Dynamic Cost–Loss Ratio Decision-making Model with an Autocorrelated Climate Variable

RICHARD W. KATZ

Environmental and Societal Impacts Group, National Center for Atmospheric Research, Boulder, Colorado*

(Manuscript received 9 July 1991, in final form 15 April 1992)

ABSTRACT

A dynamic decision-making problem is considered involving the use of information about the autocorrelation of a climate variable. Specifically, an infinite horizon, discounted version of the dynamic cost–loss ratio model is treated, in which only two states of weather ("adverse" or "not adverse") are possible and only two actions are permitted ("protect" or "do not protect"). To account for the temporal dependence of the sequence of states of the occurrence (or nonoccurrence) of adverse weather, a Markov chain model is employed. It is shown that knowledge of this autocorrelation has potential economic value to a decision maker, even without any genuine forecasts being available. Numerical examples are presented to demonstrate that a decision maker who erroneously follows a suboptimal strategy based on the belief that the climate variable is temporally independent could incur unnecessary expense. This approach also provides a natural framework for extension to the situation in which forecasts are available for an autocorrelated climate variable.

1. Introduction

Most studies of the economic value of weather and climate forecasts have tacitly assumed that the variable being forecast does not possess any temporal dependence. Nevertheless, it is well known that many climate variables are significantly autocorrelated (typically, a positive correlation termed "persistence") over a wide range of time scales. For instance, the time series of daily mean temperature, daily occurrence (or nonoccurrence) of precipitation, and monthly or seasonal Southern Oscillation index all possess substantial autocorrelation. In the absence of any forecasting system, it stands to reason that information about the degree of persistence of a climate variable may be of potential economic value to a decision maker. Moreover, it might be reasonable to employ persistence forecasts, instead of assuming an independent climate variable, as a standard of comparison for any forecasts that might be available to a decision maker.

If the decision-making situation being modeled were static (i.e., the action taken on the present occasion has no bearing on any actions taken on future occasions), then the issue of autocorrelation could be treated in a relatively straightforward manner. Knowing that a climate variable is autocorrelated could be viewed as simply a special case of having forecasts of higher quality than those based on the assumption of independence. Most decision-making problems, however, are dynamic in nature (e.g., Katz et al. 1982; Brown et al. 1986; Mjelde et al. 1988). Because the action taken by a decision maker on the current occasion does have an effect on any actions to be taken on future occasions, the presence of autocorrelation creates a decision-making problem that is inherently more complex.

A few attempts to allow for autocorrelation when assessing the economic value of information about weather or climate have been made. For instance, Epstein and Murphy (1988) considered the value of multiple-period forecasts with an autocorrelated climate variable, dealing with a very specialized dynamic decision-making problem having only a three-period time horizon. Wilks (1991) treated a broader type of dynamic decision-making problem with a finite time horizon and a persistent climate variable. The complexity of the problem forced him to resort to the computational approach of stochastic dynamic programming, only generating output for particular numerical examples. His results indicate that information about autocorrelation is itself of potential economic value to a decision maker, as well as that its presence has an effect on assessments of the value of any forecasts.

In the present paper, the focus is on the economic value of information about the temporal dependence of a climate variable, not on any genuine forecasts that might be available to the decision maker. The objective is to develop an analytical framework on which the

* The National Center for Atmospheric Research is sponsored by the National Science Foundation.

Corresponding author address: Dr. Richard W. Katz, Environmental and Societal Impacts Group, National Center for Atmospheric Research, P.O. Box 3000, Boulder, CO 80307.

extension to forecasts (whether for just the next occasion, or even for all occasions indefinitely into the future) could naturally be based. More specifically, the dynamic cost–loss ratio decision-making model (Katz and Murphy 1990) is applied to the situation in which the temporal dependence of the climate variable is modeled by a Markov chain. For the infinite horizon, discounted version of this particular problem, it is still possible to obtain analytical expressions for the optimal actions and minimal expected expenses associated with an autocorrelated climate. Of primary interest is how these results differ from the suboptimal strategy, based on the erroneous belief that the climate variable is temporally independent.

In section 2, the notation and basic concepts are introduced for the dynamic cost–loss ratio decision-making problem. Analytical results concerning the structure of the optimal policy and the corresponding minimal expected expenses for this model are presented in section 3, with technical derivations being relegated to the Appendix. Section 4 contains some numerical examples, illustrating these analytical results. The extension of this approach to the more general problem of assessing the economic value of actual forecasts is explored in section 5.

2. Notation and concepts

a. *Decision-making model*

In the cost–loss ratio model, a decision maker must choose on each occasion between two possible actions: either (i) protect; or (ii) do not protect. The climate variable has only two possible states: either (i) adverse weather (with the indicator variable $\theta = 1$); or (ii) not adverse weather (with $\theta = 0$). If protective action is taken, then the decision maker incurs a cost C. On the other hand, if protective action is not taken and adverse weather does occur, a loss L is incurred, $0 < C < L$. The static version of this decision-making problem, in which the decision maker simply minimizes expected expense for a single occasion at a time, was treated by Katz and Murphy (1987).

The model is made dynamic by assuming that the loss L can be incurred at most once. Now the goal of the decision maker is to minimize the expected expense, not for a single occasion, but totaled over an infinite horizon (i.e., indefinitely into the future). Future expenses are discounted, the rationale being that a dollar today is worth more than a dollar in the future because of uncertainty and opportunities for investment and consumption. Discounting accounts for an individual's time preference for money.

Specifically, an expense V incurred on the next occasion has a "present value" of only αV, where α, $0 < \alpha < 1$, is a "discount factor." In relative terms, this "future value" V is diminished (or discounted) at the rate $r = (V - \alpha V)/\alpha V = (1 - \alpha)/\alpha$, termed the "discount rate" (for a more detailed explanation, see Alchian and Allen 1972, chapter 11). The finite horizon, undiscounted version of the dynamic cost–loss ratio decision-making problem was first introduced by Murphy et al. (1985), and additional results were obtained by Krzysztofowicz and Long (1990).

b. *Climatological information*

The expected expenses are calculated with respect to some probability distribution for the climate variable θ. Instead of taking the sequence of climate states $\{\theta_t: t = 0, 1, \cdots\}$ to be independent, it is now assumed to constitute a two-state, first-order Markov chain. This model has been applied, for instance, to the daily occurrence of precipitation (Gabriel and Neumann 1962). The parameters of the Markov chain include the unconditional (or marginal) probability of adverse weather

$$p_\theta = \Pr\{\theta_t = 1\}, \qquad (1)$$

$0 < p_\theta < 1$, and transition probabilities relating the states at times $t - 1$ and t

$$P_{ij} = \Pr\{\theta_t = j \,|\, \theta_{t-1} = i\}, \qquad (2)$$

$i, j = 0, 1$.

Two parameters are sufficient to characterize the Markov chain, and it is convenient to work with the climatological probability of adverse weather p_θ and the first-order autocorrelation coefficient

$$d = \operatorname{corr}(\theta_{t-1}, \theta_t) = P_{11} - P_{01}. \qquad (3)$$

The transition probabilities can be expressed as

$$P_{01} = p_\theta(1 - d), \quad P_{00} = 1 - P_{01},$$
$$P_{10} = (1 - p_\theta)(1 - d), \quad P_{11} = 1 - P_{10}. \qquad (4)$$

It is assumed that the autocorrelation parameter is positive (i.e., $0 < d < 1$), as is typical for climate variables. Now the autocorrelation function for the $\{\theta_t\}$ process decreases at a geometric rate toward zero; that is,

$$\operatorname{corr}(\theta_{t-k}, \theta_t) = d^k, \qquad (5)$$

$k = 1, 2, \cdots$, where "corr" denotes the correlation coefficient. Consequently, this form of climatological information actually provides "forecasts" over the entire future horizon, whose skill cascades down to zero. The limiting case of $d = 0$ corresponds to an independent climate variable, as was treated for the infinite horizon, discounted version of the dynamic cost–loss ratio model by Katz and Murphy (1990). The other extreme of $d = 1$ corresponds to a perfectly persistent climate.

c. *Optimal policy and expected expense*

A rule that specifies the action to be taken as a function of the information available to the decision maker (e.g., the most recent state of the weather) is termed a "policy." Among all possible policies, the "optimal policy" is the one that minimizes the expected expense, denoted by E_c, discounted and totaled over the infinite future horizon. The assumption that all future climate states are correlated with the present state [see (5)], in combination with the dynamics of the cost–loss ratio decision-making model, makes the problem of determining these expected expenses relatively complex.

Because of the assumed Markovian property of the climate variable, it simplifies matters to condition on the initial climate state. The expense E_c can be expressed as

$$E_c = (1 - p_\theta)E_c(0) + p_\theta E_c(1), \tag{6}$$

where $E_c(i)$ denotes the minimal expected expense given $\theta_0 = i$, $i = 0, 1$. In other words, E_c represents the unconditional or weighted average expense. Although the primary focus is on expense E_c, it turns out to be easier to first derive expressions for the two conditional expenses, $E_c(0)$ and $E_c(1)$, and then obtain an expression for E_c by substitution into (6).

It is of interest to compare the minimal expected expense for the optimal policy with the expense corresponding to erroneously following the suboptimal policy based on temporal independence. This difference in expected expenses is one reasonable way to quantify the economic value to the decision maker of taking into account the autocorrelation of the climate variable.

3. Analytical results

a. Optimal policy

Attention can be restricted to so-called "stationary policies." These are strategies that depend only on the information currently available and not in any other way on which particular occasion the decision is being made (Ross 1970, chapter 6). The past history of weather states (i.e., $\theta_0, \theta_1, \ldots, \theta_{t-1}$) is available to the decision maker prior to the tth occasion, of which only the most recent state θ_{t-1} needs to be retained because of the imposition of the Markovian property on the climate variable. Consequently, four different forms of strategy (i.e., two states of weather in combination with two possible actions) must be considered:

Strategy (i). Always protect (i.e., protect on tth occasion whether $\theta_{t-1} = 0$ or $\theta_{t-1} = 1$);

Strategy (ii). Never protect (i.e., do not protect on tth occasion whether $\theta_{t-1} = 0$ or $\theta_{t-1} = 1$);

Strategy (iii). Protect on tth occasion if $\theta_{t-1} = 1$, do not protect if $\theta_{t-1} = 0$; and

Strategy (iv). Protect on tth occasion if $\theta_{t-1} = 0$, do not protect if $\theta_{t-1} = 1$.

Because the persistence parameter d is assumed positive, strategy (iv) can never be optimal and is eliminated from consideration.

The structure of the optimal policy for dependent climatology [i.e., conditions under which strategies (i)–(iii) are optimal] is derived in the Appendix. First, the results for the two limiting circumstances of $d = 0$ and $d = 1$ are presented.

Independent climate ($d = 0$). Because using the past history of θ makes no difference economically when $d = 0$, strategy (iii) cannot be optimal. As shown by Katz and Murphy (1990) and as follows from a special case of the result (A6) derived in the Appendix, strategy (i) is optimal if $p_\theta > p_\theta^*$, where

$$p_\theta^* = \frac{C/L}{1 - \alpha^*(C/L)} \tag{7}$$

and the reciprocal of the discount rate r is

$$\alpha^* = \frac{\alpha}{1 - \alpha}. \tag{8}$$

In other words, strategy (i) of always protecting is optimal provided the climatological probability of adverse weather p_θ is large enough relative to the cost–loss ratio C/L, adjusted for the discount factor α. Otherwise, strategy (ii) of never protecting is optimal.

Perfectly persistent climate ($d = 1$). Because it is not economical to protect when it is known that adverse weather will not occur (i.e., if adverse weather did not occur previously), strategy (i) of always protecting cannot be optimal. As is easy to show directly and as also follows from the result (A8) derived in the Appendix, strategy (iii) of only protecting when given that adverse weather has just occurred is optimal provided

$$C/L < 1 - \alpha, \tag{9}$$

no matter what the probability of adverse weather p_θ. Otherwise, strategy (ii) of never protecting is optimal. Condition (9) can be interpreted as requiring that future expenses be discounted at a rate r that is large enough, relative to the cost–loss ratio C/L.

General case ($0 < d < 1$). It is shown in the Appendix that the strategy (i) of always protecting is optimal for $0 < d < d^*$, where the critical persistence threshold is

$$d^* = 1 - \frac{C/L}{[1 - \alpha^*(C/L)]p_\theta}. \tag{10}$$

This threshold is well defined (that is, d^* is contained within the unit interval [0, 1]), provided $p_\theta > p_\theta^*$ [see (7)]. It represents the break-even point between strategy (i) or (iii) being optimal, and is directly related to the relative magnitudes of the cost–loss ratio C/L and the probability of adverse weather p_θ, except for an additional factor that reflects the discounting of future expenses. The interval (0, d^*) for which always protecting is the optimal strategy becomes wider (i.e., d^* increases) as C/L decreases, as p_θ increases, or as the discount rate r increases.

Obtaining the range of values of d for which strategy (ii) of never protecting is optimal [and, by the process of elimination, the range for which strategy (iii) is optimal] is a much more difficult problem. It is shown in the Appendix that the critical values of d, specifying when strategy (ii) becomes optimal, satisfy a quadratic equation in d,

$$(1 - p_\theta)\alpha d^2 - \{1 - p_\theta[1 + \alpha - \alpha^*(C/L)]\}d + \{(C/L) - p_\theta[1 - \alpha^*(C/L)]\} = 0, \quad (11)$$

with roots d^{**} and d^{***}. These roots specify the transition between strategy (ii) or (iii) being optimal. Depending on the numerical values of C/L, α, and p_θ, neither, one, or both of the roots may fall within the unit interval [0, 1] of relevance for climate variables. Rather than attempting to provide any general interpretation of the structure of the optimal policy (11), some numerical examples are presented in section 4.

b. Expected expense

Using the idea of "backwards induction" (Ross 1983), the conditional expected expenses satisfy a simple stochastic dynamic programming recursion

$$E_c(i) = \min(\{C + \alpha[P_{i0}E_c(0) + P_{i1}E_c(1)]\}, \{P_{i1}L + \alpha P_{i0}E_c(0)\}), \quad i = 0, 1. \quad (12)$$

The first (second) term (in curly brackets) on the right-hand side of (12) is the expected expense over the future horizon if protective action is (is not) taken on the current occasion. The Appendix contains a simple probabilistic argument by which this recursion may be obtained. All of the expressions for the expected expenses to be utilized, as well as for the structure of the optimal policy already stated in section 3a, follow directly from (12). This fundamental recursion specifies a system of two equations in two unknowns, $E_c(0)$ and $E_c(1)$. Once these two solutions are obtained, the desired expected expense E_c can be determined by (6).

The minimal expected expenses for strategies (i)–(iii) (derived in the Appendix) are of the following form:

(i) Strategy (i) (always protect)

$$E_c = \frac{C}{1 - \alpha}. \quad (13)$$

(ii) Strategy (ii) (never protect)

$$E_c = \frac{(1 - \alpha d)p_\theta L}{1 - \alpha[1 - p_\theta(1 - d)]}. \quad (14)$$

(iii) Strategy (iii) (protect on tth occasion if $\theta_{t-1} = 1$)

$$E_c = \left[\frac{\{1 - \alpha[1 - p_\theta(1 - d)]\}C + (1 - p_\theta)(1 - \alpha d)(1 - d)L}{\{1 - \alpha[1 - p_\theta(1 - d)]\}\{1 - \alpha[1 - (1 - p_\theta)(1 - d)]\}}\right]p_\theta. \quad (15)$$

Conditions that specify which of these three expected expenses is minimal were presented in section 3a. Nevertheless, (13)–(15) hold even when the strategy adopted by the decision maker is not optimal, an issue that arises when defining the economic value of any information about the persistence of climate (see section 4). To help interpret these expressions for the minimal expected expense, again the two limiting cases of $d = 0$ and $d = 1$ are examined.

Independent climate ($d = 0$). For the strategy (i) of always protecting, the expected expense (13) remains unchanged because it is independent of d. This expression simply represents the present value of incurring a cost C on the current and on every future occasion. The expected expense (14) for strategy (ii) of never protecting reduces to

$$E_c = \frac{p_\theta L}{1 - \alpha(1 - p_\theta)} \quad (16)$$

(in agreement with Katz and Murphy 1990). In (16), the immediate expected expense for never protecting of $p_\theta L$ is adjusted by a factor arising to allow for the possibility that the loss L is incurred on some future occasion and for the discounting of future expenses. Finally, the expected expense (15) for strategy (iii) of protecting on the tth occasion if $\theta_{t-1} = 1$ reduces to

$$E_c = \left\{\frac{[1 - \alpha(1 - p_\theta)]C + (1 - p_\theta)L}{[1 - \alpha(1 - p_\theta)](1 - \alpha p_\theta)}\right\}p_\theta \quad (17)$$

[but recall that strategy (iii) cannot be optimal if $d = 0$]. Not surprisingly, (17) is very nearly just a weighted average of (13) and (16).

Perfectly persistent climate ($d = 1$). As already observed, the expected expense (13) for strategy (i) of always protecting remains unchanged [but recall that strategy (i) cannot be optimal if $d = 1$]. The expected expense (14) for strategy (ii) of never protecting simplifies to $E_c = p_\theta L$, or just the same as the immediate expected expense of never protecting. Finally, the expected expense (15) for strategy (iii) of protecting on the tth occasion if $\theta_{t-1} = 1$ reduces to $E_c = p_\theta C/(1 - \alpha)$, or just the present value of always incurring expense C weighted by p_θ.

General case ($0 < d < 1$). Like (16), the expected expense (14) for strategy (ii) of never protecting has the same interpretation, except that now the adjustment factor involves the persistence parameter d as well (because this parameter affects the future probability of incurring the loss L). Analogous to (17), the expected expense (15) for strategy (iii) resembles a weighted average of (13) and (14). To illustrate how these expected expenses vary as a function of d, some numerical examples are presented in section 4.

4. Numerical examples

Before presenting some examples, the concept of economic value needs to be more fully explained (e.g., Winkler and Murphy 1985). First, suppose that the decision maker naïvely follows the suboptimal policy based on the erroneous belief that the climate variable is temporally independent. Nevertheless, the expected expenses incurred by the decision maker should be determined from the expressions (13)–(15) that correctly take into account dependence. The rationale is that, even though the decision maker is not necessarily aware of the existence of autocorrelation, the actual expenditures will still reflect the persistence of climate.

Formally, the economic value of information about the temporal dependence of a climate variable is defined as

> The reduction in expected expense when, instead of the suboptimal policy based on the erroneous assumption of an independent climate, the optimal policy is adopted by the decision maker.

It is an immediate consequence of this definition that information about the temporal dependence of climate will have positive economic value to the decision maker only if the suboptimal policy based on independence actually differs from the optimal policy under dependence.

In each of the three following examples, the numerical values of the two economic parameters, the cost–loss ratio C/L and the discount factor α, and of the climate parameter, the unconditional probability of adverse weather p_θ, are specified. Attention is focused on the other climate parameter, the first-order autocorrelation coefficient d, which is allowed to vary all the way from independence (i.e., $d = 0$) to perfect persistence (i.e., $d = 1$). Without loss in generality, it is assumed that the loss $L = 1$ (i.e., $C/L = C$). Hence, the minimal expected expenses can be viewed as scaled to fall within the unit interval [0, 1].

a. Example 1 ($C/L = 0.01$, $\alpha = 0.98$, $p_\theta = 0.025$)

This example concerns a situation in which the decision maker incurs a relatively large penalty for naïvely ignoring autocorrelation, provided the dependence parameter d is sufficiently greater than zero. The numerical value of the discount factor α happens to be a reasonable choice for a seasonal or longer time scale (i.e., a discount rate of $r \approx 0.02$).

Optimal policy. The analytical results presented in section 3a can be employed to establish the exact form of optimal policy. Because $C/L = 0.01 < 0.02 = 1 - \alpha$, (9) implies that strategy (iii) of protecting given only that adverse weather has just occurred is optimal for $d = 1$. Moreover, because $p_\theta = 0.025 > 0.0196 \approx p_\theta^*$ [by (7)], strategy (i) of always protecting must be optimal for $d = 0$. The transition between strategy (i) or (iii) being optimal occurs at the persistence

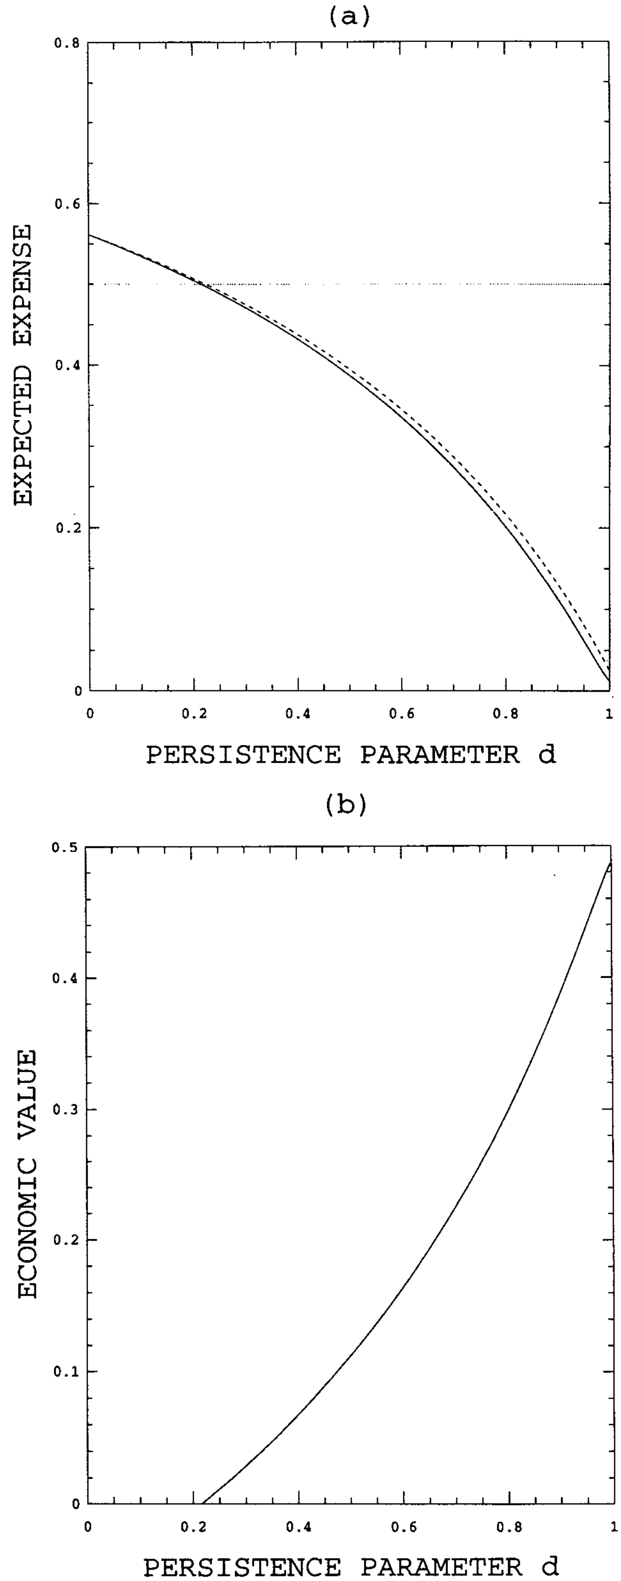

FIG. 1. Example 1 ($C/L = 0.01$, $\alpha = 0.98$, $p_\theta = 0.025$): (a) Expected expense for strategy (i) (dotted line), strategy (ii) (dashed line), and strategy (iii) (solid line) versus persistence parameter d; and (b) economic value versus d.

threshold of $d^* \approx 0.2157$ [by (10)]. It turns out that both of the roots of the quadratic equation (11) are outside of the interval [0, 1], implying that strategy (ii) of never protecting cannot be optimal. Unaware of the presence of autocorrelation, the decision maker would naïvely follow strategy (i) because it is optimal for $d = 0$. But this strategy would be suboptimal for $0.2157 < d < 1$.

Expected expense and economic value. The expected expense curves for strategies (i)–(iii) [based on (13)–(15)] are shown in Fig. 1a. As d increases above the threshold value of 0.2157, the difference between the expected expense for strategy (i) (i.e., the suboptimal strategy that would be naïvely followed by the decision maker) and for strategy (iii) (i.e., the optimal strategy) steadily increases. Figure 1b gives this difference in expected expense, termed the economic value of persistence. The curve reaches a maximum of about 0.49 (or nearly one-half of the normalized loss $L = 1$) at $d = 1$. It is convex over most of the range of d, becoming concave only for values of d near one (i.e., above $d \approx 0.97$). This result is in slight contrast to those obtained by Katz and Murphy (1990) for the same dynamic cost–loss ratio decision-making model. With an independent climate (i.e., $d = 0$) and imperfect forecasts, they showed that the quality/value curve must be convex. In other words, an increasing degree of persistence is not completely analogous to an improvement in the quality of imperfect forecasts, insofar as economic value is concerned.

b. Example 2 ($C/L = 0.02$, $\alpha = 0.98$, $p_\theta = 0.05$)

This example concerns a situation in which the decision maker, who naïvely does not follow the optimal policy, still incurs virtually no penalty in terms of increased expected expense. The same numerical value of the discount factor α is employed as in example 1.

Optimal policy. Since $C/L = 0.02 = 1 - \alpha$, (9) implies that the decision maker should be indifferent between strategy (ii) and (iii) in the extreme case of $d = 1$. Moreover, because $p_\theta = 0.05 < 1 = p_\theta^*$ [by (7)], strategy (ii) must be optimal when $d = 0$. The threshold d^* [see (10)] falls outside the interval [0, 1], implying that strategy (i) is never optimal. On the other hand, one of the roots of the quadratic equation (11) is $d^{**} \approx 0.0204$ and the other is $d^{***} = 1$ [note that it is easy to show that $d = 1$ must be a root of (11) when $C/L = 1 - \alpha$]. So strategy (ii) is optimal for $0 < d < 0.0204$, whereas strategy (iii) is optimal for $0.0204 < d < 1$. The decision maker would naïvely select strategy (ii), because it is optimal for $d = 0$. But this strategy would be suboptimal for nearly all values of d (i.e., $0.0204 < d < 1$).

Expected expense and economic value. The expected expense curves for strategies (ii) and (iii) are shown in Fig. 2a [curve for strategy (i) not shown because it

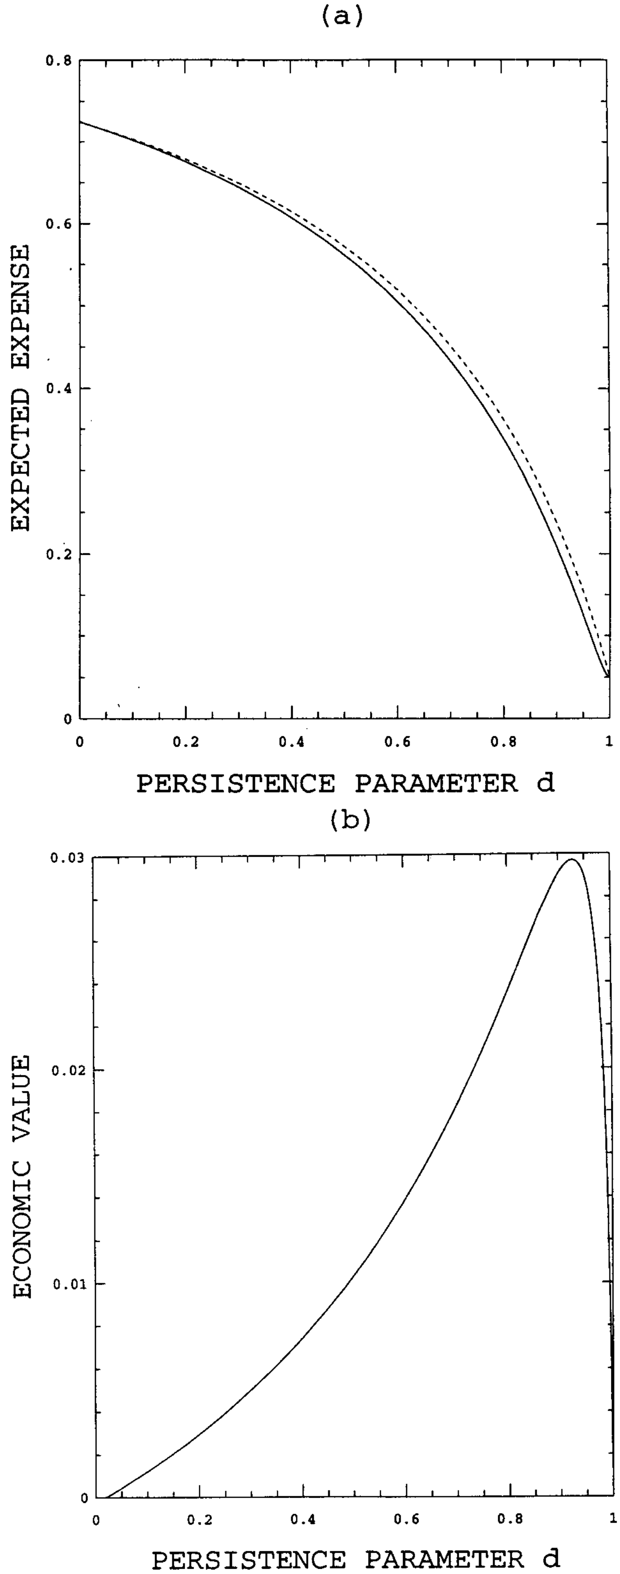

FIG. 2. Example 2 ($C/L = 0.02$, $\alpha = 0.98$, $p_\theta = 0.05$): (a) Expected expense for strategy (ii) (dashed line) and strategy (iii) (solid line) versus persistence parameter d; and (b) economic value versus d.

cannot be optimal]. Very little difference in expected expense exists between the strategy (ii) and (iii) curves, irrespective of which strategy is optimal. Surprisingly, what small differences do occur gradually increase up to $d \approx 0.93$, but then decrease back down to zero at $d = 1$. Although both expense curves monotonically decrease as the persistence parameter d increases, the difference in their relative rates of decrease suffices to explain this result. The fact that economic value (Fig. 2b) actually decreases for large d might not have been anticipated. For a fixed form of decision-making model, the economic value cannot decrease as the quality of imperfect forecasts increases [provided the measure of quality is consistent with the so-called "sufficiency" relation (Ehrendorfer and Murphy 1988)] (Hilton 1981). Again, this result confirms the danger of drawing an analogy between improved forecasts and increased persistence.

c. *Example 3 ($C/L = 0.15$, $\alpha = 0.9$, $p_\theta = 0.2$)*

This example concerns a situation in which some unanticipated results are obtained. In particular, the expected expense for one strategy actually increases as the persistence parameter d approaches one. The numerical value of the discount factor α is a reasonable choice for an annual or longer time scale (i.e., a discount rate of $r \approx 0.11$).

Optimal policy. Because $C/L = 0.15 > 0.1 = 1 - \alpha$, (9) implies that strategy (ii) is optimal when $d = 1$. Moreover, because the persistence threshold d^* [see (10)] falls outside the interval [0, 1], strategy (i) can never be optimal. By the process of elimination, strategy (ii) must also be optimal when $d = 0$. On the other hand, because $C/L > 1 - \alpha$, the quadratic equation (11) has both roots within the interval [0, 1]. For the parameter values specified in this example, $d^{**} \approx 0.3416$ and $d^{***} \approx 0.8945$. In other words, strategy (ii) is optimal for both small values of d, namely, $0 < d < 0.3416$, and large values of d, namely, $0.8945 < d < 1$, whereas strategy (iii) is optimal for $0.3416 < d < 0.8945$. The decision maker would naïvely follow strategy (ii), because it is optimal for $d = 0$. But this strategy would be suboptimal for $0.3416 < d < 0.8945$.

Expected expense and economic value. The expected expense curves for strategies (ii) and (iii) are shown in Fig. 3a [curve for strategy (i) not shown because it cannot be optimal]. While the expected expense curve for strategy (ii) steadily decreases as the persistence parameter d increases, that for strategy (iii) decreases to a minimum at $d \approx 0.98$, then actually increases as d approaches one. Consequently, the economic value of persistence (Fig. 3b) first increases, reaching a maximum at about $d \approx 0.73$, and then decreases back down to zero. Although it might not have been anticipated that the expected expense for a given strategy could increase as d increases, changing the degree of

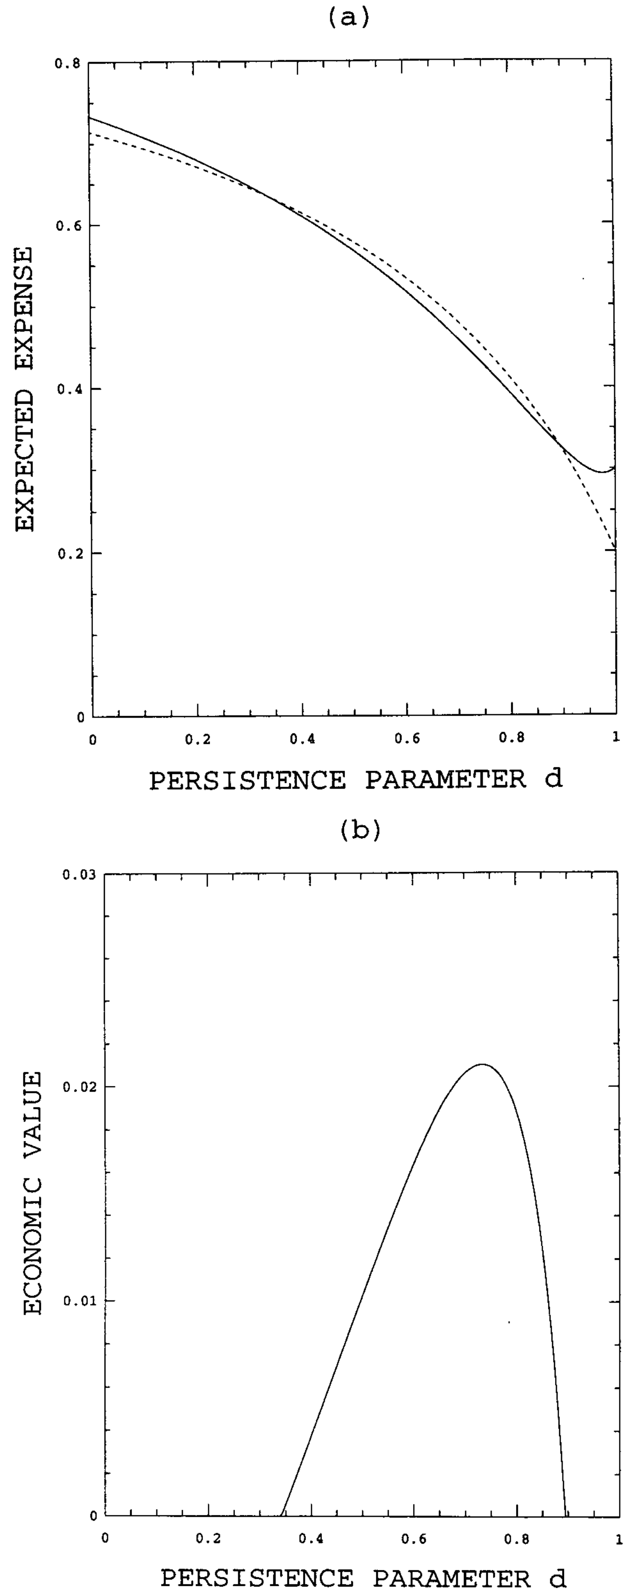

FIG. 3. Example 3 ($C/L = 0.15$, $\alpha = 0.9$, $p_\theta = 0.2$): Same format as Fig. 2.

persistence of a climate variable is tantamount to changing the climate. This factor complicates the interpretation of any results that involve comparisons as the parameter d varies. It is also consistent with the nonexistence of any general results concerning how the value of information changes as the other characteristics of the decision-making situation change (Hilton 1981).

5. Implications and extensions

For an infinite horizon, discounted version of the dynamic cost–loss ratio decision-making model, the situation in which the sequence of occurrences (or nonoccurrences) of adverse weather is modeled as a two-state Markov chain has been treated. It has been established for this model under precisely what conditions a decision maker who naïvely ignores the autocorrelation of the climate variable would be following a suboptimal policy. Moreover, through specific numerical examples, it has been demonstrated to what extent expense could be avoided by a decision maker who takes into account persistence. Drawing an analogy between a greater degree of persistence and better forecasts, some of these properties could have been anticipated, but others turn out to be somewhat surprising.

The notion that a relatively simple form of statistical information, characterizing the tendency of climate to persist, can still be of economic value to a decision maker could be further explored. Such information ordinarily would not be viewed as sophisticated enough to constitute a bona fide forecasting system. Although real-world decision-making situations are more complex than the one treated here, perhaps some decision makers have evolved strategies that take into account such persistence, especially if actual forecasts were either unavailable or ignored because of their perceived inaccuracy. For instance, the fact that the structure of the optimal policy depends on climatological persistence could be employed to speculate about how individual decision makers might adapt their strategies to a changing climate. A justification of the need for information—not just about how the probability of adverse weather would change, but about how the autocorrelation of the climate variable would change as well—has been provided. Nevertheless, this issue has received very little attention so far in studies of global climate change.

Originally, it was the intention to quantify the economic value of weather or climate forecasts for multiple future periods. The complexity of this problem made a compelling argument for first considering, as the appropriate standard of comparison, the situation in which information about the autocorrelation of the climate variable is available to the decision maker. It has been seen that the simplest form of first-order dependence (namely, a two-state Markov chain) produces statistical forecasts for all future time horizons, with the skill cascading toward zero at a geometric rate. If genuine forecasts were available, then the same sort of issues would arise in interpreting lead time and limits to predictability. For instance, actual forecasts for only one occasion ahead, in combination with an autocorrelated climate variable, would produce "forecasts" for all future time horizons whose skill would be greater than that for the autocorrelation alone.

To treat these issues in a systematic fashion, it would be necessary to model the joint time series of forecasts and observed climate as a bivariate stochastic process. It might be reasonable to assume that the sequence of forecasts of adverse weather (or not adverse weather) is likewise a two-state Markov chain. Other assumptions can be devised by thinking in terms of the two limiting circumstances of either (i) a forecasting system with no skill over that for the autocorrelation of the climate variable; or (ii) perfect forecasts, not just for the next occasion, but for every future occasion. It would be plausible to require that the probability of a forecast of adverse weather be identical to the climatological probability of occurrence of adverse weather, and that the first-order autocorrelation coefficient of the time series of forecasts be identical to that for the climate variable. This approach would be analogous to that taken in research on teleconnections, in which to identify the nature of the relationship between variables, one is driven to define the concepts of "leading," "lagging," and "feedback" by means of measures of predictability (Katz 1988).

Acknowledgments. I thank Barbara G. Brown and Mary W. Downton for help in drafting figures, Kathleen A. Miller for supplying references, and Allan H. Murphy and Daniel S. Wilks for providing comments. I am grateful to one of the referees for detecting an error in one equation. This research was supported in part by the National Science Foundation (Division of Atmospheric Sciences) under Grant ATM-8714108. A portion of this research was conducted while the author was a sabbatical visitor in the Department of Land, Air and Water Resources at the University of California, Davis.

APPENDIX

Stochastic Dynamic Programming

1. Dynamic programming recursion

The derivation of (12) follows from a purely probabilistic argument. Given an initial weather state of $\theta_0 = i$, the decision maker is faced with the choice on the first occasion of either (i) protecting or (ii) not protecting.

(*i*) *Protect*. If the decision maker takes protective action on the first occasion, then an immediate expense C is incurred. Moreover, on the second occasion the decision maker will be given weather state $\theta_1 = j$ with probability P_{ij}, $j = 0, 1$. In these circumstances, the decision-making process can be viewed as probabilistically equivalent to starting over again with an initial state of $\theta_0 = j$. By definition, the future expected expense, totaled and discounted over the still infinite horizon from the second occasion onwards, is $E_c(j)$. Consequently, the expected expense over this future horizon is

$$P_{i0}E_c(0) + P_{i1}E_c(1). \quad \text{(A1)}$$

Because future expenses are being discounted, this expression (A1) must be multiplied by the discount factor α. Combining these results yields the first expression (in curly brackets) on the right-hand side of (12).

(*ii*) *Do not protect*. If the decision maker does not take protective action on the first occasion, then with probability P_{i1} an immediate expense L is incurred and the decision-making process terminates (i.e., no future expenses are incurred). On the other hand, with probability P_{i0} no expense is incurred on the first occasion. In these circumstances, the decision-making process can be viewed as starting over again on the second occasion with an initial state of $\theta_0 = 0$. Consequently, the discounted expected expense over the infinite future horizon starting with the second occasion is $\alpha E_c(0)$. Combining these results yields the second expression (in curly brackets) on the right-hand side of (12).

2. Expected expense

Strategy (i). If this strategy is selected by the decision maker, then the dynamic programming recursion (12) reduces to

$$E_c(i) = C + \alpha[P_{i0}E_c(0) + P_{i1}E_c(1)], \quad \text{(A2)}$$

$i = 0, 1$. Multiplying both sides of (A2) for $i = 0$ by $1 - p_\theta$ and (A2) for $i = 1$ by p_θ, adding these two new equations together, and making use of (4) produces a single equation in the weighted expense E_c [by (6)]. The solution to this equation is (13).

Strategy (ii). If this strategy is selected, then (12) reduces to

$$E_c(i) = P_{i1}L + \alpha P_{i0}E_c(0), \quad \text{(A3)}$$

$i = 0, 1$. Because (A3) for $i = 0$ is an equation in only one unknown, $E_c(0)$, first this equation is solved. Next, this solution for $E_c(0)$ is substituted into (A3) for $i = 1$ to obtain a solution for $E_c(1)$ as well. Finally, these two solutions for the conditional expenses are inserted into (6) to produce a solution for the weighted expense E_c, which simplifies to (14) through use of (4).

Strategy (iii). If this strategy is selected, then (12) reduces to

$$E_c(0) = P_{01}L + \alpha P_{00}E_c(0),$$
$$E_c(1) = C + \alpha[P_{10}E_c(0) + P_{11}E_c(1)]. \quad \text{(A4)}$$

The exact same approach applied to obtain (14) from (A3) also works to obtain (15) from (A4).

3. Optimal policy

Strategy (i). It is sufficient to consider conditions for which protection is optimal given an initial (or current) weather state of $\theta_0 = 0$ [because this is the only situation in which strategy (i) differs from strategy (iii)]. From the dynamic programming recursion (12), protection is optimal provided

$$C + \alpha P_{01}E_c(1) < P_{01}L. \quad \text{(A5)}$$

Using the solution to (A2) for $E_c(1)$ and applying (4) gives the critical threshold C^* for the cost–loss ratio C/L

$$C^* = \frac{p_\theta(1-d)}{1 + \alpha^* p_\theta(1-d)}. \quad \text{(A6)}$$

Equating the expression (A6) for C^* with C/L and solving for d yields the critical value (10) for d^*.

Strategy (ii). It is sufficient to consider conditions for which protection is not optimal given $\theta_0 = 1$ [because this is the only situation in which strategy (ii) differs from strategy (iii)]. From (12), protection is not optimal provided

$$C + \alpha P_{11}E_c(1) > P_{11}L. \quad \text{(A7)}$$

Using the solution to (A3) for $E_c(1)$ and applying (4) gives the critical threshold C^{**} for C/L

$$C^{**} = \frac{(1-\alpha D)[1-(1-p_\theta)(1-d)]}{1+\alpha^* p_\theta(1-d)}. \quad \text{(A8)}$$

Equating the expression (A8) with C/L yields the quadratic equation (11) in d.

REFERENCES

Alchian, A. A., and W. R. Allen, 1972: *University Economics: Elements of Inquiry*, third ed. Wadsworth, 857 pp.

Brown, B. G., R. W. Katz, and A. H. Murphy, 1986: On the economic value of seasonal-precipitation forecasts: The fallowing/planting problem. *Bull. Amer. Meteor. Soc.*, **67**, 833–841.

Ehrendorfer, M., and A. H. Murphy, 1988: Comparative evaluation of weather forecasting systems: Sufficiency, quality, and accuracy. *Mon. Wea. Rev.*, **116**, 1757–1770.

Epstein, E. S., and A. H. Murphy, 1988: Use and value of multiple-period forecasts in a dynamic model of the cost–loss ratio situation. *Mon. Wea. Rev.*, **116**, 746–761.

Gabriel, K. R., and J. Neumann, 1962: A Markov chain model for daily rainfall occurrence at Tel Aviv. *Quart. J. Roy. Meteor. Soc.,* **88,** 90–95.

Hilton, R. W., 1981: The determination of information value: Synthesizing some general results. *Management Sci.,* **27,** 57–64.

Katz, R. W., 1988: Use of cross correlations in the search for teleconnections. *J. Climatol.,* **8,** 241–253.

——, and A. H. Murphy, 1987: Quality/value relationship for imperfect information in the umbrella problem. *Amer. Statist.,* **41,** 187–189.

——, and ——, 1990: Quality/value relationships for imperfect weather forecasts in a prototype multistage decision-making model. *J. Forecasting,* **9,** 75–86.

——, ——, and R. L. Winkler, 1982: Assessing the value of frost forecasts to orchardists: A dynamic decision-making approach. *J. Appl. Meteor.,* **21,** 518–531.

Krzysztofowicz, R., and D. Long, 1990: To protect or not to protect: Bayes decisions with forecasts. *Eur. J. Oper. Res.,* **44,** 319–330.

Mjelde, J. W., S. T. Sonka, B. L. Dixon, and P. J. Lamb, 1988: Valuing forecast characteristics in a dynamic agricultural production system. *Amer. J. Agric. Econ.,* **70,** 674–684.

Murphy, A. H., R. W. Katz, R. L. Winkler, and W.-R. Hsu, 1985: Repetitive decision making and the value of forecasts in the cost–loss ratio situation: A dynamic model. *Mon. Wea. Rev.,* **113,** 801–813.

Ross, S. M., 1970: *Applied Probability Models with Optimization Applications.* Holden-Day, 198 pp.

——, 1983: *Introduction to Stochastic Dynamic Programming.* Academic Press, 164 pp.

Wilks, D. S., 1991: Representing serial correlation of meteorological events and forecasts in dynamic decision-analytic models. *Mon. Wea. Rev.,* **119,** 1640–1662.

Winkler, R. L., and A. H. Murphy, 1985: Decision analysis. *Probability, Statistics, and Decision Making in the Atmospheric Sciences,* A. H. Murphy and R. W. Katz, Eds., Westview Press, 493–524.