecutive trials, incrementing the last trial T from 160 to 480. Effective learning is indicated by rapid convergence of the coefficient estimates β_{kT} (as T increases) to the objective values $\ln r_k$ listed in Table 2. Obstacles to learning are suggested by slow convergence, convergence to some other value, or divergence of the coefficient estimates. This empirical approach embodies some of the theoretical ideas on learning in Marcet and Sargent (1989) as explained in KF98.

Figure 7 presents two examples. The intercept coefficient β_{oT} is constrained to its objective value 0 to reduce clutter and to improve statistical efficiency. The dependent variable in Fig. 7A uses the continuous responses generated by the Bayes.0 normative model. All eight coefficient estimates indeed converge closely to the objective values in the last third of the trials. Figure 7B uses the actual continuous choices of a high-scoring subject. The coefficient estimates maintain essentially the same ordering as the objective values, but the subject appears to overrespond to the more informative symptoms (denoted $k = 5$ and $k = 8$ for low and high blood pressure) and perhaps underresponds to the least informative symptoms. Other subjects exhibit similar but noisier rolling regressions.

A more parsimonious representation (a direct analogue of the central equation in KF98) is obtained by exploiting regularities in the objective log odds ratios $w_k = \ln r_k$. Recall from Table 2 that $w_1 = -w_4 \approx 1.39$ while $w_2 = -w_3 \approx 0.41 \approx 0.3w_1$, and $-w_5 = w_8 \approx 2.30$ while $-w_6 = w_7 \approx 0.69 \approx 0.3w_8$. Hence essentially all of the information embodied in the eight symptom indicators is captured in the two variables $x_1(s) = I_1(s) + 0.3I_2(s) - 0.3I_3(s) - I_4(s)$ for all levels of the first symptom

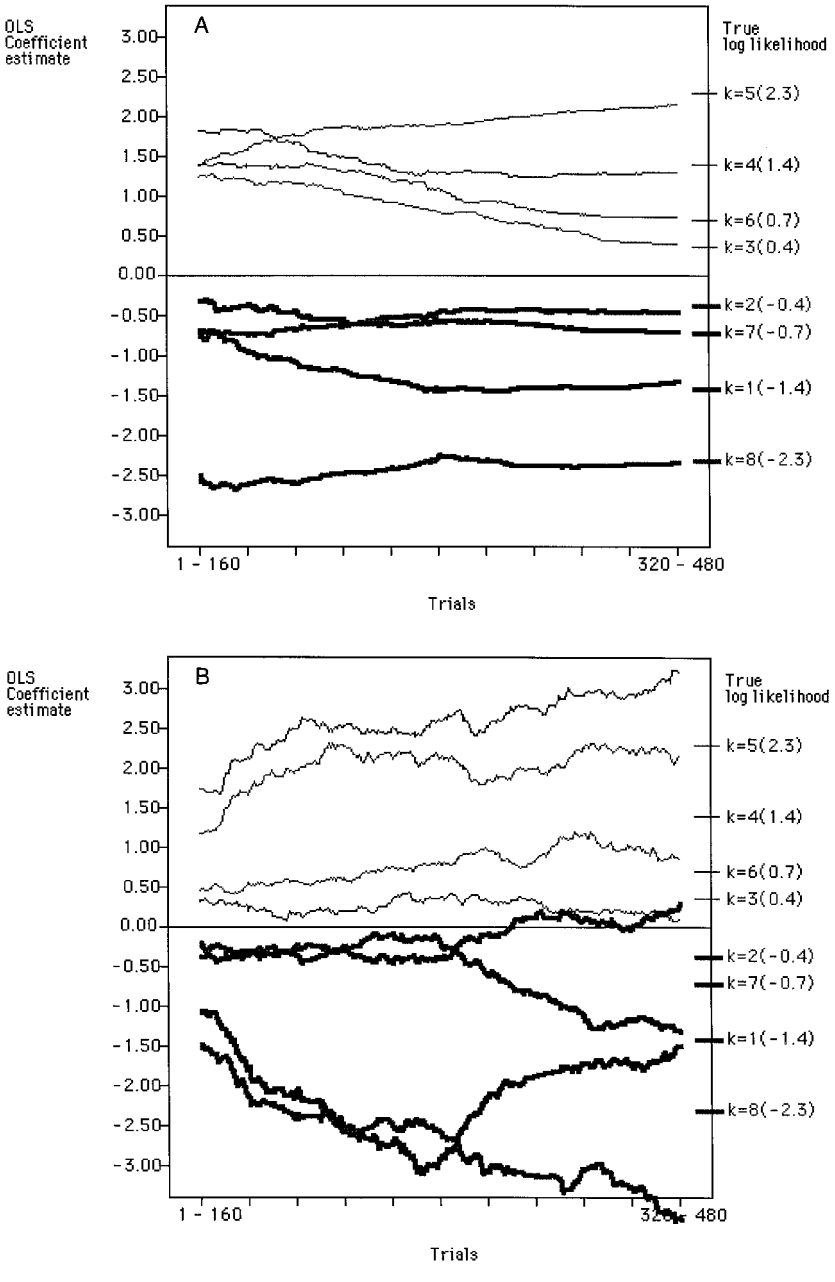

FIG. 7. Rolling regressions for eight indicator variables: (A) Bayes.0 simulation; (B) subject #28.

(temperature) and $x_2(s) = I_5(s) + 0.3I_6(s) - 0.3I_7(s) - I_8(s)$ for all levels of the second symptom (blood pressure). The rolling regression

$$L(c_t) = \alpha_{oT} + \alpha_{1T}x_1(s_t) + \alpha_{2T}x_2(s_t) + \varepsilon_t \tag{9}$$

estimated over the same windows ending at $T = 160$–480 thus shows effective learning to the extent that the estimated coefficients α_{iT} converge to the objective values

$\alpha_1 \approx 1.4$ and $\alpha_2 \approx -2.3$. Figure 8 presents the same two examples (again constraining the intercept to its objective value of 0). Here the Bayes.0 simulation finds the objective values right away and hardly waivers. The high-scoring subject somewhat overresponds. KF98 show that this subject is quite typical–on average subjects in the same treatment cell (History with Score but no pay) overrespond to about the same degree as in Fig. 8B.

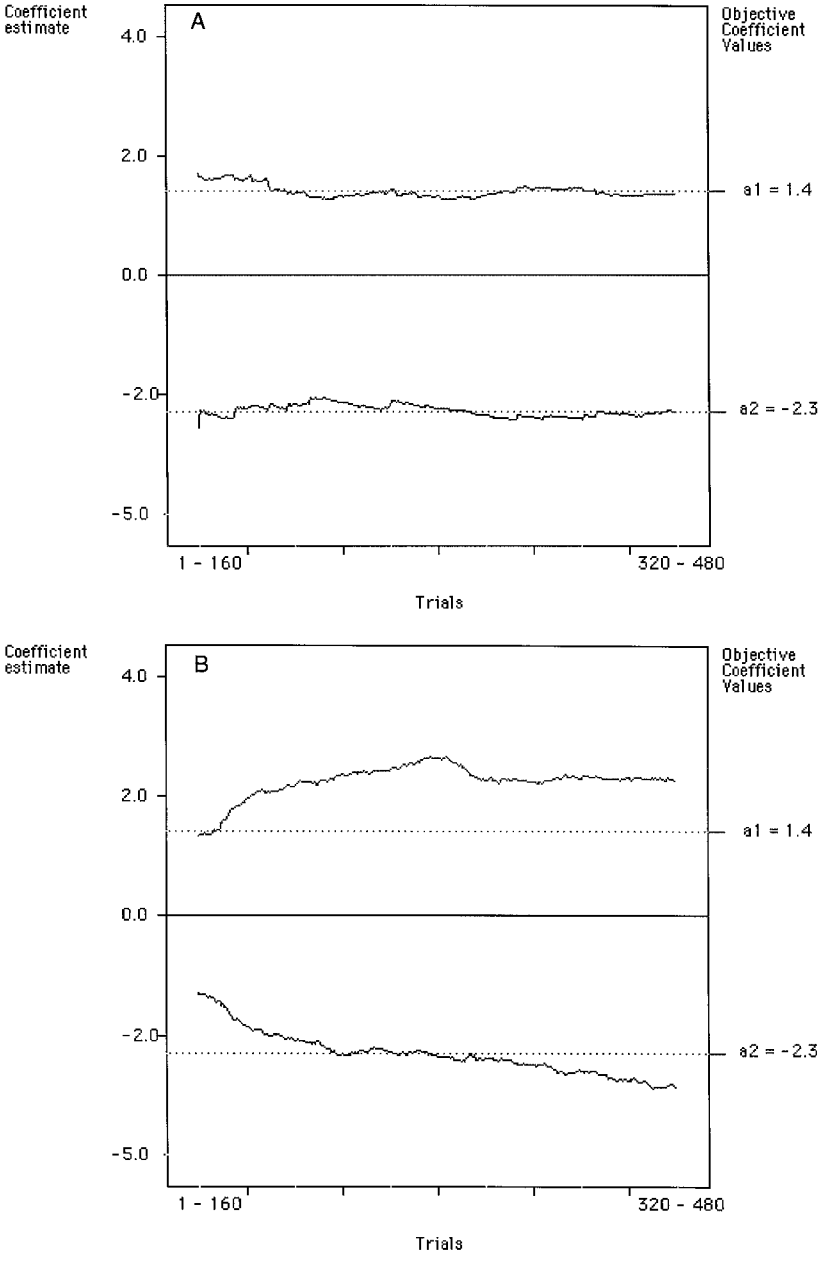

FIG. 8. Rolling regressions for two composite symptom variables: (A) Bayes.0 simulation; (B) subject #28.

DISCUSSION

Our results point to the following conclusions.

1. Subjects do learn in our broadened medical diagnosis task. Even during the first half of the 480 trials they usually choose the more likely disease and their average continuous response reflects fairly closely the true reliability of the symptom configurations. They do even better in the second half of the trials.

2. Subjects are quite heterogeneous. Some of them earn much higher overall scores than others in the task. Fitted learning rate coefficients vary widely across subjects in each learning model. For example, the Bayes.1 model indicates that some subjects hold strongly to initial beliefs and respond only slightly to the feedback in each trial, while other subjects have weaker priors and respond much more strongly to evidence in early trials.

3. Some of the variation across subjects comes from differences in experimental treatments such as Score (whether or not a quadratic scoring rule is used to rate performance trial by trial) and History (whether or not subjects have access to a summary of symptoms and diseases seen in earlier trials). In a companion paper we found statistically significant impacts of the Score treatment, and in the present paper we find that with the History treatment the subjects earned significantly higher average scores and tended to behave more as predicted by the learning models. Nevertheless, it is clear that subjects vary more within a given treatment than across treatments.

4. Although subjects differ in many ways, they appear to be remarkably similar in the underlying way that they process information. We considered five quite different learning models featured in the recent literature. Contrary to our expectations, the same model accounts best for information processing in virtually all subjects. Again contrary to the expectations of the majority of authors, the best predictor of our subjects' behavior is Bayes.1, a normative model of optimal information processing with a single free parameter to capture individual differences in prior information.

5. Even our best fitting model left unexplained considerable behavioral variability, with prediction errors (RMSD) of 10 to 26% for most subjects. Using new rolling regression techniques (borrowed from Kelley and Friedman, 1998) we found that the systematic behavior of most subjects moved over time towards the normative (optimal) value, but even at the end of the session, subjects still tended to overreact to current evidence.

Many question remain open. Of course, the general issue of what sort of tasks favor different competing learning models will remain active for many years. Nearer term, we hope to see progress on two puzzles that remain in our work.

1. The Bayesian model (and also the RL model) has a learning rate that decays as $1/t$; the weight accorded new evidence deadlines proportionately to the accumulation of old evidence. Yet allowing the learning rate to decline (at a faster or slower rate than $1/t$) in the error based learning models such as FLMP17.2 did not much improve their explanations of the data.

2. FM97 summarizes conservatism (or underconfidence) in terms of a logistic parameter. A simple transformation of the continuous choice data based on fitted values

of that parameter brings most subjects' averaged choices (across trials with a given symptom configuration) closer to the objective posterior probabilities. Yet the transformation does not consistently improve the trial-by-trial fits of the learning models. We suspect the apparent inconsistency may be related to overresponse to current symptoms, but we are unsure of how best to model the combination of conservatism and overresponse.

The findings on the predictive superiority of Bayes.1 must be interpreted with caution. For several decades, cognitive psychologists have constructed and tested all sorts of non-Bayesian models precisely because of shortcomings in Bayesian model as a description of human behavior. Edwards (1968) is a delightful survey of early results, mostly for urn problems. For example, what is the probability that a sample of 8 red and 4 blue balls is drawn from an urn that has 70% red balls rather than one with 70% blue balls? The Bayesian posterior probability (under the usual conventions) is 0.97 but subjects' responses typically are in the 0.7 to 0.8 range, a bias called conservatism. However, Edwards notes that the bias is greatly reduced in more ecological tasks (eg, men's and women's heights, rather than two urns) and by training and experience. See Camerer (1995) for a recent survey of cognitive biases.

With feedback in 480 trials and an ecologically interpretable task, our medical diagnosis data should be expected to exhibit relatively small biases. Moreover, biases are present in our data, including a form of conservatism reported in MF97 as well as the persistent overreaction to current evidence mentioned in point 5 above. It just turns out that the non-Bayesian descriptive models that we collected from the recent literature don't capture the regularities in our data as effectively as Bayes.1. Many interpretations of this fact are possible. One interpretation that we find attractive is that human rationality is bounded and adaptive, and will resemble the normative Bayesian model only to the extent that the environment encourages such adaptive behavior (e.g., Gigerenzer and Hoffrage, 1995).

The belief that learning is domain specific is widespread among psychologists, as we noted in the introduction. A referee reminds us that some game theorists believe that, even within the class of two player bimatrix games, different forms of learning take place depending on whether the game is zero sum. We hope our results will encourage other investigators to reconsider domain general learning. We found that the same model predicted best across a fairly broad class of stochastic individual choice environments. Our conjecture is that a broad adaptive model that specializes in our environments to something resembling Bayes.1 will predict well across most interesting individual choice environments as well as interactive game environments.

REFERENCES

Camerer, C. (1995). Individual decision making. In J. H. Kagel & A. E. Roth, (Eds.), *The handbook of experimental economics*, pp. 587–703. Princeton, NJ: Princeton Univ. Press.

Edwards, W. (1968). Conservatism in human information processing. In B. Kleinmuntz (Ed.), *Formal representation of human judgment*, pp. 17–52. New York: Wiley.

Estes, W. K., Campbell, J. A., Hatsopoulos, N., & Hurwitz, J. B. (1989). Base-rate effects in category learning: A comparison of parallel network and memory storage-retrieval models. *Journal of Experimental Psychology*: *Learning*, *Memory*, & *Cognition*, **15**(4), 556–571.

Friedman, D., & Massaro, D. W. (1997). *Understanding variability in binary and continuous choice*. Manuscript. University of California, Santa Cruz. *Psychonomic Bulletin and Review*, to appear.

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

Gigerenzer, G., & Hoffrage, U. (1995). How to improve bayesian reasoning without instruction: Frequency formats. *Psychological Review*, **104**(4), 684–704.

Gluck, M. A., & Bower, G. H. (1988). From conditioning to category learning: An adaptive network model. *Journal of Experimental Psychology*: *General*, **117**, 225–244.

Kelley, H., & Friedman, D. (1998). *Learning to forecast price*. Working paper #394, Economics Department, University of California, Santa Cruz.

Kivenen, J., & Warmuth, M. K. (1995). Additive versus exponentiated gradient updates for linear prediction. In *Proc. 27th Annual ACM Symposium on Theory of Computing*, pp. 209–218.

Kruschke, J. (1992). ALCOVE: An exemplar-based connectionist model of category learning. *Psychological Review*, **99**(1), 22–44.

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

Marcet, A., & Sargent, T. (1989). Convergence of least squares learning mechanics in self referential linear stochastic models. *Journal of Economic Theory*, **48**, 337–68.

Massaro, D. W. (1987). *Speech perception by ear and eye*: *A paradigm for psychological inquiry*. Hillsdale, NJ: Earlbaum.

Massaro, D. W. (1989). Testing between the TRACE model and the fuzzy logical model of speech perception. *Cognitive Psychology*, **21**, 398–421.

Massaro, D. W. (1994). A pattern recognition account of decision making. *Memory & Cognition*, **22**, 616–627.

Massaro, D. W. (in press). *Perceiving talking faces*: *From speech perception to a behavioral principle*. Cambridge, MA: MIT Press.

Massaro, D. W., & Friedman, D. (1990). Model of integration given multiple sources of information. *Psychological Review*, **97**(2), 225–252.

McClelland, J. L., & Rumelhart, D. E. (1986). Distributed memory and the representation of general and specific information. *Journal of Experimental Psychology*, **114**(2), 159–188.

Medin, D. L., Altom, M. W., Edelson, S. M., & Freko, D. (1982). Correlated symptoms and simulated medical classification. *Journal of Experimental Psychology*: *Learning*, *Memory*, *and Cognition*, **8**, 37–50.

Minsky, M. L., & Papert, S. A. (1988 expanded edition; original edition 1969). *Perceptrons*. Cambridge, MA: MIT Press.

Pullum, G. K. (1996). Learnability, hyperlearning, and the poverty of the stimulus. Paper presented at the parasession on learnability. In *22nd annual meeting of the Berkeley Linguistics Society*, *Berkeley*, *CA*.

Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In A. H. Balck & W. F. Prokasy (Eds.), *Classical conditioning II*: *Current theory and research*, pp. 64–99. New York: Appleton–Century–Crofts.

Rosenblatt, F. (1958). The perceptron: A probabilistic model for information storage and organization in the brain. *Psychological Review*, **65**, 386–408.

Rumelhart, D. E., & McClelland, J. L. (Eds.) (1986). *Parallel distributed processing*, *Volume 1*: *Foundations*. Cambridge, MA: MIT Press.

Selfridge, O. G. (1959). Pandemonium: A paradigm for learning. In *Mechanization of thought processes*, pp. 511–526. London: H. M. Stationery Office.

Selten, R. (1996). Axiomatic characterization of the quadratic scoring rule. University of Bonn Discussion Paper No. B-390. *Experimental Economics*, to appear.

Shanks, D. R. (1991). A connectionist account of base-rate biases in categorization. *Connection Science*, **3**, 143–162.

Zadeh, L. A. (1965). Fuzzy sets. *Information and Control*, **8**, 338–353.

Received February 6, 1998