JOURNAL OF APPLIED BEHAVIOR ANALYSIS 1992, 25, 585–597 NUMBER 3 (FALL 1992)

GOING BEYOND WITH BAYESIAN UPDATING

BRENDAN O'FLAHERTY
COLUMBIA UNIVERSITY

AND

JUDITH L. KOMAKI
BARUCH COLLEGE, CUNY

It all started innocently enough. One spring-like winter day, I happened to ask Brendan whether economists ever dealt with escalation.

"With what?" he replied.

"A phenomenon where people keep investing in the face of continuing losses," I answered.

Then I described how industrial/organizational psychologists had become intrigued with situations in which investors seemed to throw good money after bad, how their explanations for the phenomenon centered on individual characteristics such as commitment, how Sonia Goltz (1992) used a standard bread-and-butter operant procedure—fixed and variable schedules of reinforcement—to explain their persistence, and how Goltz's experiments had shown that during the extinction phase investors even increased their investments for a while when the news was all bad.

Without a moment's hesitation, he exclaimed, "I bet I can predict the turning point."

Another arrogant-economist remark, I thought to myself. "How?"

"Bayesian updating."

Then we talked at length about Bayesian analysis techniques and how they could be used to predict the shape of extinction curves.[1] I realized that these techniques might be just what psychologists needed as an enticement to study sequences of behavior over time.

And that's how this commentary got started.

—JLK

So What Is Bayesian Updating?

Bayesian updating is both a predictive tool and a normative philosophy of decision making.

A predictive tool. As an analytical technique, Bayesian updating allows us to forecast how people will act on a trial-by-trial basis and to predict the turning point at which persons who have escalated their investments begin to taper off. For example, we can predict whether subjects in experiments (e.g., Goltz, 1992; Hantula, 1990) will invest more or less on Trial 9 than on Trial 8. Then, taking into consideration the information they obtain on Trial 9, we can predict whether subjects' investments will go up or down on Trial 10, and so on for each trial.

We are interested primarily in using Bayesian updating as an analytical technique to predict how persons will act. To make these predictions, we figure out what beliefs are warranted for the subjects (*warranted beliefs*), given the information they have received, and what probabilities are warranted by these beliefs (*warranted probabilities*). Then we use these warranted probabilities to predict what the subjects will do. Knowing the direction of

This research was partially supported by the U.S. Army Research Institute contract MDA903-91-K-0136 to the second author. The views, opinions, and findings contained in this paper are those of the authors and should not be construed as the official position of the U.S. Army Research Institute or as an official Department of the Army position, policy, or decision.

Send correspondence and reprint requests to Brendan O'Flaherty, Department of Economics, Columbia University, New York, New York 10027.

[1] In this paper, we do not forecast the level of the extinction curve, but more sophisticated extensions of Bayesian methods, including utilities (i.e., the expected payoff from investing or not investing) could do so. For instance, see the methods employed in O'Flaherty and Siow (1990).

changes in warranted probabilities will tell us the direction of changes in persons' actions.

A normative philosophy of decision making. Bayesian updating stems from a tradition dating from Port Royal logicians of 1660 and Bernoulli's *Ars Conjectandi* of 1713. For some researchers, this tradition provides a normative philosophy about the way information should be used to make decisions. Scientists are the exemplar. A scientist systematically gathers information about the situation at hand and then uses this information to calculate an optimal course of action. Once this course of action has been determined, the scientist behaves accordingly. Given new information, the scientist updates his or her optimal course of action and then adapts his or her behavior. Life, for a Bayesian, is a series of never-ending discoveries with fresh updates. Bayesian decision makers are constantly on the lookout for information. The information is not static and confined to a given trial. Rather the information is obtained on many trials over a period of time. Whatever information they receive, Bayesian decision makers use something approximating mathematical formulas to aggregate this cacophony of information to update their beliefs. This is the process Bayesians believe people should use in making decisions, *and* the results of this process are the decisions that Bayesians believe people should make. Thus, this normative philosophy covers both the process of deciding and the actual decisions themselves.

Relationship with Bayes' theorem. Bayesian updating gets its name from a famous theorem in probability theory proved by Thomas Bayes in 1763. This theorem provides the technical tools to perform the updating done by Bayesian decision makers. The use of Bayes' theorem is by itself not a distinctive component of Bayesian updating. Bayes' theorem follows from the standard definition of probability (Kolmogorov, 1950), so anyone who uses probabilities implicitly accepts Bayes' theorem. The Bayesian tradition, however, uses Bayes' theorem in a special way that is not accepted by all users of probability.

What differentiates Bayesians most distinctly from other users of probability is that they are willing to attach probabilities to propositions that others would not attach probabilities to. Consider the proposition, "Mikhail Gorbachev was born on a Thursday," as an example. A Bayesian who didn't know Gorbachev well would argue that the probability is a little bit over one seventh (because fewer births happen on the weekend). Some non-Bayesians think this position ridiculous, holding that Gorbachev has already been born and was either born on Thursday or not. Thus, the proposition is either true or false and not, as a Bayesian would believe, something in between.

Bayesians go so far as to attach probabilities to propositions about probabilities, for example, in defining warranted beliefs. A proposition about probabilities—"The probability is one seventh that Gorbachev was born on a Thursday"—can also have a probability attached to it. Call the quoted proposition "Proposition P." Suppose we knew there was a 1% probability that Stalin had instituted a bonus system for doctors who had births on Thursdays, and if such a system had been implemented that one third of all births would have taken place on Thursdays. Then we could say that the probability is 99% that Proposition P is true and 1% that it is false, a probability about a probability.

Bayesian Notions in the Literature

The Bayesian tradition has been a source of controversy in the psychology and decision-theory literature for many years. For instance, the journal *Theory and Decision* publishes frequent discussions about whether people are or should be Bayesian decision makers. Kahneman and Tversky (1982) are well known for arguing that people do not act in a Bayesian manner. Numerous experiments have been conducted to find out whether they do or not in many different contexts, for example, in criminal trials (Faigman & Baglioni, 1988) and in laboratory experiments (Birnbaum & Mellers, 1983; Geller & Pitz, 1968, 1970).

The Bayesian tradition has also influenced statisticians and inspired nonclassical techniques of parameter estimation. These techniques have been especially prominent recently in item-response theory (e.g., Frederiksen, Mislevy, & Bejar, in press; Lim & Drasgow, 1990).

These bodies of literature, however, tend to deal with decision problems that are almost static in time (all the information arrives effectively at the same time). Nontrivial uses of updating, like Goltz's (1992), require different pieces of information to be learned at different times. This type of updating, too, is well represented in the literature.

In economics. Practically all economists are Bayesians, and Bayesian updating is the established tradition in economics. This idea underlies, for instance, the efficient-markets hypothesis in finance (e.g., Fama, 1965; Malkiel, 1985; Samuelson, 1965); Bayesian mechanism-design theory in game theory (e.g., Myerson, 1985); rational expectations macroeconomics (e.g., Lucas, 1976; Muth, 1961; Sargent & Wallace, 1976); discussions of the "peso problem" (how traders react to the possibility of a currency devaluation) in international monetary theory (e.g., Lewis, 1988, 1989). Katz (1986), a labor economist, treated laid-off workers as Bayesians who get information and revise over time their beliefs about the probability that they will be recalled to their old jobs; he predicted their actions from this process.

In psychology. Bayesian updating models have helped psychologists understand a variety of different phenomena, including group polarization (Bordley, 1983), the formation of delusions (Hemsley & Garety, 1986), the stress of waiting to be served (Osuna, 1985), the aggregation of information from experts (Mendel & Sheridan, 1989), visual signal detection (Burgess, 1985), recidivism rates of drunk drivers (Hauer, 1983), and mating decisions of amphipods (Hunte, Myers, & Doyle, 1985).

In helping decision makers. A number of researchers have developed software to help decision makers be "better" Bayesians in clinical diagnosis using the MMPI (Hsu, 1988), marketing (Chatterjee, Eliashberg, & Gatignon, 1988) and military tactical intelligence (Adelman, Donnell, Phelps, & Patterson, 1982). The artificial intelligence literature (e.g., Garbolino, 1987) also considers the uses of Bayesian updating methods.

How Can Bayesian Inference Be Used to Predict Decisions?

Take a look at Figure 1. Subjects progress from Trial 1 to Trial 2 and so on. They start off each trial with *initial* or *prior beliefs* about the probability that the investment will pay off. On each trial, they have an opportunity to invest or not. They also receive information about the profitability of the investment in that period. This information is called a *realization.* The realization is used in *updating* their beliefs, with Bayes' theorem providing the technical means for this updating. The updated beliefs become the prior beliefs for the next trial, and the process repeats. The beliefs that result from carrying out these steps properly are called *warranted beliefs* (Keynes, 1921). The warranted beliefs that subjects have at any time imply a probability that the investment will be successful. This probability is called a *warranted probability.* On the basis of this warranted probability, subjects either invest or not.

Four steps are used in a Bayesian analysis: (a) select a model that reflects what subjects see, (b) define the "situation," (c) calculate warranted beliefs, and (d) calculate warranted probabilities. The first two steps are needed to determine what sort of information subjects use in formulating their warranted beliefs.

Select a model that reflects what subjects see. The first step is to select a model that a subject could have about the data he or she sees. There are many models to choose from. One of the most powerful, simple, and common models is the x-step Markov process, in which the probability of an outcome on each trial depends on the previous x realizations. For instance, in a one-step Markov process, each trial's outcome is a function solely of the previous realization. Two good examples are the alternation of night and day and the alternation in the business cycle of boom and bust. In some work, the term *Lag 1 probability* is used to describe a probability in a one-step Markov process (e.g., Komaki & Citera, 1990).

The zero-step Markov process has been used by psychologists in describing decision making and superior–subordinate interactions (e.g., Edwards,

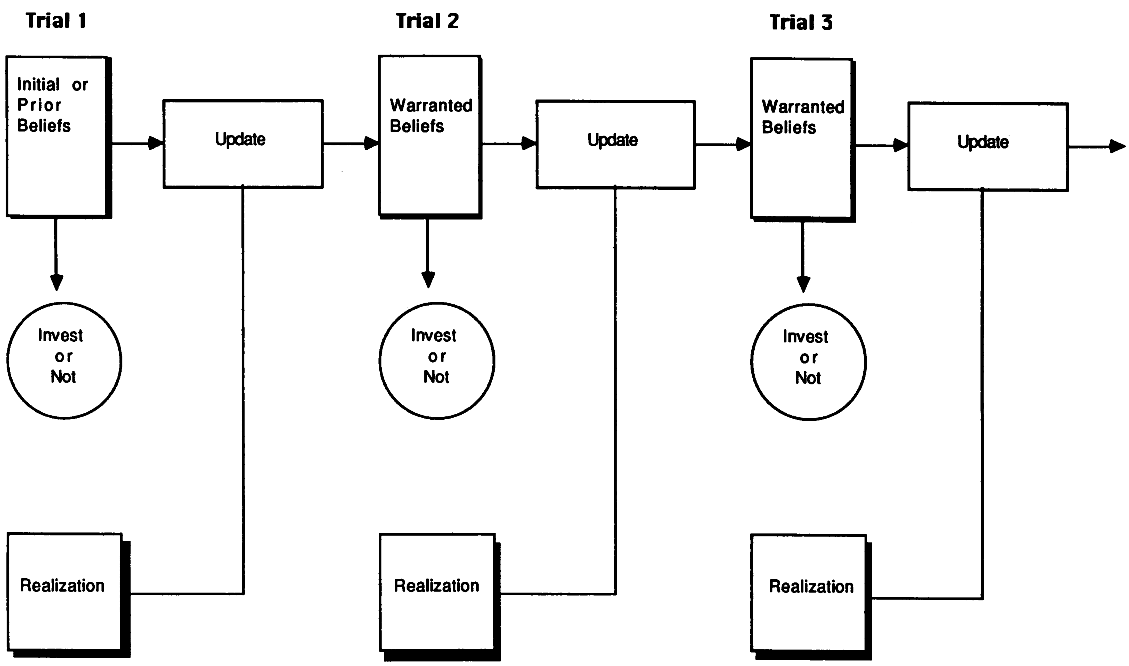

Figure 1. A schematic diagram of sequential Bayesian updating and decision making.

1955; Geller & Pitz, 1968, 1970; Komaki & Citera, 1990). In these processes, the probability of each outcome is not affected by the previous outcome or by any truncated sequence of previous outcomes. Usually probabilities in zero-step Markov processes are called *expected* or *unconditional* probabilities.

There are other models too. For instance, in "stock-flow" models, stocks of carbon dioxide, gasoline, good will, or bad luck gradually accumulate or decumulate. If one models subjects as holding the gambler's fallacy that good luck is bound to follow a string of bad luck, then one could use a stock-flow model to reflect this process.

The models that we have chosen for the Goltz (1992) experiments are the two- and three-step Markov processes, because these are the models that allow subjects as scientists to predict realizations most accurately. In practical terms, we used a rough goodness-of-fit criterion—in this case, the fit between the previous and future returns on investment. (This is the same process as Komaki and Citera, 1990, used more formally to reject the zero-step model—expected probability—in favor of the one-step model—Lag 1 probability.)

Define the situation. Once we decide on a model, the definition of a *situation* follows. For instance, with a two-step Markov process, the situation is the previous two realizations, in order. In the Goltz (1992) experiment, each realization can be categorized as either good (g)—a gain of $30 per $100 invested—or bad ($b$)—a loss of $10 per $100 invested. If we think about a situation as a two-step history, then there are four possible situation: (a) both previous realizations might be good (gg), (b) both might be bad (bb), (c) a good realization might precede a bad (gb), or (d) a bad might precede a good (bg).

Calculate warranted beliefs. Warranted beliefs are the "best" possible beliefs you can have about what is going on, given the information you have at the moment. (For a discussion of the still ongoing debate about what constitutes the "best" possible belief, refer to Levi, 1991.) For example, given the information we have about New York highways, the shortest drive on highways between Manhattan and Saratoga Springs is on Route 87. Some travelers may actually believe that the shortest drive is on Route 9 rather than on Route 87. Actual beliefs are those travelers have, whereas warranted ones are those they *should* have. Bayesians assume travelers will make inferences optimally, and figure out the very shortest drive.

Calculating warranted beliefs uses the beliefs warranted at the beginning of the experiment (initial beliefs), the information the subject receives on each trial (realization), and the process of updating.

In Goltz's (1992) experiments, warranted beliefs were calculated as a beta distribution. (To see how the warranted beliefs were calculated in the Goltz experiment, refer to the appendix available from the first author.) This distribution is typically used for processes like Goltz's, for which the information received in each trial is binary (e.g., in O'Flaherty & Siow, 1990). The use of beta distributions is based on a classical principle of epistemology, the principle of "insufficient reason"; on Bayes' theorem; and on the requirement that the probabilities be between zero and one. The advantage of calculating warranted beliefs this way is that they take on a very simple form.

Calculate warranted probabilities. The formula for warranted probability is $(s + 1)/(n + 2)$, where n stands for the number of times you have observed a particular situation before and s stands for the number of times the next realization after that particular situation was good. As n gets large, the warranted probability approaches the empirical frequency—the actual proportion of times a good realization followed a particular situation. A good realization increases both the numerator and denominator by one; a bad realization increases only the denominator.

This formula for warranted probabilities is a short answer to the question of how Bayesian inference can be used to predict decisions. To match the Goltz (1992) data, which are dollar amounts invested, we make the reasonable assumption that the amounts invested are positively related to warranted probabilities.

Table 1
Data for Calculating Warranted Probabilities with Two-step Histories

Two-step history	Number of occurrences (n)	Number of good realizations following it (s)
gg	2	0
gb	4	2
bg	5	3
bb	3	2

Making Predictions

With a two-step Markov process. Let's illustrate how to make predictions using a two-step process. Consider Goltz's (1992) 16-trial partial-variable acquisition phase[2] in terms of good and bad realizations: *b g g b b b g b g g b b g b g g*. We can go through this phase, find out the number of times each two-step history has occurred, and count the number of times a good realization has followed each two-step history.

The data are shown in Table 1. The last two realizations of the acquisition phase are both good, so at the start of the extinction phase the situation is *gg*. Twice this situation has been encountered before, and both times the realization following this history was bad. So $s = 0$, $n = 2$, and using the formula, the warranted probability of a good realization for the first extinction trial is $(s + 1)/(n + 2) = (0 + 1)/(2 + 2) = 1/4$.

On the second trial of the extinction phase, the history is *gb*. From Table 1, $s = 2$, $n = 4$, so at

[2] Technically speaking, Goltz does not use a classic operant paradigm. In her fixed and variable schedules of reinforcement, subjects did not receive reinforcement on every nth response. Rather, on each trial, subjects chose to respond or not and received good or bad news about investment outcomes, whether or not they invested. In mathematical terms, the Goltz paradigm presents a simple discrete-time Bayesian updating problem, whereas the classic operant experiment would present a noticeably more difficult, but not impossible, problem—the so-called two-armed bandit problem in continuous time (Gittins & Jones, 1974).

Table 2
Predictions with a Two-step History

Trial in extinction phase	Warranted probability of a good realization
1	1/4
2	1/2
3	3/5
4	1/2
5	3/7
6	3/8
7	1/3
8	3/10
9	3/11
10	1/4
11	3/13
12	3/14

Table 3
Data for Calculating Warranted Probabilities with Three-step Histories

Three-step history	Number of occurrences (n)	Number of good realizations following it (s)
ggg	0	0
ggb	2	0
gbg	2	2
gbb	2	1
bgg	2	0
bgb	2	2
bbg	2	0
bbb	1	1

this trial the warranted probability is $(s + 1)/(n + 2) = (2 + 1)/(4 + 2) = 3/6 = 1/2$. The warranted probability goes up from the first trial to the second. This gives us our first prediction: Investment should go up from the first extinction trial to the second.

On the third trial of the extinction phase, the history is *bb*. From Table 1, $(s + 1)/(n + 2) = (2 + 1)/(3 + 2) = 3/5$. Once again, warranted probability goes up, so we have another prediction: Investment should also go up from the second extinction trial to the third.

The fourth trial requires updating from Table 1. That's because the history at the third extinction trial was *bb* and the realization was bad. So on the fourth trial, history *bb* has occurred four times ($n = 4$) with good realizations twice ($s = 2$). Thus, warranted probability is $(s + 1)/(n + 2) = (2 + 1)/(4 + 2) = 3/6$. Now warranted probability is going down, so we predict that investment will fall from the third to the fourth trial.

Proceeding in this fashion, we can compile a list of warranted probabilities throughout the extinction phase (Table 2). This sequence replicates many features of Goltz's (1992) data (refer to her Figure 2). Like hers, it increases to the third trial and tapers off thereafter; so the sequence predicts the turning point. As with Goltz's data, the second and fourth trials are about the same. These warranted probabilities differ from Goltz's data, however, in that the bump at Trial 8 doesn't appear, and the first trial is lower than the fifth. In short, two-step histories, although not perfectly in line, give a fairly accurate prediction of the pattern of investment subjects make when they have gone through Goltz's particular partial-variable acquisition and extinction phases.

Making predictions with three-step Markov processes. An alternative approach is to use three-step Markov processes. This approach is very attractive because, as shown in Table 3, the fit during the acquisition phase is extremely good: The three-step history is almost a perfect predictor of the next realization. For almost every history, the realization that followed it was always the same. For instance, whenever the history *ggb* occurred, the next realization was bad; whenever *gbg* occurred, the next

Table 4
Predictions with a Three-step History

Trial in extinction phase	Warranted probability of a good realization
1	1/4
2	1/4
3	1/2
4	1/3
5	1/4
6	1/5
7	1/6
8	1/7
9	1/8
10	1/9
11	1/10
12	1/11

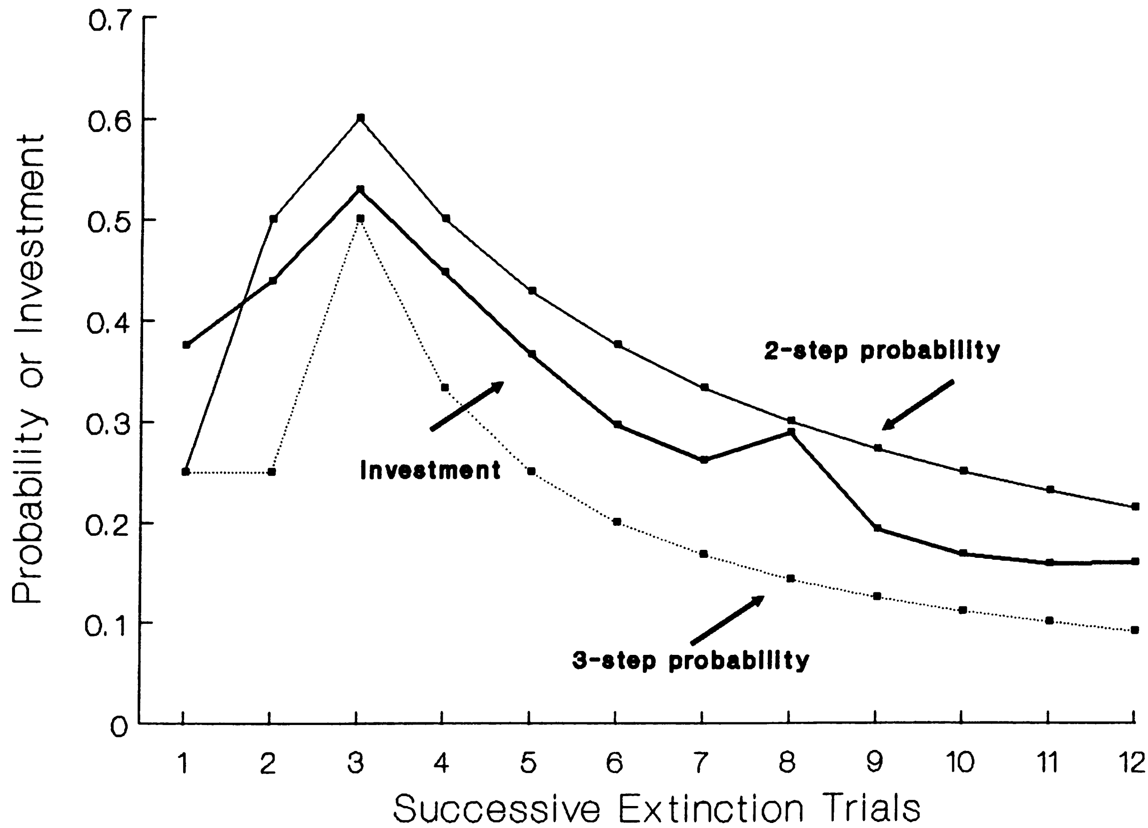

Figure 2. Comparison of warranted probabilities and investments.

realization was good. Therefore, almost always in the acquisition phase, the realization is precisely determined by the three-step history. Such excellent predictive power is one reason why three-step histories might be attractive.

Proceeding in the same way we did with the two-step process, we can construct a table of warranted probabilities, as shown in Table 4. Thus, three-step histories predict the turning point just as well as two-step histories. Although three-step histories predict the relationship between Trials 2 and 4 a little worse, they predict the relationship between Trials 1 and 5 much better than two-step histories. Once again, the Bayesian method gives reasonably good predictions.

The results for the two- and three-step histories are graphed in Figure 2. Figure 2 plots trials of the extinction phase on the abscissa and both warranted probability and investments from Goltz (1992) on the ordinate. To make investment comparable, we divided the investment by $12,000; obviously this normalization was arbitrary. Figure 2 gives a visual demonstration that the two-step and three-step histories predict the pattern of investment reasonably well.

Note that these techniques can also be applied to partial-fixed schedules. Using two-step histories gives warranted probabilities that agree with the Goltz (1992) data in being monotonically declining. They also agree with her data in that the partial-variable probabilities are higher than partial-fixed probabilities on the third and all subsequent trials and lower on the first. However, on the second trial, Goltz's data show less investment with partial-fixed schedules, but warranted probabilities are the same (two-step) or lower (three-step). Except for this small discrepancy, which can be resolved with more sophisticated approaches, Bayesian methods also work reasonably well with this type of acquisition schedule.

IMPLICATIONS

A Grasp of the Larger Context

By seeing how Bayesian analysis techniques can be used to predict behavior, we can better see how actions are a function of the larger context of an evolving sequence of realizations rather than of the realization on any one trial. The two or three immediately preceding realizations are in some sense "most important" because they indicate what situation applies. However, the derivation of warranted probabilities can bring the entire previous history of the subject involved to bear on any particular trial. This is because warranted probabilities are calculated using n, the number of times the particular situation occurred before, and s, the number of times a good realization followed that situation. In this manner, all of the past affects the present.

Glimpsing the Rationality Behind Seemingly Irrational Actions

By looking at the larger context, we can see how subjects systematically use the information they have received. Although their behavior might initially appear to be in error to someone who takes a shortsighted view and concentrates only on the immediately preceding realizations, we can actually see how "rationally" they are behaving in the larger context.

The phenomenon of escalation as a decision error may in fact exist (Brilmayer, 1983; Kahneman & Tversky, 1983). Be that as it may, the people in Goltz's (1992) experiments did not make this error on average. In this sense, her experiments support Bowen (1987), who distinguishes between decisions as dilemmas and decisions as errors.

Bayesians would not recommend training to prevent escalation of this nature. What is most startling about Goltz's (1992) subjects is how closely they seemed (on average at least) to abide by the precepts of Bayesian inference. Given the information Goltz's subjects had, no one—not even the world's most able statistician with an arbitrarily long period for reflection—could have done better. Short of giving them foreknowledge, nobody could have taught them anything (on average at least). Hence, recommendations for training to prevent escalation are not supported by this experiment. (An exception might be made for individual outliers who egregiously fail to reason in the Bayesian manner; such outliers might benefit from training in Bayesian inference.)

FUTURE DIRECTIONS

Incorporating Time in the Explanans *and* Explanandum

Although the field of applied behavior analysis often contains experiments in which the *explanandum* (that which is to be explained) is a sequence of events over time, almost always the *explanans* (that which explains it) is a single event. The most common question raised is whether or not the principle of positive reinforcement can be successfully applied with yet another set of subjects on yet another behavior in yet another setting. This is particularly true of the work in organizational behavior management (Komaki, Coombs, & Schepman, 1991). Fox, Hopkins, and Anger (1987), for example, showed that a token economy program was successful in reducing mining accidents. Although this study is noteworthy because of the subjects and setting (underground coal miners in two mines in Arizona and Wyoming) and the length of the interventions (from 11 to 12 years), it nonetheless did not look at how miners' accident rates varied over time or at the reinforcement schedule. Thus, it was a joy to see Goltz's (1992) work in which both the *explanandum* (sequence of investments over 12 extinction-phase trials) and the *explanans* (acquisition schedule of reinforcement over 16 trials) were sequences of events.

A time series as explanandum. Bayesian updating techniques open up for study a number of research topics that have previously eluded psychologists. Rarely do psychologists look at data recorded on a daily, weekly, monthly, or yearly basis. In fact, one-shot pretests and posttests are the norm. Many possibilities to use a time series as explanandum exist (McGrath & Kelly, 1986): plateaus in learning, procrastination, changes in the strategies of groups over project life cycles (Gersick, 1988), memory failure, and changes in the performance of individuals over time (Hoffman, Jacobs, & Gerras, 1992).

A time series as explanans. It has been even rarer to see a time series as *explanans*. This, too, should change. We could look at patterns of training over an employee's tenure at a firm, schedules of disciplinary actions, timing of pay raises, or arrangement of negative consequences. We could even think of an individual's life history as an *explanans*, as Skinner (1984) does in his autobiography: "I have learned to accept my mistakes by referring them to a personal history. . . . My behavior at the Royal Society dinner, for example, was, to say the least, unfortunate, but I could reflect that my early life was very different from that of Lord Adrian's or most of the other guests'" (pp. 407–408).

If we could chart a person's history, we could better predict what he or she might do. Let's say we wanted to predict who would be likely to collaborate in writing an article like this one. Taking a Bayesian approach, we could identify and chart similar events. For example, we might chart (a) activities involving unusual combinations such as a PhD candidate running an election campaign or a theoretical economist studying homelessness, (b) rare and time-consuming activities such as running 50-mile ultra marathons, and (c) opportunities taken to teach mathematical topics to neophyte audiences. We could tally these activities for the 20-year period preceding the possible collaboration. An informal tally from 1972 through 1991 for one of the authors showed a frequency that looked like this: 1 2 1 2 2 0 2 4 4 2 2 5 4 3 8 4 3 7 11 2. An individual who had frequently had good experiences with similar undertakings, and rarely had had bad ones, would be likely to try a collaboration like this.

Getting Back to the Individual

Although only aggregate data were used in the present case, the Bayesian approach can also predict behavior on an individual-by-individual basis. To do so, we need to change our assessment of initial beliefs. In using the principle of insufficient reason to establish initial beliefs in the aggregate, we took a rather extreme position within the Bayesian community. Even among users of warranted probability, such an application is not entirely uncontroversial; it implies that all subjects enter the experiment entirely devoid of relevant prior experience, either with stocks or experiments.

The reason we used the principle was to show that even with one hand tied behind our backs, we could make interesting aggregate predictions. Researchers who relax the assumption will be able to analyze individual data by identifying an individual subject's initial beliefs from his or her behavior in the acquisition phase and then predicting during the extinction phase each individual's behavior. This might be an exciting project for technically trained individuals familiar with Bayesian techniques. For researchers who do not want to worry about initial beliefs, there is also a way to go. No matter what initial beliefs are, warranted probabilities approach empirical frequencies if subjects have enough data. Researchers who do not want to worry about initial beliefs should construct very long acquisition phases, considerably longer than the 16 Goltz (1992) used.

INTERACTIONS AMONG DISCIPLINES

Issues in Learning

Testing sequential effects in resistance to extinction. Bayesian methods can be used to test hypotheses in learning. Capaldi (1966), for example, posited that "sequential variables" are the major determinants of resistance to extinction. These variables can be conveniently described as properties of Markov processes. Using the Bayesian approach, investigators can come closer to specifying "the exact magnitude of the response on such and such a trial," what Capaldi (1966) refers to as "the ultimate aim of behavior theory" (pp. 472–473).

Predicting peaks in extinction curves. Bayesian methods provide a different way to describe extinction curves. Killeen (1982), for instance, developed an exponentially weighted moving average (EWMA) analysis which, he showed, implies that in certain kinds of experiments extinction curves should be logistic curves. Killeen's analysis, like ours, accounted for patterns of behavior occurring over time and considered at each time the full history of the experiment up to that point. The EWMA analysis was not entirely satisfactory in his judgment (P. R. Killeen, personal communication, March 17, 1992). It had trouble with the "transition state" at the beginning of the extinction phase when the responses escalated. The EWMA prediction is a slow but monotonic tapering off of responses. With some nontrivial modifications, Bayesian methods can be used to forecast behavior during these transition states. The actual procedure will have to be fairly complex, because subjects in classic operant experiments determine by their responses how much information they get.

Issues in Industrial/Organizational (I/O) Psychology

Looking at escalation as well as persistence. In the I/O psychology community, there is considerable interest in how long persons will persist and when they will escalate (Staw & Ross, 1987). Unfortunately, virtually all of the research has examined persistence. Using Bayesian procedures, one can start to look at when and for how long persons will escalate. (One can also use approaches taken by others, e.g., Wallsten, 1976, in a non-Bayesian framework.)

Identifying impact of the timing of losses. One can also begin to test the influence of the timing of losses (Brockner, 1992). Staw and Ross (1989) suggested that "if economic losses are large and they occur early in a project's life cycle, withdrawal may well be the dominant response. However, if losses do not appear until later in the process, then persistence could be the typical response" (p. 219). They suggest, based on a study by Brockner and Rubin (1985), that bad news might have more of an impact when it is introduced early in the cycle rather than later when resources have already been committed. Both hypotheses can be tested using Bayesian methods, but to a Bayesian, these predictions are obvious. A Bayesian considers n—how much learning has taken place. When n is small (i.e., early in the cycle), one bad realization causes a big change in the warranted probability. When n is large, one bad realization has only minor impact.

Assessing impact of individual differences. Not all subjects undergo acquisition phases from which we can make inferences about their warranted probabilities. Instead, some are thought to have a certain amount of "inertia" (e.g., Geller & Pitz, 1968). In related experiments, subjects were allowed to choose whether to become involved in a project or not, and these subjects were thought to be "committed" to that particular project (e.g., McCain, 1986). The subjects who chose a project initially were those who held more favorable initial beliefs. These subjects' posterior beliefs continued to be more favorable than those held by people who chose to reject a project or those who might have been randomly assigned; thus, they were likely to continue investing in the project longer. In these cases, we need to relax the assumption that all subjects enter the experiment with the same initial beliefs and to deal with an issue economists refer to as "unobserved heterogeneity" (for an example, refer to Heckman & Willis, 1977).

Issues in Economics

Choosing a model. Perhaps the most far-reaching research direction is the possibility of using the Bayesian approach to construct a priori a theory of the dynamics of behavior (Killeen, 1982). What we did with the Goltz (1992) data was somewhat ad hoc. Although we showed that a Markov three-step process fits Goltz's extinction curve data, we did this by trial and error. For economists, a major task still to be done is to find for each acquisition phase a mathematical model that describes how subjects process data during the acquisition phase and gives good predictions of how they will act.

Shocks and government intervention. The psychological methods used by Goltz (1992) and her resulting findings can also be used profitably to address some of the deepest, most important, and most controversial issues in economics, like "rational expectations" in macroeconomics. (For ideas from economics to psychology, see Hursh, 1984, and Kagel & Winkler, 1972.)

Before rational expectations writers appeared on the scene in the mid 1970s (Lucas, 1976; Sargent & Wallace, 1976), the Keynesian orthodoxy in macroeconomics held that "shocks"—unexpected changes to the economic environment, such as oil embargoes, new diseases, amazing inventions, or possibly waves of pessimism or optimism—could cause unemployment or recessions while people adjusted their wage and price demands to the new levels appropriate to the new conditions. This adjustment process could be long, painful, and needlessly wasteful of resources, but governments could and should ease adjustment by manipulating the money supply and possibly government spending.

Rational expectations writers challenged this tradition by arguing that people were at least as smart as governments and economists and thus could adjust at least as well without assistance as with it. In particular, if governments could figure out the prices and wages to which the economy would eventually converge after the shock, then economic actors could too, and so should go to these prices and wages immediately without any intervening period of pain and unemployment. Therefore, government intervention is at least unnecessary.

If governments did intervene, moreover, according to some rule like "increase money supply when unemployment is high," people would be able to see this policy, predict it, and anticipate its implementation when they set wages and prices. But anticipated policies can have no impact.

To see this, suppose, for instance, the government followed a policy of increasing money supply by 5% whenever unemployment was high, in the hope of reducing real wages and prices and reducing unemployment. If people knew what the government was going to do, they would raise their wages and prices by 5% in anticipation, and the policy would be ineffective. In the rational expectations view, people's ability to figure out what the world is like makes government intervention both unnecessary and ineffective.

Whether rational expectations gives a correct view of the economy—and therefore whether government intervention makes sense—depends on how, and how well, people can learn about changing circumstances and adjust to them. Economists have studied this question in some depth (e.g., Bray & Kreps, 1986; Frydman & Phelps, 1983). Goltz's (1992) experiments give some direct evidence on this fundamental and very relevant question for economics. How much government intervention is appropriate depends on how well economic actors can figure out on their own that shocks have occurred. Going from an acquisition phase to an extinction phase is a shock. Goltz is looking at how well people learn about shocks. Economists have much to learn from this line of research. In short, Bayesian analysis can tell us a lot about how people responded in Goltz's experiments. More importantly, it can reveal many important directions for future research.

REFERENCES

Adelman, L., Donnell, M. L., Phelps, R. H., & Patterson, J. F. (1982). An iterative Bayesian decision aid: Toward improving the user-aid and user-organization interfaces. *IEEE Transactions on Systems, Man and Cybernetics,* **12,** 733–743.

Bernoulli, J. (1899). *Ars conjectandi.* Leipzig: Engelman. (originally published in 1713)

Birnbaum, M. H., & Mellers, B. A. (1983). Bayesian inference: Combining base rates with opinions of sources who vary in credibility. *Journal of Personality and Social Psychology,* **45,** 792–804.

Bordley, R. F. (1983). A Bayesian model of group polarization. *Organizational Behavior and Human Performance,* **32,** 262–274.

Bowen, M. G. (1987). The escalation phenomenon reconsidered: Decision dilemmas or decision errors? *Academy of Management Review,* **12**(1), 52–66.

Bray, M. M., & Kreps, D. (1986). Rational learning and rational expectations. In W. Heller, R. Starr, & D. Starrett (Eds.), *Essays in honour of K. J. Arrow.* Cambridge, England: Cambridge University Press.

Brilmayer, L. (1983). Discrepancies between human behavior and formal theories of rationality: The incompleteness of Bayesian probability logic. *Behavioral and Brain Sciences,* **6,** 488–489.

Brockner, J. (1992). The escalation of commitment to a failing course of action: Toward theoretical progress. *Academy of Management,* **17**(1), 39–61.

Brockner, J., & Rubin, J. (1985). *Entrapment in escalating conflicts.* New York: Springer-Verlag.

Burgess, A. (1985). Visual signal detection: III. On Bayesian use of prior knowledge and cross correlation. *Journal of the Optical Society of America-A,* **2,** 1498–1507.

Capaldi, E. J. (1966). Partial reinforcement: A hypothesis of sequential effects. *Psychological Review,* **73,** 459–477.

Chatterjee, R., Eliashberg, J., & Gatignon, H. (1988). A practical Bayesian approach to selection of optimal market testing strategies. *Journal of Marketing Research,* **25,** 363–375.

Edwards, W. (1955). The prediction of decisions among bets. *Journal of Experimental Psychology,* **50,** 201–214.

Faigman, D. L., & Baglioni, A. J. (1988). Bayes' theorem in the trial process: Instructing jurors on the value of statistical evidence. *Law and Human Behavior,* **12,** 1–17.

Fama, E. (1965). The behavior of stock market prices. *Journal of Business,* **38,** 34–105.

Fox, D. K., Hopkins, B. L., & Anger, W. K. (1987). The long-term effects of a token economy on safety performance in open-pit mining. *Journal of Applied Behavior Analysis,* **20,** 215–224.

Frederiksen, N., Mislevy, R., & Bejar, I. I. (Eds.). (in press). *Test theory for a new generation of tests.* Hillsdale, NJ: Erlbaum.

Frydman, R., & Phelps, E. S. (Eds.). (1983). *Individual forecasting and aggregate outcomes.* Cambridge, England: Cambridge University Press.

Garbolino, P. (1987). Bayesian theory and artificial intelligence: The quarrelsome marriage. *International Journal of Man Machine Studies,* **27,** 729–742.

Geller, E. S., & Pitz, G. (1968). Confidence and decision speed in the revision of opinion. *Organizational Behavior and Human Performance,* **3,** 190–201.

Geller, E. S., & Pitz, G. F. (1970). Effects of prediction, probability, and run length on choice reaction speed. *Journal of Experimental Psychology,* **84,** 361–367.

Gersick, C. J. G. (1988). Time and transition in work teams: Toward a new model of group development. *Academy of Management Journal,* **31,** 9–41.

Gittens, J. C., & Jones, D. M. (1974). A dynamic allocation index for the design of experiments. In J. Gani (Ed.), *Progress in statistics* (pp. 241–266). Amsterdam: North Holland.

Goltz, S. (1992). A sequential learning analysis of decisions in organizations to escalate investments despite continuing costs or losses. *Journal of Applied Behavior Analysis,* **25,** 561–574.

Hantula, D. A. (1990, August). *The role of an alternative course of action in sustaining and reversing escalation and persistence.* Paper presented at the Academy of Management, San Francisco, CA.

Hauer, E. (1983). An application of the likelihood/Bayes approach to the estimation of safety countermeasure effectiveness. *Accident Analysis and Prevention,* **15,** 287–298.

Heckman, J. J., & Willis, R. (1977). A beta-logistic model for the analysis of sequential labor force participation by married women. *Journal of Political Economy,* **85,** 27–58.

Hemsley, D. R., & Garety, P. A. (1986). The formation of delusions: A Bayesian analysis. *British Journal of Psychiatry,* **12,** 1–17.

Hoffman, D. A., Jacobs, R., & Gerras, S. J. (1992). Mapping individual performance over time. *Journal of Applied Psychology,* **77**(2), 185–195.

Hsu, L. M. (1988). Fixed versus flexible MMPI diagnostic rules. *Journal of Consulting and Clinical Psychology,* **56,** 458–462.

Hunte, W., Myers, R. A., & Doyle, R. W. (1985). Bayesian mating decisions in an amphipod, *Gammarus lawrencianus* Bousfield. *Animal Behavior,* **33,** 366–372.

Hursh, S. R. (1984). Behavioral economics. *Journal of the Experimental Analysis of Behavior,* **42,** 435–452.

Kagel, J. H., & Winkler, R. C. (1972). Behavioral economics: Areas of cooperative research between economics and applied behavioral analysis. *Journal of Applied Behavior Analysis,* **3,** 335–342.

Kahneman, D., & Tversky, A. (Eds.). (1982). *Judgment under uncertainty: Heuristics and biases.* New York: Cambridge University Press.

Kahneman, D., & Tversky, A. (1983). Can irrationality be intelligently discussed? *Behavioral and Brain Sciences,* **6,** 509–510.

Katz, L. (1986). *Layoffs, recalls, and the duration of unemployment* (National Bureau of Economic Research Working paper, number 1825). Cambridge, MA: National Bureau of Economic Research.

Keynes, J. M. (1921). *A treatise on probability.* New York: Harper and Row.

Killeen, P. R. (1982). Incentive theory. In D. J. Bernstein (Ed.), *Nebraska symposium on motivation, 1981: Response structure and organization* (pp. 169–216). Lincoln: University of Nebraska Press.

Kolmogorov, A. N. (1950). *Foundations of probability theory.* New York: Chelsea Publishing Co. [Original article: Grundbegriffe der Wahrscheinlichkeitsrechneung, *Erglle Mathematica,* 1933, **2**(3)]

Komaki, J. L., & Citera, M. (1990). Beyond effective supervision: Identifying key interactions between superior and subordinate. *Leadership Quarterly,* **1,** 91–105.

Komaki, J. L., Coombs, T., & Schepman, S. (1991). Motivational implications of reinforcement theory. In R. M. Steers & L. W. Porter (Eds.), *Motivation and work behavior* (pp. 87–107). New York: McGraw-Hill.

Levi, I. (1991). *The fixation of belief and its undoing.* New York: Cambridge University Press.

Lewis, K. K. (1988). The persistence of the peso problem when policy is noisy. *Journal of International Money and Finance,* **7,** 1–11.

Lewis, K. K. (1989). Changing beliefs and systematic rational forecast errors with evidence from foreign exchange. *American Economic Review,* **79,** 621–636.

Lim, R. G., & Drasgow, F. (1990). Evaluation of two methods for estimating item response theory parameters when assessing differential item functioning. *Journal of Applied Psychology,* **75,** 164–174.

Lucas, R. (1976). Econometric policy evaluation: A critique. In K. Brunner & A. Meltzer (Eds.), *The Phillips curve and the market* (pp. 19–46). Amsterdam: North Holland.

Malkiel, B. (1985). *A random walk down Wall Street* (4th ed.). New York: Norton.

McCain, B. (1986). Continuing investment under conditions of failure: A laboratory study of the limits to escalation. *Journal of Applied Psychology,* **71,** 280–284.

McGrath, J. E., & Kelly, J. R. (1986). *Time and human interaction: Toward a social psychology of time.* New York: Guilford.

Mendel, M. B., & Sheridan, T. B. (1989). Filtering information from human experts. *IEEE Transactions on Systems, Man and Cybernetics,* **19**, 6–16.

Muth, J. F. (1961). Rational expectations and the theory of price movements. *Econometrica,* **29**, 315–335.

Myerson, R. (1985). Bayesian equilibrium and incentive compatibility: An introduction. In L. Hurwicz, D. Schmeidler, & H. Sonnenshein (Eds.), *Social goals and social organization* (pp. 229–260). New York: Cambridge University Press.

O'Flaherty, B., & Siow, A. (1990, July). *Up or out rules in the market for lawyers.* Paper presented at the National Bureau of Economic Research Conference on Empirical Methods in Industrial Organization, Cambridge, MA.

Osuna, E. E. (1985). The psychological cost of waiting. *Journal of Mathematical Psychology,* **29**, 82–105.

Samuelson, P. (1965). Proof that properly anticipated prices fluctuate randomly. *Industrial Management Review,* **6**(2), 41–49.

Sargent, T. J., & Wallace, N. (1976). Rational expectations and the theory of economic policy. *Journal of Monetary Economics,* **2**(2), 169–183.

Skinner, B. F. (1984). *A matter of consequences.* New York: New York University Press.

Staw, B. M., & Ross, J. (1987). Behavior in escalation situations: Antecedents, prototypes and solutions. *Research in Organizational Behavior,* **9**, 39–78.

Staw, B. M., & Ross, J. (1989). Understanding behavior in escalation situations. *Science,* **246**, 216–220.

Wallsten, T. S. (1976). Using conjoint-measurement models to investigate a theory about probabilistic information processing. *Journal of Mathematical Psychology,* **14**, 144–185.

Received May 25, 1992
Initial editorial decision June 7, 1992
Revision received June 15, 1992
Final acceptance June 18, 1992
Action Editor, E. Scott Geller